GNAT Reference Manual

GNAT, The GNU Ada 95 Compiler
GNAT Version for GCC 3.3.6

Ada Core Technologies, Inc.

Copyright © 1995-2001, Free Software Foundation

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with the Invariant Sections being "GNU Free Documentation License", with the Front-Cover Texts being "GNAT Reference Manual", and with no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

About This Guide

This manual contains useful information in writing programs using the GNAT compiler. It includes information on implementation dependent characteristics of GNAT, including all the information required by Annex M of the standard.

Ada 95 is designed to be highly portable,and guarantees that, for most programs, Ada 95 compilers behave in exactly the same manner on different machines. However, since Ada 95 is designed to be used in a wide variety of applications, it also contains a number of system dependent features to Functbe used in interfacing to the external world.

Note: Any program that makes use of implementation-dependent features may be non-portable. You should follow good programming practice and isolate and clearly document any sections of your program that make use of these features in a non-portable manner.

What This Reference Manual Contains

This reference manual contains the following chapters:

- Chapter 1 [Implementation Defined Pragmas], page 3 lists GNAT implementation-dependent pragmas, which can be used to extend and enhance the functionality of the compiler.
- Chapter 2 [Implementation Defined Attributes], page 41 lists GNAT implementation-dependent attributes which can be used to extend and enhance the functionality of the compiler.
- Chapter 3 [Implementation Advice], page 51 provides information on generally desirable behavior which are not requirements that all compilers must follow since it cannot be provided on all systems, or which may be undesirable on some systems.
- Chapter 4 [Implementation Defined Characteristics], page 75 provides a guide to minimizing implementation dependent features.
- Chapter 5 [Intrinsic Subprograms], page 99 describes the intrinsic subprograms implemented by GNAT, and how they can be imported into user application programs.
- Chapter 6 [Representation Clauses and Pragmas], page 103 describes in detail the way that GNAT represents data, and in particular the exact set of representation clauses and pragmas that is accepted.
- Chapter 7 [Standard Library Routines], page 125 provides a listing of packages and a brief description of the functionality that is provided by Ada's extensive set of standard library routines as implemented by GNAT.
- Chapter 8 [The Implementation of Standard I/O], page 133 details how the GNAT implementation of the input-output facilities.
- Chapter 10 [Interfacing to Other Languages], page 159 describes how programs written in Ada using GNAT can be interfaced to other programming languages.
- Chapter 14 [Specialized Needs Annexes], page 167 describes the GNAT implementation of all of the special needs annexes.
- Chapter 15 [Compatibility Guide], page 169 includes sections on compatibility of GNAT with other Ada 83 and Ada 95 compilation systems, to assist in porting code from other environments.

This reference manual assumes that you are familiar with Ada 95 language, as described in the International Standard ANSI/ISO/IEC-8652:1995, Jan 1995.

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`', '`button` `names`', and '`field` `names`'.
- `Variables`.
- *Emphasis.*
- [optional information or parameters]
- Examples are described by text

  ```
  and then shown this way.
  ```

Commands that are entered by the user are preceded in this manual by the characters '`$` ' (dollar sign followed by space). If your system uses this sequence as a prompt, then the commands will appear exactly as you see them in the manual. If your system uses some other prompt, then the command will appear with the '`$`' replaced by whatever prompt character you are using.

Related Information

See the following documents for further information on GNAT:

- *GNAT User's Guide*, which provides information on how to use the GNAT compiler system.
- *Ada 95 Reference Manual*, which contains all reference material for the Ada 95 programming language.
- *Ada 95 Annotated Reference Manual*, which is an annotated version of the standard reference manual cited above. The annotations describe detailed aspects of the design decision, and in particular contain useful sections on Ada 83 compatibility.
- *DEC Ada, Technical Overview and Comparison on DIGITAL Platforms*, which contains specific information on compatibility between GNAT and DEC Ada 83 systems.
- *DEC Ada, Language Reference Manual, part number AA-PYZAB-TK* which describes in detail the pragmas and attributes provided by the DEC Ada 83 compiler system.

1 Implementation Defined Pragmas

Ada 95 defines a set of pragmas that can be used to supply additional information to the compiler. These language defined pragmas are implemented in GNAT and work as described in the Ada 95 Reference Manual.

In addition, Ada 95 allows implementations to define additional pragmas whose meaning is defined by the implementation. GNAT provides a number of these implementation-dependent pragmas which can be used to extend and enhance the functionality of the compiler. This section of the GNAT Reference Manual describes these additional pragmas.

Note that any program using these pragmas may not be portable to other compilers (although GNAT implements this set of pragmas on all platforms). Therefore if portability to other compilers is an important consideration, the use of these pragmas should be minimized.

`pragma Abort_Defer`

Syntax:

```
pragma Abort_Defer;
```

This pragma must appear at the start of the statement sequence of a handled sequence of statements (right after the `begin`). It has the effect of deferring aborts for the sequence of statements (but not for the declarations or handlers, if any, associated with this statement sequence).

`pragma Ada_83`

Syntax:

```
pragma Ada_83;
```

A configuration pragma that establishes Ada 83 mode for the unit to which it applies, regardless of the mode set by the command line switches. In Ada 83 mode, GNAT attempts to be as compatible with the syntax and semantics of Ada 83, as defined in the original Ada 83 Reference Manual as possible. In particular, the new Ada 95 keywords are not recognized, optional package bodies are allowed, and generics may name types with unknown discriminants without using the `(<>)` notation. In addition, some but not all of the additional restrictions of Ada 83 are enforced.

Ada 83 mode is intended for two purposes. Firstly, it allows existing legacy Ada 83 code to be compiled and adapted to GNAT with less effort. Secondly, it aids in keeping code backwards compatible with Ada 83. However, there is no guarantee that code that is processed correctly by GNAT in Ada 83 mode will in fact compile and execute with an Ada 83 compiler, since GNAT does not enforce all the additional checks required by Ada 83.

`pragma Ada_95`

Syntax:

```
pragma Ada_95;
```

A configuration pragma that establishes Ada 95 mode for the unit to which it applies, regardless of the mode set by the command line switches. This mode is set automatically for the `Ada` and `System` packages and their children, so you need not specify it in these contexts. This pragma is useful when writing a

reusable component that itself uses Ada 95 features, but which is intended to be usable from either Ada 83 or Ada 95 programs.

`pragma Annotate`

Syntax:

```
pragma Annotate (IDENTIFIER {, ARG});

ARG ::= NAME | EXPRESSION
```

This pragma is used to annotate programs. *identifier* identifies the type of annotation. GNAT verifies this is an identifier, but does not otherwise analyze it. The *arg* argument can be either a string literal or an expression. String literals are assumed to be of type `Standard.String`. Names of entities are simply analyzed as entity names. All other expressions are analyzed as expressions, and must be unambiguous.

The analyzed pragma is retained in the tree, but not otherwise processed by any part of the GNAT compiler. This pragma is intended for use by external tools, including ASIS.

`pragma Assert`

Syntax:

```
pragma Assert (
  boolean_EXPRESSION
  [, static_string_EXPRESSION])
```

The effect of this pragma depends on whether the corresponding command line switch is set to activate assertions. The pragma expands into code equivalent to the following:

```
if assertions-enabled then
   if not boolean_EXPRESSION then
      System.Assertions.Raise_Assert_Failure
        (string_EXPRESSION);
   end if;
end if;
```

The string argument, if given, is the message that will be associated with the exception occurrence if the exception is raised. If no second argument is given, the default message is '*file*`:`*nnn*', where *file* is the name of the source file containing the assert, and *nnn* is the line number of the assert. A pragma is not a statement, so if a statement sequence contains nothing but a pragma assert, then a null statement is required in addition, as in:

```
...
if J > 3 then
   pragma Assert (K > 3, "Bad value for K");
   null;
end if;
```

Note that, as with the `if` statement to which it is equivalent, the type of the expression is either `Standard.Boolean`, or any type derived from this standard type.

If assertions are disabled (switch `-gnata` not used), then there is no effect (and in particular, any side effects from the expression are suppressed). More precisely it is not quite true that the pragma has no effect, since the expression

is analyzed, and may cause types to be frozen if they are mentioned here for the first time.

If assertions are enabled, then the given expression is tested, and if it is `False` then `System.Assertions.Raise_Assert_Failure` is called which results in the raising of `Assert_Failure` with the given message.

If the boolean expression has side effects, these side effects will turn on and off with the setting of the assertions mode, resulting in assertions that have an effect on the program. You should generally avoid side effects in the expression arguments of this pragma. However, the expressions are analyzed for semantic correctness whether or not assertions are enabled, so turning assertions on and off cannot affect the legality of a program.

`pragma Ast_Entry`

Syntax:

```
pragma AST_Entry (entry_IDENTIFIER);
```

This pragma is implemented only in the OpenVMS implementation of GNAT. The argument is the simple name of a single entry; at most one `AST_Entry` pragma is allowed for any given entry. This pragma must be used in conjunction with the `AST_Entry` attribute, and is only allowed after the entry declaration and in the same task type specification or single task as the entry to which it applies. This pragma specifies that the given entry may be used to handle an OpenVMS asynchronous system trap (`AST`) resulting from an OpenVMS system service call. The pragma does not affect normal use of the entry. For further details on this pragma, see the DEC Ada Language Reference Manual, section 9.12a.

`pragma C_Pass_By_Copy`

Syntax:

```
pragma C_Pass_By_Copy
  ([Max_Size =>] static_integer_EXPRESSION);
```

Normally the default mechanism for passing C convention records to C convention subprograms is to pass them by reference, as suggested by RM B.3(69). Use the configuration pragma `C_Pass_By_Copy` to change this default, by requiring that record formal parameters be passed by copy if all of the following conditions are met:

- The size of the record type does not exceed *static_integer_expression*.
- The record type has `Convention C`.
- The formal parameter has this record type, and the subprogram has a foreign (non-Ada) convention.

If these conditions are met the argument is passed by copy, i.e. in a manner consistent with what C expects if the corresponding formal in the C prototype is a struct (rather than a pointer to a struct).

You can also pass records by copy by specifying the convention `C_Pass_By_Copy` for the record type, or by using the extended `Import` and `Export` pragmas, which allow specification of passing mechanisms on a parameter by parameter basis.

pragma Comment

Syntax:

```
pragma Comment (static_string_EXPRESSION);
```

This is almost identical in effect to pragma `Ident`. It allows the placement of a comment into the object file and hence into the executable file if the operating system permits such usage. The difference is that `Comment`, unlike `Ident`, has no limit on the length of the string argument, and no limitations on placement of the pragma (it can be placed anywhere in the main source unit).

pragma Common_Object

Syntax:

```
pragma Common_Object (
     [Internal =>] LOCAL_NAME,
  [, [External =>] EXTERNAL_SYMBOL]
  [, [Size     =>] EXTERNAL_SYMBOL] )

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION
```

This pragma enables the shared use of variables stored in overlaid linker areas corresponding to the use of `COMMON` in Fortran. The single object *local_name* is assigned to the area designated by the *External* argument. You may define a record to correspond to a series of fields. The *size* argument is syntax checked in GNAT, but otherwise ignored.

`Common_Object` is not supported on all platforms. If no support is available, then the code generator will issue a message indicating that the necessary attribute for implementation of this pragma is not available.

pragma Complex_Representation

Syntax:

```
pragma Complex_Representation
        ([Entity =>] LOCAL_NAME);
```

The *Entity* argument must be the name of a record type which has two fields of the same floating-point type. The effect of this pragma is to force gcc to use the special internal complex representation form for this record, which may be more efficient. Note that this may result in the code for this type not conforming to standard ABI (application binary interface) requirements for the handling of record types. For example, in some environments, there is a requirement for passing records by pointer, and the use of this pragma may result in passing this type in floating-point registers.

pragma Component_Alignment

Syntax:

```
pragma Component_Alignment (
     [Form =>] ALIGNMENT_CHOICE
  [, [Name =>] type_LOCAL_NAME]);

ALIGNMENT_CHOICE ::=
  Component_Size
| Component_Size_4
| Storage_Unit
```

```
| Default
```

Specifies the alignment of components in array or record types. The meaning of the *Form* argument is as follows:

`Component_Size`
: Aligns scalar components and subcomponents of the array or record type on boundaries appropriate to their inherent size (naturally aligned). For example, 1-byte components are aligned on byte boundaries, 2-byte integer components are aligned on 2-byte boundaries, 4-byte integer components are aligned on 4-byte boundaries and so on. These alignment rules correspond to the normal rules for C compilers on all machines except the VAX.

`Component_Size_4`
: Naturally aligns components with a size of four or fewer bytes. Components that are larger than 4 bytes are placed on the next 4-byte boundary.

`Storage_Unit`
: Specifies that array or record components are byte aligned, i.e. aligned on boundaries determined by the value of the constant `System.Storage_Unit`.

`Default`
: Specifies that array or record components are aligned on default boundaries, appropriate to the underlying hardware or operating system or both. For OpenVMS VAX systems, the `Default` choice is the same as the `Storage_Unit` choice (byte alignment). For all other systems, the `Default` choice is the same as `Component_Size` (natural alignment).

If the `Name` parameter is present, *type_local_name* must refer to a local record or array type, and the specified alignment choice applies to the specified type. The use of `Component_Alignment` together with a pragma `Pack` causes the `Component_Alignment` pragma to be ignored. The use of `Component_Alignment` together with a record representation clause is only effective for fields not specified by the representation clause.

If the `Name` parameter is absent, the pragma can be used as either a configuration pragma, in which case it applies to one or more units in accordance with the normal rules for configuration pragmas, or it can be used within a declarative part, in which case it applies to types that are declared within this declarative part, or within any nested scope within this declarative part. In either case it specifies the alignment to be applied to any record or array type which has otherwise standard representation.

If the alignment for a record or array type is not specified (using pragma `Pack`, pragma `Component_Alignment`, or a record rep clause), the GNAT uses the default alignment as described previously.

`pragma Convention_Identifier`
: Syntax:

```
pragma Convention_Identifier (
         [Name =>]       IDENTIFIER,
         [Convention =>] convention_IDENTIFIER);
```

This pragma provides a mechanism for supplying synonyms for existing convention identifiers. The `Name` identifier can subsequently be used as a synonym for the given convention in other pragmas (including for example pragma `Import` or another `Convention_Identifier` pragma). As an example of the use of this, suppose you had legacy code which used Fortran77 as the identifier for Fortran. Then the pragma:

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

would allow the use of the convention identifier `Fortran77` in subsequent code, avoiding the need to modify the sources. As another example, you could use this to parametrize convention requirements according to systems. Suppose you needed to use `Stdcall` on windows systems, and `C` on some other system, then you could define a convention identifier `Library` and use a single `Convention_Identifier` pragma to specify which convention would be used system-wide.

pragma CPP_Class

Syntax:

```
pragma CPP_Class ([Entity =>] LOCAL_NAME);
```

The argument denotes an entity in the current declarative region that is declared as a tagged or untagged record type. It indicates that the type corresponds to an externally declared C++ class type, and is to be laid out the same way that C++ would lay out the type.

If (and only if) the type is tagged, at least one component in the record must be of type `Interfaces.CPP.Vtable_Ptr`, corresponding to the C++ Vtable (or Vtables in the case of multiple inheritance) used for dispatching.

Types for which `CPP_Class` is specified do not have assignment or equality operators defined (such operations can be imported or declared as subprograms as required). Initialization is allowed only by constructor functions (see pragma `CPP_Constructor`).

Pragma `CPP_Class` is intended primarily for automatic generation using an automatic binding generator tool. See Section 10.2 [Interfacing to C++], page 160 for related information.

pragma CPP_Constructor

Syntax:

```
pragma CPP_Constructor ([Entity =>] LOCAL_NAME);
```

This pragma identifies an imported function (imported in the usual way with pragma `Import`) as corresponding to a C++ constructor. The argument is a name that must have been previously mentioned in a pragma `Import` with `Convention` = `CPP`, and must be of one of the following forms:

- `function` *Fname* `return` *T*`'Class`
- `function` *Fname* `(...) return` *T*`'Class`

where *T* is a tagged type to which the pragma `CPP_Class` applies.

The first form is the default constructor, used when an object of type T is created on the Ada side with no explicit constructor. Other constructors (including the copy constructor, which is simply a special case of the second form in which the one and only argument is of type T), can only appear in two contexts:

- On the right side of an initialization of an object of type T.
- In an extension aggregate for an object of a type derived from T.

Although the constructor is described as a function that returns a value on the Ada side, it is typically a procedure with an extra implicit argument (the object being initialized) at the implementation level. GNAT issues the appropriate call, whatever it is, to get the object properly initialized.

In the case of derived objects, you may use one of two possible forms for declaring and creating an object:

- `New_Object : Derived_T`
- `New_Object : Derived_T := (`*`constructor-function-call with ...`*`)`

In the first case the default constructor is called and extension fields if any are initialized according to the default initialization expressions in the Ada declaration. In the second case, the given constructor is called and the extension aggregate indicates the explicit values of the extension fields.

If no constructors are imported, it is impossible to create any objects on the Ada side. If no default constructor is imported, only the initialization forms using an explicit call to a constructor are permitted.

Pragma `CPP_Constructor` is intended primarily for automatic generation using an automatic binding generator tool. See Section 10.2 [Interfacing to C++], page 160 for more related information.

`pragma CPP_Virtual`

Syntax:

```
pragma CPP_Virtual
     [Entity     =>] ENTITY,
  [, [Vtable_Ptr =>] vtable_ENTITY,]
  [, [Position   =>] static_integer_EXPRESSION])
```

This pragma serves the same function as pragma `Import` in that case of a virtual function imported from C++. The *Entity* argument must be a primitive subprogram of a tagged type to which pragma `CPP_Class` applies. The *Vtable_Ptr* argument specifies the Vtable_Ptr component which contains the entry for this virtual function. The *Position* argument is the sequential number counting virtual functions for this Vtable starting at 1.

The `Vtable_Ptr` and `Position` arguments may be omitted if there is one Vtable_Ptr present (single inheritance case) and all virtual functions are imported. In that case the compiler can deduce both these values.

No `External_Name` or `Link_Name` arguments are required for a virtual function, since it is always accessed indirectly via the appropriate Vtable entry.

Pragma `CPP_Virtual` is intended primarily for automatic generation using an automatic binding generator tool. See Section 10.2 [Interfacing to C++], page 160 for related information.

`pragma CPP_Vtable`

Syntax:

```
pragma CPP_Vtable (
  [Entity      =>] ENTITY,
  [Vtable_Ptr  =>] vtable_ENTITY,
  [Entry_Count =>] static_integer_EXPRESSION);
```

Given a record to which the pragma `CPP_Class` applies, this pragma can be specified for each component of type `CPP.Interfaces.Vtable_Ptr`. *Entity* is the tagged type, *Vtable_Ptr* is the record field of type `Vtable_Ptr`, and *Entry_Count* is the number of virtual functions on the C++ side. Not all of these functions need to be imported on the Ada side.

You may omit the `CPP_Vtable` pragma if there is only one `Vtable_Ptr` component in the record and all virtual functions are imported on the Ada side (the default value for the entry count in this case is simply the total number of virtual functions).

Pragma `CPP_Vtable` is intended primarily for automatic generation using an automatic binding generator tool. See Section 10.2 [Interfacing to C++], page 160 for related information.

`pragma Debug`

Syntax:

```
pragma Debug (PROCEDURE_CALL_WITHOUT_SEMICOLON);

PROCEDURE_CALL_WITHOUT_SEMICOLON ::=
  PROCEDURE_NAME
| PROCEDURE_PREFIX ACTUAL_PARAMETER_PART
```

The argument has the syntactic form of an expression, meeting the syntactic requirements for pragmas.

If assertions are not enabled on the command line, this pragma has no effect. If asserts are enabled, the semantics of the pragma is exactly equivalent to the procedure call statement corresponding to the argument with a terminating semicolon. Pragmas are permitted in sequences of declarations, so you can use pragma `Debug` to intersperse calls to debug procedures in the middle of declarations.

`pragma Elaboration_Checks`

Syntax:

```
pragma Elaboration_Checks (RM | Static);
```

This is a configuration pragma that provides control over the elaboration model used by the compilation affected by the pragma. If the parameter is RM, then the dynamic elaboration model described in the Ada Reference Manual is used, as though the `-gnatE` switch had been specified on the command line. If the parameter is Static, then the default GNAT static model is used. This configuration pragma overrides the setting of the command line. For full details on the elaboration models used by the GNAT compiler, see section “Elaboration Order Handling in GNAT” in the *GNAT User's Guide*.

`pragma Eliminate`

Syntax:

```
pragma Eliminate (
    [Unit_Name =>] IDENTIFIER |
                   SELECTED_COMPONENT);

pragma Eliminate (
    [Unit_Name       =>]  IDENTIFIER |
                          SELECTED_COMPONENT,
    [Entity          =>]  IDENTIFIER |
                          SELECTED_COMPONENT |
                          STRING_LITERAL
  [,[Parameter_Types =>]  PARAMETER_TYPES]
  [,[Result_Type     =>]  result_SUBTYPE_NAME]
  [,[Homonym_Number  =>]  INTEGER_LITERAL]);

PARAMETER_TYPES ::= (SUBTYPE_NAME {, SUBTYPE_NAME})
SUBTYPE_NAME    ::= STRING_LITERAL
```

This pragma indicates that the given entity is not used outside the compilation unit it is defined in. The entity may be either a subprogram or a variable.

If the entity to be eliminated is a library level subprogram, then the first form of pragma `Eliminate` is used with only a single argument. In this form, the `Unit_Name` argument specifies the name of the library level unit to be eliminated.

In all other cases, both `Unit_Name` and `Entity` arguments are required. item is an entity of a library package, then the first argument specifies the unit name, and the second argument specifies the particular entity. If the second argument is in string form, it must correspond to the internal manner in which GNAT stores entity names (see compilation unit Namet in the compiler sources for details).

The remaining parameters are optionally used to distinguish between overloaded subprograms. There are two ways of doing this.

Use `Parameter_Types` and `Result_Type` to specify the profile of the subprogram to be eliminated in a manner similar to that used for the extended `Import` and `Export` pragmas, except that the subtype names are always given as string literals, again corresponding to the internal manner in which GNAT stores entity names.

Alternatively, the `Homonym_Number` parameter is used to specify which overloaded alternative is to be eliminated. A value of 1 indicates the first subprogram (in lexical order), 2 indicates the second etc.

The effect of the pragma is to allow the compiler to eliminate the code or data associated with the named entity. Any reference to an eliminated entity outside the compilation unit it is defined in, causes a compile time or link time error.

The parameters of this pragma may be given in any order, as long as the usual rules for use of named parameters and position parameters are used.

The intention of pragma `Eliminate` is to allow a program to be compiled in a system independent manner, with unused entities eliminated, without the requirement of modifying the source text. Normally the required set of `Eliminate` pragmas is constructed automatically using the gnatelim tool. Elimination of unused entities local to a compilation unit is automatic, without requiring the use of pragma `Eliminate`.

Note that the reason this pragma takes string literals where names might be expected is that a pragma `Eliminate` can appear in a context where the relevant names are not visible.

`pragma Export_Exception`

Syntax:

```
pragma Export_Exception (
     [Internal =>] LOCAL_NAME,
  [, [External =>] EXTERNAL_SYMBOL,]
  [, [Form     =>] Ada | VMS]
  [, [Code     =>] static_integer_EXPRESSION]);

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION
```

This pragma is implemented only in the OpenVMS implementation of GNAT. It causes the specified exception to be propagated outside of the Ada program, so that it can be handled by programs written in other OpenVMS languages. This pragma establishes an external name for an Ada exception and makes the name available to the OpenVMS Linker as a global symbol. For further details on this pragma, see the DEC Ada Language Reference Manual, section 13.9a3.2.

`pragma Export_Function ...`

Syntax:

```
pragma Export_Function (
     [Internal         =>] LOCAL_NAME,
  [, [External         =>] EXTERNAL_SYMBOL]
  [, [Parameter_Types  =>] PARAMETER_TYPES]
  [, [Result_Type      =>] result_SUBTYPE_MARK]
  [, [Mechanism        =>] MECHANISM]
  [, [Result_Mechanism =>] MECHANISM_NAME]);

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION

PARAMETER_TYPES ::=
  null
| SUBTYPE_MARK {, SUBTYPE_MARK}

MECHANISM ::=
  MECHANISM_NAME
| (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
  [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::=
  Value
| Reference
| Descriptor [([Class =>] CLASS_NAME)]

CLASS_NAME ::= ubs | ubsb | uba | s | sb | a | nca
```

Use this pragma to make a function externally callable and optionally provide information on mechanisms to be used for passing parameter and result values. We recommend, for the purposes of improving portability, this pragma always be used in conjunction with a separate pragma `Export`, which must precede the pragma `Export_Function`. GNAT does not require a separate pragma `Export`, but if none is present, `Convention Ada` is assumed, which is usually not what is wanted, so it is usually appropriate to use this pragma in conjunction with a `Export` or `Convention` pragma that specifies the desired foreign convention. Pragma `Export_Function` (and `Export`, if present) must appear in the same declarative region as the function to which they apply.

internal_name must uniquely designate the function to which the pragma applies. If more than one function name exists of this name in the declarative part you must use the `Parameter_Types` and `Result_Type` parameters is mandatory to achieve the required unique designation. *subtype_ marks* in these parameters must exactly match the subtypes in the corresponding function specification, using positional notation to match parameters with subtype marks. Passing by descriptor is supported only on the OpenVMS ports of GNAT.

`pragma Export_Object ...`

Syntax:

```
pragma Export_Object
      [Internal =>] LOCAL_NAME,
   [, [External =>] EXTERNAL_SYMBOL]
   [, [Size     =>] EXTERNAL_SYMBOL]

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION
```

This pragma designates an object as exported, and apart from the extended rules for external symbols, is identical in effect to the use of the normal `Export` pragma applied to an object. You may use a separate Export pragma (and you probably should from the point of view of portability), but it is not required. *Size* is syntax checked, but otherwise ignored by GNAT.

`pragma Export_Procedure ...`

Syntax:

```
pragma Export_Procedure (
     [Internal        =>] LOCAL_NAME
  [, [External        =>] EXTERNAL_SYMBOL]
  [, [Parameter_Types =>] PARAMETER_TYPES]
  [, [Mechanism       =>] MECHANISM]);

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION

PARAMETER_TYPES ::=
  null
| SUBTYPE_MARK {, SUBTYPE_MARK}

MECHANISM ::=
  MECHANISM_NAME
```

```
| (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
  [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::=
  Value
| Reference
| Descriptor [([Class =>] CLASS_NAME)]

CLASS_NAME ::= ubs | ubsb | uba | s | sb | a | nca
```

This pragma is identical to `Export_Function` except that it applies to a procedure rather than a function and the parameters `Result_Type` and `Result_Mechanism` are not permitted. GNAT does not require a separate pragma `Export`, but if none is present, `Convention Ada` is assumed, which is usually not what is wanted, so it is usually appropriate to use this pragma in conjunction with a `Export` or `Convention` pragma that specifies the desired foreign convention.

`pragma Export_Valued_Procedure`

Syntax:

```
pragma Export_Valued_Procedure (
     [Internal        =>] LOCAL_NAME
  [, [External        =>] EXTERNAL_SYMBOL]
  [, [Parameter_Types =>] PARAMETER_TYPES]
  [, [Mechanism       =>] MECHANISM]);

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION

PARAMETER_TYPES ::=
  null
| SUBTYPE_MARK {, SUBTYPE_MARK}

MECHANISM ::=
  MECHANISM_NAME
| (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
  [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::=
  Value
| Reference
| Descriptor [([Class =>] CLASS_NAME)]

CLASS_NAME ::= ubs | ubsb | uba | s | sb | a | nca
```

This pragma is identical to `Export_Procedure` except that the first parameter of *local_name*, which must be present, must be of mode `OUT`, and externally the subprogram is treated as a function with this parameter as the result of the function. GNAT provides for this capability to allow the use of `OUT` and `IN OUT` parameters in interfacing to external functions (which are not permitted in Ada functions). GNAT does not require a separate pragma `Export`, but if none

is present, `Convention Ada` is assumed, which is almost certainly not what is wanted since the whole point of this pragma is to interface with foreign language functions, so it is usually appropriate to use this pragma in conjunction with a `Export` or `Convention` pragma that specifies the desired foreign convention.

`pragma Extend_System`

Syntax:

```
pragma Extend_System ([Name =>] IDENTIFIER);
```

This pragma is used to provide backwards compatibility with other implementations that extend the facilities of package `System`. In GNAT, `System` contains only the definitions that are present in the Ada 95 RM. However, other implementations, notably the DEC Ada 83 implementation, provide many extensions to package `System`.

For each such implementation accommodated by this pragma, GNAT provides a package `Aux_xxx`, e.g. `Aux_DEC` for the DEC Ada 83 implementation, which provides the required additional definitions. You can use this package in two ways. You can `with` it in the normal way and access entities either by selection or using a `use` clause. In this case no special processing is required.

However, if existing code contains references such as `System.xxx` where *xxx* is an entity in the extended definitions provided in package `System`, you may use this pragma to extend visibility in `System` in a non-standard way that provides greater compatibility with the existing code. Pragma `Extend_System` is a configuration pragma whose single argument is the name of the package containing the extended definition (e.g. `Aux_DEC` for the DEC Ada case). A unit compiled under control of this pragma will be processed using special visibility processing that looks in package `System.Aux_xxx` where `Aux_xxx` is the pragma argument for any entity referenced in package `System`, but not found in package `System`.

You can use this pragma either to access a predefined `System` extension supplied with the compiler, for example `Aux_DEC` or you can construct your own extension unit following the above definition. Note that such a package is a child of `System` and thus is considered part of the implementation. To compile it you will have to use the appropriate switch for compiling system units. See the GNAT User's Guide for details.

`pragma External`

Syntax:

```
pragma External (
  [   Convention    =>] convention_IDENTIFIER,
  [   Entity        =>] local_NAME
  [, [External_Name =>] static_string_EXPRESSION ]
  [, [Link_Name     =>] static_string_EXPRESSION ]);
```

This pragma is identical in syntax and semantics to pragma `Export` as defined in the Ada Reference Manual. It is provided for compatibility with some Ada 83 compilers that used this pragma for exactly the same purposes as pragma `Export` before the latter was standardized.

`pragma External_Name_Casing`

Syntax:

```
pragma External_Name_Casing (
  Uppercase | Lowercase
  [, Uppercase | Lowercase | As_Is]);
```

This pragma provides control over the casing of external names associated with Import and Export pragmas. There are two cases to consider:

Implicit external names

Implicit external names are derived from identifiers. The most common case arises when a standard Ada 95 Import or Export pragma is used with only two arguments, as in:

```
pragma Import (C, C_Routine);
```

Since Ada is a case insensitive language, the spelling of the identifier in the Ada source program does not provide any information on the desired casing of the external name, and so a convention is needed. In GNAT the default treatment is that such names are converted to all lower case letters. This corresponds to the normal C style in many environments. The first argument of pragma `External_Name_Casing` can be used to control this treatment. If `Uppercase` is specified, then the name will be forced to all uppercase letters. If `Lowercase` is specified, then the normal default of all lower case letters will be used.

This same implicit treatment is also used in the case of extended DEC Ada 83 compatible Import and Export pragmas where an external name is explicitly specified using an identifier rather than a string.

Explicit external names

Explicit external names are given as string literals. The most common case arises when a standard Ada 95 Import or Export pragma is used with three arguments, as in:

```
pragma Import (C, C_Routine, "C_routine");
```

In this case, the string literal normally provides the exact casing required for the external name. The second argument of pragma `External_Name_Casing` may be used to modify this behavior. If `Uppercase` is specified, then the name will be forced to all uppercase letters. If `Lowercase` is specified, then the name will be forced to all lowercase letters. A specification of `As_Is` provides the normal default behavior in which the casing is taken from the string provided.

This pragma may appear anywhere that a pragma is valid. In particular, it can be used as a configuration pragma in the '`gnat.adc`' file, in which case it applies to all subsequent compilations, or it can be used as a program unit pragma, in which case it only applies to the current unit, or it can be used more locally to control individual Import/Export pragmas.

It is primarily intended for use with OpenVMS systems, where many compilers convert all symbols to upper case by default. For interfacing to such compilers (e.g. the DEC C compiler), it may be convenient to use the pragma:

```
pragma External_Name_Casing (Uppercase, Uppercase);
```

to enforce the upper casing of all external symbols.

`pragma Finalize_Storage_Only`

Syntax:

```
pragma Finalize_Storage_Only (first_subtype_LOCAL_NAME);
```

This pragma allows the compiler not to emit a Finalize call for objects defined at the library level. This is mostly useful for types where finalization is only used to deal with storage reclamation since in most environments it is not necessary to reclaim memory just before terminating execution, hence the name.

`pragma Float_Representation`

Syntax:

```
pragma Float_Representation (FLOAT_REP);

FLOAT_REP ::= VAX_Float | IEEE_Float
```

This pragma is implemented only in the OpenVMS implementation of GNAT. It allows control over the internal representation chosen for the predefined floating point types declared in the packages `Standard` and `System`. For further details on this pragma, see the DEC Ada Language Reference Manual, section 3.5.7a. Note that to use this pragma, the standard runtime libraries must be recompiled. See the description of the `GNAT LIBRARY` command in the OpenVMS version of the GNAT Users Guide for details on the use of this command.

`pragma Ident`

Syntax:

```
pragma Ident (static_string_EXPRESSION);
```

This pragma provides a string identification in the generated object file, if the system supports the concept of this kind of identification string. The maximum permitted length of the string literal is 31 characters. This pragma is allowed only in the outermost declarative part or declarative items of a compilation unit. On OpenVMS systems, the effect of the pragma is identical to the effect of the DEC Ada 83 pragma of the same name.

`pragma Import_Exception`

Syntax:

```
pragma Import_Exception (
     [Internal =>] LOCAL_NAME,
  [, [External =>] EXTERNAL_SYMBOL,]
  [, [Form     =>] Ada | VMS]
  [, [Code     =>] static_integer_EXPRESSION]);

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION
```

This pragma is implemented only in the OpenVMS implementation of GNAT. It allows OpenVMS conditions (for example, from OpenVMS system services or other OpenVMS languages) to be propagated to Ada programs as Ada exceptions. The pragma specifies that the exception associated with an exception declaration in an Ada program be defined externally (in non-Ada code). For

further details on this pragma, see the DEC Ada Language Reference Manual, section 13.9a.3.1.

`pragma Import_Function ...`

Syntax:

```
pragma Import_Function (
     [Internal                 =>] LOCAL_NAME,
  [, [External                 =>] EXTERNAL_SYMBOL]
  [, [Parameter_Types          =>] PARAMETER_TYPES]
  [, [Result_Type              =>] SUBTYPE_MARK]
  [, [Mechanism                =>] MECHANISM]
  [, [Result_Mechanism         =>] MECHANISM_NAME]
  [, [First_Optional_Parameter =>] IDENTIFIER]);

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION

PARAMETER_TYPES ::=
  null
| SUBTYPE_MARK {, SUBTYPE_MARK}

MECHANISM ::=
  MECHANISM_NAME
| (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
  [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::=
  Value
| Reference
| Descriptor [([Class =>] CLASS_NAME)]

CLASS_NAME ::= ubs | ubsb | uba | s | sb | a | nca
```

This pragma is used in conjunction with a pragma `Import` to specify additional information for an imported function. The pragma `Import` (or equivalent pragma `Interface`) must precede the `Import_Function` pragma and both must appear in the same declarative part as the function specification.

The *Internal_Name* argument must uniquely designate the function to which the pragma applies. If more than one function name exists of this name in the declarative part you must use the `Parameter_Types` and *Result_Type* parameters to achieve the required unique designation. Subtype marks in these parameters must exactly match the subtypes in the corresponding function specification, using positional notation to match parameters with subtype marks.

You may optionally use the *Mechanism* and *Result_Mechanism* parameters to specify passing mechanisms for the parameters and result. If you specify a single mechanism name, it applies to all parameters. Otherwise you may specify a mechanism on a parameter by parameter basis using either positional or named notation. If the mechanism is not specified, the default mechanism is used.

Passing by descriptor is supported only on the to OpenVMS ports of GNAT.

`First_Optional_Parameter` applies only to OpenVMS ports of GNAT. It specifies that the designated parameter and all following parameters are optional, meaning that they are not passed at the generated code level (this is distinct from the notion of optional parameters in Ada where the parameters are passed anyway with the designated optional parameters). All optional parameters must be of mode `IN` and have default parameter values that are either known at compile time expressions, or uses of the `'Null_Parameter` attribute.

`pragma Import_Object`

Syntax:

```
pragma Import_Object
     [Internal =>] LOCAL_NAME,
  [, [External =>] EXTERNAL_SYMBOL],
  [, [Size     =>] EXTERNAL_SYMBOL])

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION
```

This pragma designates an object as imported, and apart from the extended rules for external symbols, is identical in effect to the use of the normal `Import` pragma applied to an object. Unlike the subprogram case, you need not use a separate `Import` pragma, although you may do so (and probably should do so from a portability point of view). *size* is syntax checked, but otherwise ignored by GNAT.

`pragma Import_Procedure`

Syntax:

```
pragma Import_Procedure (
     [Internal                 =>] LOCAL_NAME,
  [, [External                 =>] EXTERNAL_SYMBOL]
  [, [Parameter_Types          =>] PARAMETER_TYPES]
  [, [Mechanism                =>] MECHANISM]
  [, [First_Optional_Parameter =>] IDENTIFIER]);

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION

PARAMETER_TYPES ::=
  null
| SUBTYPE_MARK {, SUBTYPE_MARK}

MECHANISM ::=
  MECHANISM_NAME
| (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
  [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::=
  Value
| Reference
| Descriptor [([Class =>] CLASS_NAME)]

CLASS_NAME ::= ubs | ubsb | uba | s | sb | a | nca
```

This pragma is identical to `Import_Function` except that it applies to a procedure rather than a function and the parameters `Result_Type` and `Result_Mechanism` are not permitted.

`pragma Import_Valued_Procedure ...`

Syntax:

```
pragma Import_Valued_Procedure (
     [Internal                 =>] LOCAL_NAME,
  [, [External                 =>] EXTERNAL_SYMBOL]
  [, [Parameter_Types          =>] PARAMETER_TYPES]
  [, [Mechanism                =>] MECHANISM]
  [, [First_Optional_Parameter =>] IDENTIFIER]);

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION

PARAMETER_TYPES ::=
  null
| SUBTYPE_MARK {, SUBTYPE_MARK}

MECHANISM ::=
  MECHANISM_NAME
| (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
  [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::=
  Value
| Reference
| Descriptor [([Class =>] CLASS_NAME)]

CLASS_NAME ::= ubs | ubsb | uba | s | sb | a | nca
```

This pragma is identical to `Import_Procedure` except that the first parameter of *local_name*, which must be present, must be of mode `OUT`, and externally the subprogram is treated as a function with this parameter as the result of the function. The purpose of this capability is to allow the use of `OUT` and `IN OUT` parameters in interfacing to external functions (which are not permitted in Ada functions). You may optionally use the `Mechanism` parameters to specify passing mechanisms for the parameters. If you specify a single mechanism name, it applies to all parameters. Otherwise you may specify a mechanism on a parameter by parameter basis using either positional or named notation. If the mechanism is not specified, the default mechanism is used.

Note that it is important to use this pragma in conjunction with a separate pragma Import that specifies the desired convention, since otherwise the default convention is Ada, which is almost certainly not what is required.

`pragma Initialize_Scalars`

Syntax:

```
pragma Initialize_Scalars;
```

This pragma is similar to `Normalize_Scalars` conceptually but has two important differences. First, there is no requirement for the pragma to be used

uniformly in all units of a partition, in particular, it is fine to use this just for some or all of the application units of a partition, without needing to recompile the run-time library.

In the case where some units are compiled with the pragma, and some without, then a declaration of a variable where the type is defined in package Standard or is locally declared will always be subject to initialization, as will any declaration of a scalar variable. For composite variables, whether the variable is initialized may also depend on whether the package in which the type of the variable is declared is compiled with the pragma.

The other important difference is that there is control over the value used for initializing scalar objects. At bind time, you can select whether to initialize with invalid values (like Normalize_Scalars), or with high or low values, or with a specified bit pattern. See the users guide for binder options for specifying these cases.

This means that you can compile a program, and then without having to recompile the program, you can run it with different values being used for initializing otherwise uninitialized values, to test if your program behavior depends on the choice. Of course the behavior should not change, and if it does, then most likely you have an erroneous reference to an uninitialized value.

Note that pragma `Initialize_Scalars` is particularly useful in conjunction with the enhanced validity checking that is now provided in GNAT, which checks for invalid values under more conditions. Using this feature (see description of the `-gnatv` flag in the users guide) in conjunction with pragma `Initialize_Scalars` provides a powerful new tool to assist in the detection of problems caused by uninitialized variables.

`pragma Inline_Always`

Syntax:

```
pragma Inline_Always (NAME [, NAME]);
```

Similar to pragma `Inline` except that inlining is not subject to the use of option `-gnatn` for inter-unit inlining.

`pragma Inline_Generic`

Syntax:

```
pragma Inline_Generic (generic_package_NAME)
```

This is implemented for compatibility with DEC Ada 83 and is recognized, but otherwise ignored, by GNAT. All generic instantiations are inlined by default when using GNAT.

`pragma Interface`

Syntax:

```
pragma Interface (
     [Convention    =>] convention_identifier,
     [Entity =>] local_name
  [, [External_Name =>] static_string_expression],
  [, [Link_Name     =>] static_string_expression]);
```

This pragma is identical in syntax and semantics to the standard Ada 95 pragma `Import`. It is provided for compatibility with Ada 83. The definition is upwards

compatible both with pragma `Interface` as defined in the Ada 83 Reference Manual, and also with some extended implementations of this pragma in certain Ada 83 implementations.

`pragma Interface_Name`

Syntax:

```
pragma Interface_Name (
     [Entity        =>] LOCAL_NAME
   [, [External_Name =>] static_string_EXPRESSION]
   [, [Link_Name     =>] static_string_EXPRESSION]);
```

This pragma provides an alternative way of specifying the interface name for an interfaced subprogram, and is provided for compatibility with Ada 83 compilers that use the pragma for this purpose. You must provide at least one of *External_Name* or *Link_Name*.

`pragma License`

Syntax:

```
pragma License (Unrestricted | GPL | Modified_GPL | Restricted);
```

This pragma is provided to allow automated checking for appropriate license conditions with respect to the standard and modified GPL. A pragma `License`, which is a configuration pragma that typically appears at the start of a source file or in a separate '`gnat.adc`' file, specifies the licensing conditions of a unit as follows:

- Unrestricted This is used for a unit that can be freely used with no license restrictions. Examples of such units are public domain units, and units from the Ada Reference Manual.
- GPL This is used for a unit that is licensed under the unmodified GPL, and which therefore cannot be `with`'ed by a restricted unit.
- Modified_GPL This is used for a unit licensed under the GNAT modified GPL that includes a special exception paragraph that specifically permits the inclusion of the unit in programs without requiring the entire program to be released under the GPL. This is the license used for the GNAT run-time which ensures that the run-time can be used freely in any program without GPL concerns.
- Restricted This is used for a unit that is restricted in that it is not permitted to depend on units that are licensed under the GPL. Typical examples are proprietary code that is to be released under more restrictive license conditions. Note that restricted units are permitted to `with` units which are licensed under the modified GPL (this is the whole point of the modified GPL).

Normally a unit with no `License` pragma is considered to have an unknown license, and no checking is done. However, standard GNAT headers are recognized, and license information is derived from them as follows.

A GNAT license header starts with a line containing 78 hyphens. The following comment text is searched for the appearence of any of the following strings.

If the string "GNU General Public License" is found, then the unit is assumed to have GPL license, unless the string "As a special exception" follows, in which case the license is assumed to be modified GPL.

If one of the strings "This specification is adapated from the Ada Semantic Interface" or "This specification is derived from the Ada Reference Manual" is found then the unit is assumed to be unrestricted.

These default actions means that a program with a restricted license pragma will automatically get warnings if a GPL unit is inappropriately `with`'ed. For example, the program:

```
with Sem_Ch3;
with GNAT.Sockets;
procedure Secret_Stuff is
...
end Secret_Stuff
```

if compiled with pragma `License` (`Restricted`) in a '`gnat.adc`' file will generate the warning:

```
1.  with Sem_Ch3;
           |
   >>> license of withed unit "Sem_Ch3" is incompatible

2.  with GNAT.Sockets;
3.  procedure Secret_Stuff is
```

Here we get a warning on `Sem_Ch3` since it is part of the GNAT compiler and is licensed under the GPL, but no warning for `GNAT.Sockets` which is part of the GNAT run time, and is therefore licensed under the modified GPL.

`pragma Link_With`

Syntax:

```
pragma Link_With (static_string_EXPRESSION {,static_string_EXPRESSION});
```

This pragma is provided for compatibility with certain Ada 83 compilers. It has exactly the same effect as pragma `Linker_Options` except that spaces occurring within one of the string expressions are treated as separators. For example, in the following case:

```
pragma Link_With ("-labc -ldef");
```

results in passing the strings `-labc` and `-ldef` as two separate arguments to the linker. In addition pragma Link_With allows multiple arguments, with the same effect as successive pragmas.

`pragma Linker_Alias`

Syntax:

```
pragma Linker_Alias (
  [Entity =>] LOCAL_NAME
  [Alias  =>] static_string_EXPRESSION);
```

This pragma establishes a linker alias for the given named entity. For further details on the exact effect, consult the GCC manual.

`pragma Linker_Section`

Syntax:

```
pragma Linker_Section (
  [Entity  =>] LOCAL_NAME
  [Section =>] static_string_EXPRESSION);
```

This pragma specifies the name of the linker section for the given entity. For further details on the exact effect, consult the GCC manual.

`pragma No_Run_Time`

Syntax:

```
pragma No_Run_Time;
```

This is a configuration pragma that makes sure the user code does not use nor need anything from the GNAT run time. This is mostly useful in context where code certification is required. Please consult the *GNAT Pro High-Integrity Edition User's Guide* for additional information.

`pragma Normalize_Scalars`

Syntax:

```
pragma Normalize_Scalars;
```

This is a language defined pragma which is fully implemented in GNAT. The effect is to cause all scalar objects that are not otherwise initialized to be initialized. The initial values are implementation dependent and are as follows:

`Standard.Character`
: Objects whose root type is Standard.Character are initialized to Character'Last. This will be out of range of the subtype only if the subtype range excludes this value.

`Standard.Wide_Character`
: Objects whose root type is Standard.Wide_Character are initialized to Wide_Character'Last. This will be out of range of the subtype only if the subtype range excludes this value.

`Integer types`
: Objects of an integer type are initialized to base_type'First, where base_type is the base type of the object type. This will be out of range of the subtype only if the subtype range excludes this value. For example, if you declare the subtype:

  ```
  subtype Ityp is integer range 1 .. 10;
  ```

  then objects of type x will be initialized to Integer'First, a negative number that is certainly outside the range of subtype `Ityp`.

`Real types`
: Objects of all real types (fixed and floating) are initialized to base_type'First, where base_Type is the base type of the object type. This will be out of range of the subtype only if the subtype range excludes this value.

`Modular types`
: Objects of a modular type are initialized to typ'Last. This will be out of range of the subtype only if the subtype excludes this value.

`Enumeration types`

Objects of an enumeration type are initialized to all one-bits, i.e. to the value `2 ** typ'Size - 1`. This will be out of range of the enumeration subtype in all cases except where the subtype contains exactly 2**8, 2**16, or 2**32 elements.

`pragma Long_Float`

Syntax:

```
pragma Long_Float (FLOAT_FORMAT);

FLOAT_FORMAT ::= D_Float | G_Float
```

This pragma is implemented only in the OpenVMS implementation of GNAT. It allows control over the internal representation chosen for the predefined type `Long_Float` and for floating point type representations with `digits` specified in the range 7 through 15. For further details on this pragma, see the *DEC Ada Language Reference Manual*, section 3.5.7b. Note that to use this pragma, the standard runtime libraries must be recompiled. See the description of the `GNAT LIBRARY` command in the OpenVMS version of the GNAT User's Guide for details on the use of this command.

`pragma Machine_Attribute ...`

Syntax:

```
pragma Machine_Attribute (
  [Attribute_Name =>] string_EXPRESSION,
  [Entity         =>] LOCAL_NAME);
```

Machine dependent attributes can be specified for types and/or declarations. Currently only subprogram entities are supported. This pragma is semantically equivalent to `__attribute__((`*`string_expression`*`))` in GNU C, where *`string_expression`* is recognized by the GNU C macros `VALID_MACHINE_TYPE_ATTRIBUTE` and `VALID_MACHINE_DECL_ATTRIBUTE` which are defined in the configuration header file '`tm.h`' for each machine. See the GCC manual for further information.

`pragma Main_Storage`

Syntax:

```
pragma Main_Storage
  (MAIN_STORAGE_OPTION [, MAIN_STORAGE_OPTION]);

MAIN_STORAGE_OPTION ::=
  [WORKING_STORAGE =>] static_SIMPLE_EXPRESSION
| [TOP_GUARD       =>] static_SIMPLE_EXPRESSION
```

This pragma is provided for compatibility with OpenVMS Vax Systems. It has no effect in GNAT, other than being syntax checked. Note that the pragma also has no effect in DEC Ada 83 for OpenVMS Alpha Systems.

`pragma No_Return`

Syntax:

```
pragma No_Return (procedure_LOCAL_NAME);
```

procedure_local_NAME must refer to one or more procedure declarations in the current declarative part. A procedure to which this pragma is applied

may not contain any explicit `return` statements, and also may not contain any implicit return statements from falling off the end of a statement sequence. One use of this pragma is to identify procedures whose only purpose is to raise an exception.

Another use of this pragma is to suppress incorrect warnings about missing returns in functions, where the last statement of a function statement sequence is a call to such a procedure.

`pragma Passive`

Syntax:

```
pragma Passive ([Semaphore | No]);
```

Syntax checked, but otherwise ignored by GNAT. This is recognized for compatibility with DEC Ada 83 implementations, where it is used within a task definition to request that a task be made passive. If the argument `Semaphore` is present, or no argument is omitted, then DEC Ada 83 treats the pragma as an assertion that the containing task is passive and that optimization of context switch with this task is permitted and desired. If the argument `No` is present, the task must not be optimized. GNAT does not attempt to optimize any tasks in this manner (since protected objects are available in place of passive tasks).

`pragma Polling`

Syntax:

```
pragma Polling (ON | OFF);
```

This pragma controls the generation of polling code. This is normally off. If `pragma Polling (ON)` is used then periodic calls are generated to the routine `Ada.Exceptions.Poll`. This routine is a separate unit in the runtime library, and can be found in file '`a-excpol.adb`'.

Pragma `Polling` can appear as a configuration pragma (for example it can be placed in the '`gnat.adc`' file) to enable polling globally, or it can be used in the statement or declaration sequence to control polling more locally.

A call to the polling routine is generated at the start of every loop and at the start of every subprogram call. This guarantees that the `Poll` routine is called frequently, and places an upper bound (determined by the complexity of the code) on the period between two `Poll` calls.

The primary purpose of the polling interface is to enable asynchronous aborts on targets that cannot otherwise support it (for example Windows NT), but it may be used for any other purpose requiring periodic polling. The standard version is null, and can be replaced by a user program. This will require re-compilation of the `Ada.Exceptions` package that can be found in files '`a-except.ads`' and '`a-except.adb`'.

A standard alternative unit (in file '`4wexcpol.adb`' in the standard GNAT distribution) is used to enable the asynchronous abort capability on targets that do not normally support the capability. The version of `Poll` in this file makes a call to the appropriate runtime routine to test for an abort condition.

Note that polling can also be enabled by use of the `-gnatP` switch. See the *GNAT User's Guide* for details.

pragma Propagate_Exceptions

Syntax:

```
pragma Propagate_Exceptions (subprogram_LOCAL_NAME);
```

This pragma indicates that the given entity, which is the name of an imported foreign-language subprogram may receive an Ada exception, and that the exception should be propagated. It is relevant only if zero cost exception handling is in use, and is thus never needed if the alternative `longjmp` / `setjmp` implementation of exceptions is used (although it is harmless to use it in such cases).

The implementation of fast exceptions always properly propagates exceptions through Ada code, as described in the Ada Reference Manual. However, this manual is silent about the propagation of exceptions through foreign code. For example, consider the situation where `P1` calls `P2`, and `P2` calls `P3`, where `P1` and `P3` are in Ada, but `P2` is in C. `P3` raises an Ada exception. The question is whether or not it will be propagated through `P2` and can be handled in `P1`.

For the `longjmp` / `setjmp` implementation of exceptions, the answer is always yes. For some targets on which zero cost exception handling is implemented, the answer is also always yes. However, there are some targets, notably in the current version all x86 architecture targets, in which the answer is that such propagation does not happen automatically. If such propagation is required on these targets, it is mandatory to use `Propagate_Exceptions` to name all foreign language routines through which Ada exceptions may be propagated.

pragma Psect_Object

Syntax:

```
pragma Psect_Object
     [Internal =>] LOCAL_NAME,
  [, [External =>] EXTERNAL_SYMBOL]
  [, [Size     =>] EXTERNAL_SYMBOL]

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION
```

This pragma is identical in effect to pragma `Common_Object`.

pragma Pure_Function

Syntax:

```
pragma Pure_Function ([Entity =>] function_LOCAL_NAME);
```

This pragma appears in the same declarative part as a function declaration (or a set of function declarations if more than one overloaded declaration exists, in which case the pragma applies to all entities). If specifies that the function `Entity` is to be considered pure for the purposes of code generation. This means that the compiler can assume that there are no side effects, and in particular that two calls with identical arguments produce the same result. It also means that the function can be used in an address clause.

Note that, quite deliberately, there are no static checks to try to ensure that this promise is met, so `Pure_Function` can be used with functions that are conceptually pure, even if they do modify global variables. For example, a square root function that is instrumented to count the number of times it is

called is still conceptually pure, and can still be optimized, even though it modifies a global variable (the count). Memo functions are another example (where a table of previous calls is kept and consulted to avoid re-computation).

Note: Most functions in a `Pure` package are automatically pure, and there is no need to use pragma `Pure_Function` for such functions. An exception is any function that has at least one formal of type `System.Address` or a type derived from it. Such functions are not considered pure by default, since the compiler assumes that the `Address` parameter may be functioning as a pointer and that the referenced data may change even if the address value does not. The use of pragma `Pure_Function` for such a function will override this default assumption, and cause the compiler to treat such a function as pure.

Note: If pragma `Pure_Function` is applied to a renamed function, it applies to the underlying renamed function. This can be used to disambiguate cases of overloading where some but not all functions in a set of overloaded functions are to be designated as pure.

`pragma Ravenscar`

Syntax:

```
pragma Ravenscar
```

A configuration pragma that establishes the following set of restrictions:

`No_Abort_Statements`
: [RM D.7] There are no abort_statements, and there are no calls to Task_Identification.Abort_Task.

`No_Select_Statements`
: There are no select_statements.

`No_Task_Hierarchy`
: [RM D.7] All (non-environment) tasks depend directly on the environment task of the partition.

`No_Task_Allocators`
: [RM D.7] There are no allocators for task types or types containing task subcomponents.

`No_Dynamic_Priorities`
: [RM D.7] There are no semantic dependencies on the package Dynamic_Priorities.

`No_Terminate_Alternatives`
: [RM D.7] There are no selective_accepts with terminate_alternatives

`No_Dynamic_Interrupts`
: There are no semantic dependencies on Ada.Interrupts.

`No_Protected_Type_Allocators`
: There are no allocators for protected types or types containing protected subcomponents.

`No_Local_Protected_Objects`
: Protected objects and access types that designate such objects shall be declared only at library level.

`No_Requeue`
: Requeue statements are not allowed.

`No_Calendar`
: There are no semantic dependencies on the package Ada.Calendar.

`No_Relative_Delay`
: There are no delay_relative_statements.

`No_Task_Attributes`
: There are no semantic dependencies on the Ada.Task_Attributes package and there are no references to the attributes Callable and Terminated [RM 9.9].

`Static_Storage_Size`
: The expression for pragma Storage_Size is static.

`Boolean_Entry_Barriers`
: Entry barrier condition expressions shall be boolean objects which are declared in the protected type which contains the entry.

`Max_Asynchronous_Select_Nesting = 0`
: [RM D.7] Specifies the maximum dynamic nesting level of asynchronous_selects. A value of zero prevents the use of any asynchronous_select.

`Max_Task_Entries = 0`
: [RM D.7] Specifies the maximum number of entries per task. The bounds of every entry family of a task unit shall be static, or shall be defined by a discriminant of a subtype whose corresponding bound is static. A value of zero indicates that no rendezvous are possible. For the Ravenscar pragma, the value of Max_Task_Entries is always 0 (zero).

`Max_Protected_Entries = 1`
: [RM D.7] Specifies the maximum number of entries per protected type. The bounds of every entry family of a protected unit shall be static, or shall be defined by a discriminant of a subtype whose corresponding bound is static. For the Ravenscar pragma the value of Max_Protected_Entries is always 1.

`Max_Select_Alternatives = 0`
: [RM D.7] Specifies the maximum number of alternatives in a selective_accept. For the Ravenscar pragma the value if always 0.

`No_Task_Termination`
: Tasks which terminate are erroneous.

`No_Entry_Queue`
: No task can be queued on a protected entry. Note that this restrictions is checked at run time. The violation of this restriction generates a Program_Error exception.

This set of restrictions corresponds to the definition of the "Ravenscar Profile" for limited tasking, devised and published by the *International Real-Time Ada Workshop*, 1997.

The above set is a superset of the restrictions provided by pragma `Restricted_Run_Time`, it includes six additional restrictions (`Boolean_Entry_Barriers`, `No_Select_Statements`, `No_Calendar`, `Static_Storage_Size`, `No_Relative_Delay` and `No_Task_Termination`). This means that pragma `Ravenscar`, like the pragma `Restricted_Run_Time`, automatically causes the use of a simplified, more efficient version of the tasking run-time system.

`pragma Restricted_Run_Time`

Syntax:

```
pragma Restricted_Run_Time
```

A configuration pragma that establishes the following set of restrictions:

- No_Abort_Statements
- No_Asynchronous_Control
- No_Entry_Queue
- No_Task_Hierarchy
- No_Task_Allocators
- No_Dynamic_Priorities
- No_Terminate_Alternatives
- No_Dynamic_Interrupts
- No_Protected_Type_Allocators
- No_Local_Protected_Objects
- No_Requeue
- No_Task_Attributes
- Max_Asynchronous_Select_Nesting = 0
- Max_Task_Entries = 0
- Max_Protected_Entries = 1
- Max_Select_Alternatives = 0

This set of restrictions causes the automatic selection of a simplified version of the run time that provides improved performance for the limited set of tasking functionality permitted by this set of restrictions.

`pragma Share_Generic`

Syntax:

```
pragma Share_Generic (NAME {, NAME});
```

This pragma is recognized for compatibility with other Ada compilers but is ignored by GNAT. GNAT does not provide the capability for sharing of generic

code. All generic instantiations result in making an inlined copy of the template with appropriate substitutions.

`pragma Source_File_Name`

Syntax:

```
pragma Source_File_Name (
  [Unit_Name   =>] unit_NAME,
  Spec_File_Name =>  STRING_LITERAL);

pragma Source_File_Name (
  [Unit_Name   =>] unit_NAME,
  Body_File_Name =>  STRING_LITERAL);
```

Use this to override the normal naming convention. It is a configuration pragma, and so has the usual applicability of configuration pragmas (i.e. it applies to either an entire partition, or to all units in a compilation, or to a single unit, depending on how it is used. *unit_name* is mapped to *file_name_literal*. The identifier for the second argument is required, and indicates whether this is the file name for the spec or for the body.

Another form of the `Source_File_Name` pragma allows the specification of patterns defining alternative file naming schemes to apply to all files.

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

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

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

CASING_SPEC ::= Lowercase | Uppercase | Mixedcase
```

The first argument is a pattern that contains a single asterisk indicating the point at which the unit name is to be inserted in the pattern string to form the file name. The second argument is optional. If present it specifies the casing of the unit name in the resulting file name string. The default is lower case. Finally the third argument allows for systematic replacement of any dots in the unit name by the specified string literal.

For more details on the use of the `Source_File_Name` pragma, see the sections "Using Other File Names" and "Alternative File Naming Schemes" in the *GNAT User's Guide*.

`pragma Source_Reference`

Syntax:

```
pragma Source_Reference (INTEGER_LITERAL,
                         STRING_LITERAL);
```

This pragma must appear as the first line of a source file. *integer_literal* is the logical line number of the line following the pragma line (for use in error mes-

sages and debugging information). *string_literal* is a static string constant that specifies the file name to be used in error messages and debugging information. This is most notably used for the output of `gnatchop` with the `-r` switch, to make sure that the original unchopped source file is the one referred to.

The second argument must be a string literal, it cannot be a static string expression other than a string literal. This is because its value is needed for error messages issued by all phases of the compiler.

`pragma Stream_Convert`

Syntax:

```
pragma Stream_Convert (
  [Entity =>] type_LOCAL_NAME,
  [Read   =>] function_NAME,
  [Write  =>] function NAME);
```

This pragma provides an efficient way of providing stream functions for types defined in packages. Not only is it simpler to use than declaring the necessary functions with attribute representation clauses, but more significantly, it allows the declaration to made in such a way that the stream packages are not loaded unless they are needed. The use of the Stream_Convert pragma adds no overhead at all, unless the stream attributes are actually used on the designated type.

The first argument specifies the type for which stream functions are provided. The second parameter provides a function used to read values of this type. It must name a function whose argument type may be any subtype, and whose returned type must be the type given as the first argument to the pragma.

The meaning of the *Read* parameter is that if a stream attribute directly or indirectly specifies reading of the type given as the first parameter, then a value of the type given as the argument to the Read function is read from the stream, and then the Read function is used to convert this to the required target type.

Similarly the *Write* parameter specifies how to treat write attributes that directly or indirectly apply to the type given as the first parameter. It must have an input parameter of the type specified by the first parameter, and the return type must be the same as the input type of the Read function. The effect is to first call the Write function to convert to the given stream type, and then write the result type to the stream.

The Read and Write functions must not be overloaded subprograms. If necessary renamings can be supplied to meet this requirement. The usage of this attribute is best illustrated by a simple example, taken from the GNAT implementation of package Ada.Strings.Unbounded:

```
function To_Unbounded (S : String)
           return Unbounded_String
  renames To_Unbounded_String;

pragma Stream_Convert
  (Unbounded_String, To_Unbounded, To_String);
```

The specifications of the referenced functions, as given in the Ada 95 Reference Manual are:

```
function To_Unbounded_String (Source : String)
  return Unbounded_String;

function To_String (Source : Unbounded_String)
  return String;
```

The effect is that if the value of an unbounded string is written to a stream, then the representation of the item in the stream is in the same format used for `Standard.String`, and this same representation is expected when a value of this type is read from the stream.

`pragma Style_Checks`

Syntax:

```
pragma Style_Checks (string_LITERAL | ALL_CHECKS |
                     On | Off [, LOCAL_NAME]);
```

This pragma is used in conjunction with compiler switches to control the built in style checking provided by GNAT. The compiler switches, if set provide an initial setting for the switches, and this pragma may be used to modify these settings, or the settings may be provided entirely by the use of the pragma. This pragma can be used anywhere that a pragma is legal, including use as a configuration pragma (including use in the '`gnat.adc`' file).

The form with a string literal specifies which style options are to be activated. These are additive, so they apply in addition to any previously set style check options. The codes for the options are the same as those used in the `-gnaty` switch to `gcc` or `gnatmake`. For example the following two methods can be used to enable layout checking:

```
pragma Style_Checks ("l");
gcc -c -gnatyl ...
```

The form ALL_CHECKS activates all standard checks (its use is equivalent to the use of the `gnaty` switch with no options. See GNAT User's Guide for details.

The forms with `Off` and `On` can be used to temporarily disable style checks as shown in the following example:

```
pragma Style_Checks ("k"); -- requires keywords in lower case
pragma Style_Checks (Off); -- turn off style checks
NULL;                      -- this will not generate an error message
pragma Style_Checks (On);  -- turn style checks back on
NULL;                      -- this will generate an error message
```

Finally the two argument form is allowed only if the first argument is `On` or `Off`. The effect is to turn of semantic style checks for the specified entity, as shown in the following example:

```
pragma Style_Checks ("r"); -- require consistency of identifier casing
Arg : Integer;
Rf1 : Integer := ARG;      -- incorrect, wrong case
pragma Style_Checks (Off, Arg);
Rf2 : Integer := ARG;      -- OK, no error
```

`pragma Subtitle`

Syntax:

```
pragma Subtitle ([Subtitle =>] STRING_LITERAL);
```

This pragma is recognized for compatibility with other Ada compilers but is ignored by GNAT.

`pragma Suppress_All`

Syntax:

```
pragma Suppress_All;
```

This pragma can only appear immediately following a compilation unit. The effect is to apply `Suppress (All_Checks)` to the unit which it follows. This pragma is implemented for compatibility with DEC Ada 83 usage. The use of pragma `Suppress (All_Checks)` as a normal configuration pragma is the preferred usage in GNAT.

`pragma Suppress_Initialization`

Syntax:

```
pragma Suppress_Initialization ([Entity =>] type_Name);
```

This pragma suppresses any implicit or explicit initialization associated with the given type name for all variables of this type.

`pragma Task_Info`

Syntax

```
pragma Task_Info (EXPRESSION);
```

This pragma appears within a task definition (like pragma `Priority`) and applies to the task in which it appears. The argument must be of type `System.Task_Info.Task_Info_Type`. The `Task_Info` pragma provides system dependent control over aspect of tasking implementation, for example, the ability to map tasks to specific processors. For details on the facilities available for the version of GNAT that you are using, see the documentation in the specification of package System.Task_Info in the runtime library.

`pragma Task_Name`

Syntax

```
pragma Task_Name (string_EXPRESSION);
```

This pragma appears within a task definition (like pragma `Priority`) and applies to the task in which it appears. The argument must be of type String, and provides a name to be used for the task instance when the task is created. Note that this expression is not required to be static, and in particular, it can contain references to task discriminants. This facility can be used to provide different names for different tasks as they are created, as illustrated in the example below.

The task name is recorded internally in the run-time structures and is accessible to tools like the debugger. In addition the routine `Ada.Task_Identification.Image` will return this string, with a unique task address appended.

```
--  Example of the use of pragma Task_Name
```

```
with Ada.Task_Identification;
use Ada.Task_Identification;
with Text_IO; use Text_IO;
procedure t3 is

   type Astring is access String;

   task type Task_Typ (Name : access String) is
      pragma Task_Name (Name.all);
   end Task_Typ;

   task body Task_Typ is
      Nam : constant String := Image (Current_Task);
   begin
      Put_Line ("-->" & Nam (1 .. 14) & "<--");
   end Task_Typ;

   type Ptr_Task is access Task_Typ;
   Task_Var : Ptr_Task;

begin
   Task_Var :=
     new Task_Typ (new String'("This is task 1"));
   Task_Var :=
     new Task_Typ (new String'("This is task 2"));
end;
```

`pragma Task_Storage`

Syntax:

```
pragma Task_Storage
  [Task_Type =>] LOCAL_NAME,
  [Top_Guard =>] static_integer_EXPRESSION);
```

This pragma specifies the length of the guard area for tasks. The guard area is an additional storage area allocated to a task. A value of zero means that either no guard area is created or a minimal guard area is created, depending on the target. This pragma can appear anywhere a `Storage_Size` attribute definition clause is allowed for a task type.

`pragma Time_Slice`

Syntax:

```
pragma Time_Slice (static_duration_EXPRESSION);
```

For implementations of GNAT on operating systems where it is possible to supply a time slice value, this pragma may be used for this purpose. It is ignored if it is used in a system that does not allow this control, or if it appears in other than the main program unit. Note that the effect of this pragma is identical to the effect of the DEC Ada 83 pragma of the same name when operating under OpenVMS systems.

`pragma Title`

Syntax:

```
pragma Title (TITLING_OPTION [, TITLING OPTION]);

TITLING_OPTION ::=
  [Title    =>] STRING_LITERAL,
```

```
| [Subtitle =>] STRING_LITERAL
```

Syntax checked but otherwise ignored by GNAT. This is a listing control pragma used in DEC Ada 83 implementations to provide a title and/or subtitle for the program listing. The program listing generated by GNAT does not have titles or subtitles.

Unlike other pragmas, the full flexibility of named notation is allowed for this pragma, i.e. the parameters may be given in any order if named notation is used, and named and positional notation can be mixed following the normal rules for procedure calls in Ada.

`pragma Unchecked_Union`

Syntax:

```
pragma Unchecked_Union (first_subtype_LOCAL_NAME)
```

This pragma is used to declare that the specified type should be represented in a manner equivalent to a C union type, and is intended only for use in interfacing with C code that uses union types. In Ada terms, the named type must obey the following rules:

- It is a non-tagged non-limited record type.
- It has a single discrete discriminant with a default value.
- The component list consists of a single variant part.
- Each variant has a component list with a single component.
- No nested variants are allowed.
- No component has an explicit default value.
- No component has a non-static constraint.

In addition, given a type that meets the above requirements, the following restrictions apply to its use throughout the program:

- The discriminant name can be mentioned only in an aggregate.
- No subtypes may be created of this type.
- The type may not be constrained by giving a discriminant value.
- The type cannot be passed as the actual for a generic formal with a discriminant.

Equality and inequality operations on `unchecked_unions` are not available, since there is no discriminant to compare and the compiler does not even know how many bits to compare. It is implementation dependent whether this is detected at compile time as an illegality or whether it is undetected and considered to be an erroneous construct. In GNAT, a direct comparison is illegal, but GNAT does not attempt to catch the composite case (where two composites are compared that contain an unchecked union component), so such comparisons are simply considered erroneous.

The layout of the resulting type corresponds exactly to a C union, where each branch of the union corresponds to a single variant in the Ada record. The semantics of the Ada program is not changed in any way by the pragma, i.e. provided the above restrictions are followed, and no erroneous incorrect references to fields or erroneous comparisons occur, the semantics is exactly as

described by the Ada reference manual. Pragma `Suppress (Discriminant_Check)` applies implicitly to the type and the default convention is C

`pragma Unimplemented_Unit`

Syntax:

```
pragma Unimplemented_Unit;
```

If this pragma occurs in a unit that is processed by the compiler, GNAT aborts with the message '*xxx* `not implemented`', where *xxx* is the name of the current compilation unit. This pragma is intended to allow the compiler to handle unimplemented library units in a clean manner.

The abort only happens if code is being generated. Thus you can use specs of unimplemented packages in syntax or semantic checking mode.

`pragma Unreferenced`

Syntax:

```
pragma Unreferenced (local_Name {, local_Name});
```

This pragma signals that the entities whose names are listed are deliberately not referenced. This suppresses warnings about the entities being unreferenced, and in addition a warning will be generated if one of these entities is in fact referenced.

This is particularly useful for clearly signalling that a particular parameter is not referenced in some particular subprogram implementation and that this is deliberate. It can also be useful in the case of objects declared only for their initialization or finalization side effects.

If `local_Name` identifies more than one matching homonym in the current scope, then the entity most recently declared is the one to which the pragma applies.

`pragma Unreserve_All_Interrupts`

Syntax:

```
pragma Unreserve_All_Interrupts;
```

Normally certain interrupts are reserved to the implementation. Any attempt to attach an interrupt causes Program_Error to be raised, as described in RM C.3.2(22). A typical example is the `SIGINT` interrupt used in many systems for an *`Ctrl-C`* interrupt. Normally this interrupt is reserved to the implementation, so that *`Ctrl-C`* can be used to interrupt execution.

If the pragma `Unreserve_All_Interrupts` appears anywhere in any unit in a program, then all such interrupts are unreserved. This allows the program to handle these interrupts, but disables their standard functions. For example, if this pragma is used, then pressing *`Ctrl-C`* will not automatically interrupt execution. However, a program can then handle the `SIGINT` interrupt as it chooses.

For a full list of the interrupts handled in a specific implementation, see the source code for the specification of `Ada.Interrupts.Names` in file '`a-intnam.ads`'. This is a target dependent file that contains the list of interrupts recognized for a given target. The documentation in this file also specifies what interrupts are affected by the use of the `Unreserve_All_Interrupts` pragma.

`pragma Unsuppress`

Syntax:

```
pragma Unsuppress (IDENTIFIER [, [On =>] NAME]);
```

This pragma undoes the effect of a previous pragma `Suppress`. If there is no corresponding pragma `Suppress` in effect, it has no effect. The range of the effect is the same as for pragma `Suppress`. The meaning of the arguments is identical to that used in pragma `Suppress`.

One important application is to ensure that checks are on in cases where code depends on the checks for its correct functioning, so that the code will compile correctly even if the compiler switches are set to suppress checks.

`pragma Use_VADS_Size`

Syntax:

```
pragma Use_VADS_Size;
```

This is a configuration pragma. In a unit to which it applies, any use of the 'Size attribute is automatically interpreted as a use of the 'VADS_Size attribute. Note that this may result in incorrect semantic processing of valid Ada 95 programs. This is intended to aid in the handling of legacy code which depends on the interpretation of Size as implemented in the VADS compiler. See description of the VADS_Size attribute for further details.

`pragma Validity_Checks`

Syntax:

```
pragma Validity_Checks (string_LITERAL | ALL_CHECKS | On | Off);
```

This pragma is used in conjunction with compiler switches to control the built in validity checking provided by GNAT. The compiler switches, if set provide an initial setting for the switches, and this pragma may be used to modify these settings, or the settings may be provided entirely by the use of the pragma. This pragma can be used anywhere that a pragma is legal, including use as a configuration pragma (including use in the '`gnat.adc`' file).

The form with a string literal specifies which validity options are to be activated. The validity checks are first set to include only the default reference manual settings, and then a string of letters in the string specifies the exact set of options required. The form of this string is exactly as described for the `-gnatVx` compiler switch (see the GNAT users guide for details). For example the following two methods can be used to enable validity checking for mode `in` and `in out` subprogram parameters:

```
pragma Validity_Checks ("im");
gcc -c -gnatVim ...
```

The form ALL_CHECKS activates all standard checks (its use is equivalent to the use of the `gnatva` switch.

The forms with `Off` and `On` can be used to temporarily disable validity checks as shown in the following example:

```
pragma Validity_Checks ("c"); -- validity checks for copies
pragma Validity_Checks (Off); -- turn off validity checks
```

```
A := B;                       -- B will not be validity checked
pragma Validity_Checks (On);  -- turn validity checks back on
A := C;                       -- C will be validity checked
```

`pragma Volatile`

Syntax:

```
pragma Volatile (local_NAME)
```

This pragma is defined by the Ada 95 Reference Manual, and the GNAT implementation is fully conformant with this definition. The reason it is mentioned in this section is that a pragma of the same name was supplied in some Ada 83 compilers, including DEC Ada 83. The Ada 95 implementation of pragma Volatile is upwards compatible with the implementation in Dec Ada 83.

`pragma Warnings`

Syntax:

```
pragma Warnings (On | Off [, LOCAL_NAME]);
```

Normally warnings are enabled, with the output being controlled by the command line switch. Warnings (`Off`) turns off generation of warnings until a Warnings (`On`) is encountered or the end of the current unit. If generation of warnings is turned off using this pragma, then no warning messages are output, regardless of the setting of the command line switches.

The form with a single argument is a configuration pragma.

If the *local_name* parameter is present, warnings are suppressed for the specified entity. This suppression is effective from the point where it occurs till the end of the extended scope of the variable (similar to the scope of `Suppress`).

`pragma Weak_External`

Syntax:

```
pragma Weak_External ([Entity =>] LOCAL_NAME);
```

This pragma specifies that the given entity should be marked as a weak external (one that does not have to be resolved) for the linker. For further details, consult the GCC manual.

2 Implementation Defined Attributes

Ada 95 defines (throughout the Ada 95 reference manual, summarized in annex K), a set of attributes that provide useful additional functionality in all areas of the language. These language defined attributes are implemented in GNAT and work as described in the Ada 95 Reference Manual.

In addition, Ada 95 allows implementations to define additional attributes whose meaning is defined by the implementation. GNAT provides a number of these implementation-dependent attributes which can be used to extend and enhance the functionality of the compiler. This section of the GNAT reference manual describes these additional attributes.

Note that any program using these attributes may not be portable to other compilers (although GNAT implements this set of attributes on all platforms). Therefore if portability to other compilers is an important consideration, you should minimize the use of these attributes.

`Abort_Signal`

: `Standard'Abort_Signal` (`Standard` is the only allowed prefix) provides the entity for the special exception used to signal task abort or asynchronous transfer of control. Normally this attribute should only be used in the tasking runtime (it is highly peculiar, and completely outside the normal semantics of Ada, for a user program to intercept the abort exception).

`Address_Size`

: `Standard'Address_Size` (`Standard` is the only allowed prefix) is a static constant giving the number of bits in an `Address`. It is used primarily for constructing the definition of `Memory_Size` in package `Standard`, but may be freely used in user programs and has the advantage of being static, while a direct reference to System.Address'Size is non-static because Address is a private type.

`Asm_Input`

: The `Asm_Input` attribute denotes a function that takes two parameters. The first is a string, the second is an expression of the type designated by the prefix. The first (string) argument is required to be a static expression, and is the constraint for the parameter, (e.g. what kind of register is required). The second argument is the value to be used as the input argument. The possible values for the constant are the same as those used in the RTL, and are dependent on the configuration file used to built the GCC back end. Chapter 11 [Machine Code Insertions], page 161

`Asm_Output`

: The `Asm_Output` attribute denotes a function that takes two parameters. The first is a string, the second is the name of a variable of the type designated by the attribute prefix. The first (string) argument is required to be a static expression and designates the constraint for the parameter (e.g. what kind of register is required). The second argument is the variable to be updated with the result. The possible values for constraint are the same as those used in the RTL, and are dependent on the configuration file used to build the GCC back end. If there are no output operands, then this argument may either

be omitted, or explicitly given as `No_Output_Operands`. Chapter 11 [Machine Code Insertions], page 161

`AST_Entry`

This attribute is implemented only in OpenVMS versions of GNAT. Applied to the name of an entry, it yields a value of the predefined type AST_Handler (declared in the predefined package System, as extended by the use of pragma `Extend_System (Aux_DEC)`). This value enables the given entry to be called when an AST occurs. For further details, refer to the *DEC Ada Language Reference Manual*, section 9.12a.

`Bit` *obj*`'Bit`, where *obj* is any object, yields the bit offset within the storage unit (byte) that contains the first bit of storage allocated for the object. The value of this attribute is of the type `Universal_Integer`, and is always a non-negative number not exceeding the value of `System.Storage_Unit`.

For an object that is a variable or a constant allocated in a register, the value is zero. (The use of this attribute does not force the allocation of a variable to memory).

For an object that is a formal parameter, this attribute applies to either the matching actual parameter or to a copy of the matching actual parameter.

For an access object the value is zero. Note that *obj*`.all'Bit` is subject to an `Access_Check` for the designated object. Similarly for a record component *X*`.`*C*`'Bit` is subject to a discriminant check and *X*`(`*I*`).Bit` and *X*`(`*I1*`..`*I2*`)'Bit` are subject to index checks.

This attribute is designed to be compatible with the DEC Ada 83 definition and implementation of the `Bit` attribute.

`Bit_Position`

R`.`*C*`'Bit`, where *R* is a record object and C is one of the fields of the record type, yields the bit offset within the record contains the first bit of storage allocated for the object. The value of this attribute is of the type `Universal_Integer`. The value depends only on the field *C* and is independent of the alignment of the containing record *R*.

`Code_Address`

The `'Address` attribute may be applied to subprograms in Ada 95, but the intended effect from the Ada 95 reference manual seems to be to provide an address value which can be used to call the subprogram by means of an address clause as in the following example:

```
procedure K is ...

procedure L;
for L'Address use K'Address;
pragma Import (Ada, L);
```

A call to `L` is then expected to result in a call to `K`. In Ada 83, where there were no access-to-subprogram values, this was a common work around for getting the effect of an indirect call. GNAT implements the above use of `Address` and the technique illustrated by the example code works correctly.

However, for some purposes, it is useful to have the address of the start of the generated code for the subprogram. On some architectures, this is not necessarily the same as the `Address` value described above. For example, the `Address` value may reference a subprogram descriptor rather than the subprogram itself.

The `'Code_Address` attribute, which can only be applied to subprogram entities, always returns the address of the start of the generated code of the specified subprogram, which may or may not be the same value as is returned by the corresponding `'Address` attribute.

`Default_Bit_Order`

`Standard'Default_Bit_Order` (`Standard` is the only permissible prefix), provides the value `System.Default_Bit_Order` as a `Pos` value (0 for `High_Order_First`, 1 for `Low_Order_First`). This is used to construct the definition of `Default_Bit_Order` in package `System`.

`Elaborated`

The prefix of the `'Elaborated` attribute must be a unit name. The value is a Boolean which indicates whether or not the given unit has been elaborated. This attribute is primarily intended for internal use by the generated code for dynamic elaboration checking, but it can also be used in user programs. The value will always be True once elaboration of all units has been completed.

`Elab_Body`

This attribute can only be applied to a program unit name. It returns the entity for the corresponding elaboration procedure for elaborating the body of the referenced unit. This is used in the main generated elaboration procedure by the binder and is not normally used in any other context. However, there may be specialized situations in which it is useful to be able to call this elaboration procedure from Ada code, e.g. if it is necessary to do selective re-elaboration to fix some error.

`Elab_Spec`

This attribute can only be applied to a program unit name. It returns the entity for the corresponding elaboration procedure for elaborating the specification of the referenced unit. This is used in the main generated elaboration procedure by the binder and is not normally used in any other context. However, there may be specialized situations in which it is useful to be able to call this elaboration procedure from Ada code, e.g. if it is necessary to do selective re-elaboration to fix some error.

`Emax` The `Emax` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

`Enum_Rep` For every enumeration subtype *S*, *S*`'Enum_Rep` denotes a function with the following specification:

```
function S'Enum_Rep (Arg : S'Base)
  return Universal_Integer;
```

It is also allowable to apply `Enum_Rep` directly to an object of an enumeration type or to a non-overloaded enumeration literal. In this case *S*`'Enum_Rep` is equivalent to *typ*`'Enum_Rep(`*S*`)` where *typ* is the type of the enumeration literal or object.

The function returns the representation value for the given enumeration value. This will be equal to value of the `Pos` attribute in the absence of an enumeration representation clause. This is a static attribute (i.e. the result is static if the argument is static).

S`'Enum_Rep` can also be used with integer types and objects, in which case it simply returns the integer value. The reason for this is to allow it to be used for `(<>)` discrete formal arguments in a generic unit that can be instantiated with either enumeration types or integer types. Note that if `Enum_Rep` is used on a modular type whose upper bound exceeds the upper bound of the largest signed integer type, and the argument is a variable, so that the universal integer calculation is done at run-time, then the call to `Enum_Rep` may raise `Constraint_Error`.

`Epsilon` The `Epsilon` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

`Fixed_Value`

For every fixed-point type *S*, *S*`'Fixed_Value` denotes a function with the following specification:

```
function S'Fixed_Value (Arg : Universal_Integer)
  return S;
```

The value returned is the fixed-point value *V* such that

```
V = Arg * S'Small
```

The effect is thus equivalent to first converting the argument to the integer type used to represent *S*, and then doing an unchecked conversion to the fixed-point type. This attribute is primarily intended for use in implementation of the input-output functions for fixed-point values.

`Has_Discriminants`

The prefix of the `Has_Discriminants` attribute is a type. The result is a Boolean value which is True if the type has discriminants, and False otherwise. The intended use of this attribute is in conjunction with generic definitions. If the attribute is applied to a generic private type, it indicates whether or not the corresponding actual type has discriminants.

`Img` The `Img` attribute differs from `Image` in that it may be applied to objects as well as types, in which case it gives the `Image` for the subtype of the object. This is convenient for debugging:

```
Put_Line ("X = " & X'Img);
```

has the same meaning as the more verbose:

```
Put_Line ("X = " & type'Image (X));
```

where *type* is the subtype of the object X.

`Integer_Value`

For every integer type *S*, *S*`'Integer_Value` denotes a function with the following specification:

```
function S'Integer_Value (Arg : Universal_Fixed)
  return S;
```

The value returned is the integer value *V*, such that

```
Arg = V * type'Small
```

The effect is thus equivalent to first doing an unchecked convert from the fixed-point type to its corresponding implementation type, and then converting the result to the target integer type. This attribute is primarily intended for use in implementation of the standard input-output functions for fixed-point values.

`Large` The `Large` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

`Machine_Size`
This attribute is identical to the `Object_Size` attribute. It is provided for compatibility with the DEC Ada 83 attribute of this name.

`Mantissa` The `Mantissa` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

`Max_Interrupt_Priority`
`Standard'Max_Interrupt_Priority` (`Standard` is the only permissible prefix), provides the value `System.Max_Interrupt_Priority` and is intended primarily for constructing this definition in package `System`.

`Max_Priority`
`Standard'Max_Priority` (`Standard` is the only permissible prefix) provides the value `System.Max_Priority` and is intended primarily for constructing this definition in package `System`.

`Maximum_Alignment`
`Standard'Maximum_Alignment` (`Standard` is the only permissible prefix) provides the maximum useful alignment value for the target. This is a static value that can be used to specify the alignment for an object, guaranteeing that it is properly aligned in all cases. This is useful when an external object is imported and its alignment requirements are unknown.

`Mechanism_Code`
function`'Mechanism_Code` yields an integer code for the mechanism used for the result of function, and *subprogram*`'Mechanism_Code (`*n*`)` yields the mechanism used for formal parameter number *n* (a static integer value with 1 meaning the first parameter) of *subprogram*. The code returned is:

1 by copy (value)

2 by reference

3 by descriptor (default descriptor class)

4 by descriptor (UBS: unaligned bit string)

5 by descriptor (UBSB: aligned bit string with arbitrary bounds)

6 by descriptor (UBA: unaligned bit array)

7 by descriptor (S: string, also scalar access type parameter)

8 by descriptor (SB: string with arbitrary bounds)

9 by descriptor (A: contiguous array)

10 by descriptor (NCA: non-contiguous array)

Values from 3 through 10 are only relevant to Digital OpenVMS implementations.

`Null_Parameter`

A reference *T*`'Null_Parameter` denotes an imaginary object of type or subtype *T* allocated at machine address zero. The attribute is allowed only as the default expression of a formal parameter, or as an actual expression of a subprogram call. In either case, the subprogram must be imported.

The identity of the object is represented by the address zero in the argument list, independent of the passing mechanism (explicit or default).

This capability is needed to specify that a zero address should be passed for a record or other composite object passed by reference. There is no way of indicating this without the `Null_Parameter` attribute.

`Object_Size`

The size of an object is not necessarily the same as the size of the type of an object. This is because by default object sizes are increased to be a multiple of the alignment of the object. For example, `Natural'Size` is 31, but by default objects of type `Natural` will have a size of 32 bits. Similarly, a record containing an integer and a character:

```
type Rec is record
   I : Integer;
   C : Character;
end record;
```

will have a size of 40 (that is `Rec'Size` will be 40. The alignment will be 4, because of the integer field, and so the default size of record objects for this type will be 64 (8 bytes).

The *type*`'Object_Size` attribute has been added to GNAT to allow the default object size of a type to be easily determined. For example, `Natural'Object_Size` is 32, and `Rec'Object_Size` (for the record type in the above example) will be 64. Note also that, unlike the situation with the `Size` attribute as defined in the Ada RM, the `Object_Size` attribute can be specified individually for different subtypes. For example:

```
type R is new Integer;
subtype R1 is R range 1 .. 10;
subtype R2 is R range 1 .. 10;
for R2'Object_Size use 8;
```

In this example, `R'Object_Size` and `R1'Object_Size` are both 32 since the default object size for a subtype is the same as the object size for the parent subtype. This means that objects of type `R` or `R1` will by default be 32 bits (four bytes). But objects of type `R2` will be only 8 bits (one byte), since `R2'Object_Size` has been set to 8.

`Passed_By_Reference`

type`'Passed_By_Reference` for any subtype *type* returns a value of type `Boolean` value that is `True` if the type is normally passed by reference and `False` if the type is normally passed by copy in calls. For scalar types, the result is always `False` and is static. For non-scalar types, the result is non-static.

`Range_Length`

type`'Range_Length` for any discrete type *type* yields the number of values represented by the subtype (zero for a null range). The result is static for static subtypes. `Range_Length` applied to the index subtype of a one dimensional array always gives the same result as `Range` applied to the array itself.

`Safe_Emax`

The `Safe_Emax` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

`Safe_Large`

The `Safe_Large` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

`Safe_Large`

The `Safe_Large` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

`Small` The `Small` attribute is defined in Ada 95 only for fixed-point types. GNAT also allows this attribute to be applied to floating-point types for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute when applied to floating-point types.

`Storage_Unit`

`Standard'Storage_Unit` (`Standard` is the only permissible prefix) provides the value `System.Storage_Unit` and is intended primarily for constructing this definition in package `System`.

`Tick` `Standard'Tick` (`Standard` is the only permissible prefix) provides the value of `System.Tick` and is intended primarily for constructing this definition in package `System`.

`To_Address`

The `System'To_Address` (`System` is the only permissible prefix) denotes a function identical to `System.Storage_Elements.To_Address` except that it is a static attribute. This means that if its argument is a static expression, then the result of the attribute is a static expression. The result is that such an expression can be used in contexts (e.g. preelaborable packages) which require a static expression and where the function call could not be used (since the function call is always non-static, even if its argument is static).

`Type_Class`

type`'Type_Class` for any type or subtype *type* yields the value of the type class for the full type of *type*. If *type* is a generic formal type, the value is the value for the corresponding actual subtype. The value of this attribute is of type `System.Aux_DEC.Type_Class`, which has the following definition:

```
type Type_Class is
  (Type_Class_Enumeration,
   Type_Class_Integer,
   Type_Class_Fixed_Point,
```

```
Type_Class_Floating_Point,
Type_Class_Array,
Type_Class_Record,
Type_Class_Access,
Type_Class_Task,
Type_Class_Address);
```

Protected types yield the value `Type_Class_Task`, which thus applies to all concurrent types. This attribute is designed to be compatible with the DEC Ada 83 attribute of the same name.

`UET_Address`

The `UET_Address` attribute can only be used for a prefix which denotes a library package. It yields the address of the unit exception table when zero cost exception handling is used. This attribute is intended only for use within the GNAT implementation. See the unit `Ada.Exceptions` in files '`a-except.ads`' and '`a-except.adb`' for details on how this attribute is used in the implementation.

`Universal_Literal_String`

The prefix of `Universal_Literal_String` must be a named number. The static result is the string consisting of the characters of the number as defined in the original source. This allows the user program to access the actual text of named numbers without intermediate conversions and without the need to enclose the strings in quotes (which would preclude their use as numbers). This is used internally for the construction of values of the floating-point attributes from the file '`ttypef.ads`', but may also be used by user programs.

`Unrestricted_Access`

The `Unrestricted_Access` attribute is similar to `Access` except that all accessibility and aliased view checks are omitted. This is a user-beware attribute. It is similar to `Address`, for which it is a desirable replacement where the value desired is an access type. In other words, its effect is identical to first applying the `Address` attribute and then doing an unchecked conversion to a desired access type. In GNAT, but not necessarily in other implementations, the use of static chains for inner level subprograms means that `Unrestricted_Access` applied to a subprogram yields a value that can be called as long as the subprogram is in scope (normal Ada 95 accessibility rules restrict this usage).

`VADS_Size`

The `'VADS_Size` attribute is intended to make it easier to port legacy code which relies on the semantics of `'Size` as implemented by the VADS Ada 83 compiler. GNAT makes a best effort at duplicating the same semantic interpretation. In particular, `'VADS_Size` applied to a predefined or other primitive type with no Size clause yields the Object_Size (for example, `Natural'Size` is 32 rather than 31 on typical machines). In addition `'VADS_Size` applied to an object gives the result that would be obtained by applying the attribute to the corresponding type.

`Value_Size`

`type``'Value_Size` is the number of bits required to represent a value of the given subtype. It is the same as *`type`*`'Size`, but, unlike `Size`, may be set for non-first subtypes.

`Wchar_T_Size`

`Standard'Wchar_T_Size` (`Standard` is the only permissible prefix) provides the size in bits of the C `wchar_t` type primarily for constructing the definition of this type in package `Interfaces.C`.

`Word_Size`

`Standard'Word_Size` (`Standard` is the only permissible prefix) provides the value `System.Word_Size` and is intended primarily for constructing this definition in package `System`.

3 Implementation Advice

The main text of the Ada 95 Reference Manual describes the required behavior of all Ada 95 compilers, and the GNAT compiler conforms to these requirements.

In addition, there are sections throughout the Ada 95 reference manual headed by the phrase "implementation advice". These sections are not normative, i.e. they do not specify requirements that all compilers must follow. Rather they provide advice on generally desirable behavior. You may wonder why they are not requirements. The most typical answer is that they describe behavior that seems generally desirable, but cannot be provided on all systems, or which may be undesirable on some systems.

As far as practical, GNAT follows the implementation advice sections in the Ada 95 Reference Manual. This chapter contains a table giving the reference manual section number, paragraph number and several keywords for each advice. Each entry consists of the text of the advice followed by the GNAT interpretation of this advice. Most often, this simply says "followed", which means that GNAT follows the advice. However, in a number of cases, GNAT deliberately deviates from this advice, in which case the text describes what GNAT does and why.

1.1.3(20): Error Detection

> If an implementation detects the use of an unsupported Specialized Needs Annex feature at run time, it should raise `Program_Error` if feasible.

Not relevant. All specialized needs annex features are either supported, or diagnosed at compile time.

1.1.3(31): Child Units

> If an implementation wishes to provide implementation-defined extensions to the functionality of a language-defined library unit, it should normally do so by adding children to the library unit.

Followed.

1.1.5(12): Bounded Errors

> If an implementation detects a bounded error or erroneous execution, it should raise `Program_Error`.

Followed in all cases in which the implementation detects a bounded error or erroneous execution. Not all such situations are detected at runtime.

2.8(16): Pragmas

Normally, implementation-defined pragmas should have no semantic effect for error-free programs; that is, if the implementation-defined pragmas are removed from a working program, the program should still be legal, and should still have the same semantics.

The following implementation defined pragmas are exceptions to this rule:

`Abort_Defer`
Affects semantics

`Ada_83` Affects legality

`Assert` Affects semantics

`CPP_Class`
Affects semantics

`CPP_Constructor`
Affects semantics

`CPP_Virtual`
Affects semantics

`CPP_Vtable`
Affects semantics

`Debug` Affects semantics

`Interface_Name`
Affects semantics

`Machine_Attribute`
Affects semantics

`Unimplemented_Unit`
Affects legality

`Unchecked_Union`
Affects semantics

In each of the above cases, it is essential to the purpose of the pragma that this advice not be followed. For details see the separate section on implementation defined pragmas.

2.8(17-19): Pragmas

Normally, an implementation should not define pragmas that can make an illegal program legal, except as follows:

A pragma used to complete a declaration, such as a pragma `Import`;

> A pragma used to configure the environment by adding, removing, or replacing `library_items`.

See response to paragraph 16 of this same section.

3.5.2(5): Alternative Character Sets

> If an implementation supports a mode with alternative interpretations for `Character` and `Wide_Character`, the set of graphic characters of `Character` should nevertheless remain a proper subset of the set of graphic characters of `Wide_Character`. Any character set "localizations" should be reflected in the results of the subprograms defined in the language-defined package `Characters.Handling` (see A.3) available in such a mode. In a mode with an alternative interpretation of `Character`, the implementation should also support a corresponding change in what is a legal `identifier_letter`.

Not all wide character modes follow this advice, in particular the JIS and IEC modes reflect standard usage in Japan, and in these encoding, the upper half of the Latin-1 set is not part of the wide-character subset, since the most significant bit is used for wide character encoding. However, this only applies to the external forms. Internally there is no such restriction.

3.5.4(28): Integer Types

> An implementation should support `Long_Integer` in addition to `Integer` if the target machine supports 32-bit (or longer) arithmetic. No other named integer subtypes are recommended for package `Standard`. Instead, appropriate named integer subtypes should be provided in the library package `Interfaces` (see B.2).

`Long_Integer` is supported. Other standard integer types are supported so this advice is not fully followed. These types are supported for convenient interface to C, and so that all hardware types of the machine are easily available.

3.5.4(29): Integer Types

> An implementation for a two's complement machine should support modular types with a binary modulus up to `System.Max_Int*2+2`. An implementation should support a non-binary modules up to `Integer'Last`.

Followed.

3.5.5(8): Enumeration Values

> For the evaluation of a call on *`S`*`'Pos` for an enumeration subtype, if the value of the operand does not correspond to the internal code for any enumeration literal of its type (perhaps due to an un-initialized variable), then the implementation should raise `Program_Error`. This is particularly important for enumeration types with noncontiguous internal codes specified by an enumeration_representation_clause.

Followed.

3.5.7(17): Float Types

> An implementation should support `Long_Float` in addition to `Float` if the target machine supports 11 or more digits of precision. No other named floating point subtypes are recommended for package `Standard`. Instead, appropriate named floating point subtypes should be provided in the library package `Interfaces` (see B.2).

`Short_Float` and `Long_Long_Float` are also provided. The former provides improved compatibility with other implementations supporting this type. The latter corresponds to the highest precision floating-point type supported by the hardware. On most machines, this will be the same as `Long_Float`, but on some machines, it will correspond to the IEEE extended form. The notable case is all ia32 (x86) implementations, where `Long_Long_Float` corresponds to the 80-bit extended precision format supported in hardware on this processor. Note that the 128-bit format on SPARC is not supported, since this is a software rather than a hardware format.

3.6.2(11): Multidimensional Arrays

> An implementation should normally represent multidimensional arrays in row-major order, consistent with the notation used for multidimensional array aggregates (see 4.3.3). However, if a pragma `Convention` (`Fortran`, ...) applies to a multidimensional array type, then column-major order should be used instead (see B.5, "Interfacing with Fortran").

Followed.

9.6(30-31): Duration'Small

> Whenever possible in an implementation, the value of `Duration'Small` should be no greater than 100 microseconds.

Followed. (`Duration'Small` = 10**(−9)).

> The time base for `delay_relative_statements` should be monotonic; it need not be the same time base as used for `Calendar.Clock`.

Followed.

10.2.1(12): Consistent Representation

> In an implementation, a type declared in a pre-elaborated package should have the same representation in every elaboration of a given version of the package, whether the elaborations occur in distinct executions of the same program, or in executions of distinct programs or partitions that include the given version.

Followed, except in the case of tagged types. Tagged types involve implicit pointers to a local copy of a dispatch table, and these pointers have representations which thus depend on a particular elaboration of the package. It is not easy to see how it would be possible to follow this advice without severely impacting efficiency of execution.

11.4.1(19): Exception Information

> `Exception_Message` by default and `Exception_Information` should produce information useful for debugging. `Exception_Message` should be short, about one line. `Exception_Information` can be long. `Exception_Message` should not include the `Exception_Name`. `Exception_Information` should include both the `Exception_Name` and the `Exception_Message`.

Followed. For each exception that doesn't have a specified `Exception_Message`, the compiler generates one containing the location of the raise statement. This location has the form "file:line", where file is the short file name (without path information) and line is the line number in the file. Note that in the case of the Zero Cost Exception mechanism, these messages become redundant with the Exception_Information that contains a full backtrace of the calling sequence, so they are disabled. To disable explicitly the generation of the source location message, use the Pragma `Discard_Names`.

11.5(28): Suppression of Checks

> The implementation should minimize the code executed for checks that have been suppressed.

Followed.

13.1 (21-24): Representation Clauses

> The recommended level of support for all representation items is qualified as follows:

> An implementation need not support representation items containing non-static expressions, except that an implementation should support a representation item for a given entity if each non-static expression in the representation item is a name that statically denotes a constant declared before the entity.

Followed. GNAT does not support non-static expressions in representation clauses unless they are constants declared before the entity. For example:

```
X : typ;
for X'Address use To_address (16#2000#);
```

will be rejected, since the To_Address expression is non-static. Instead write:

```
X_Address : constant Address : =
To_Address    ((16#2000#);
```

```
X : typ;
for X'Address use X_Address;
```

An implementation need not support a specification for the `Size` for a given composite subtype, nor the size or storage place for an object (including a component) of a given composite subtype, unless the constraints on the subtype and its composite subcomponents (if any) are all static constraints.

Followed. Size Clauses are not permitted on non-static components, as described above.

An aliased component, or a component whose type is by-reference, should always be allocated at an addressable location.

Followed.

13.2(6-8): Packed Types

If a type is packed, then the implementation should try to minimize storage allocated to objects of the type, possibly at the expense of speed of accessing components, subject to reasonable complexity in addressing calculations.

The recommended level of support pragma `Pack` is:
For a packed record type, the components should be packed as tightly as possible subject to the Sizes of the component subtypes, and subject to any `record_representation_clause` that applies to the type; the implementation may, but need not, reorder components or cross aligned word boundaries to improve the packing. A component whose `Size` is greater than the word size may be allocated an integral number of words.

Followed. Tight packing of arrays is supported for all component sizes up to 64-bits.

An implementation should support Address clauses for imported subprograms.

Followed.

13.3(14-19): Address Clauses

For an array *X*, `X'Address` should point at the first component of the array, and not at the array bounds.

Followed.

The recommended level of support for the `Address` attribute is:
`X'Address` should produce a useful result if X is an object that is aliased or of a by-reference type, or is an entity whose `Address` has been specified.

Followed. A valid address will be produced even if none of those conditions have been met. If necessary, the object is forced into memory to ensure the address is valid.

An implementation should support `Address` clauses for imported subprograms.

Followed.

Objects (including subcomponents) that are aliased or of a by-reference type should be allocated on storage element boundaries.

Followed.

If the `Address` of an object is specified, or it is imported or exported, then the implementation should not perform optimizations based on assumptions of no aliases.

Followed.

13.3(29-35): Alignment Clauses

The recommended level of support for the `Alignment` attribute for subtypes is:
An implementation should support specified Alignments that are factors and multiples of the number of storage elements per word, subject to the following:

Followed.

An implementation need not support specified `Alignment`s for combinations of `Size`s and `Alignment`s that cannot be easily loaded and stored by available machine instructions.

Followed.

An implementation need not support specified `Alignment`s that are greater than the maximum `Alignment` the implementation ever returns by default.

Followed.

The recommended level of support for the `Alignment` attribute for objects is:
Same as above, for subtypes, but in addition:

Followed.

For stand-alone library-level objects of statically constrained subtypes, the implementation should support all `Alignment`s supported by the target linker. For example, page alignment is likely to be supported for such objects, but not for subtypes.

Followed.

13.3(42-43): Size Clauses

The recommended level of support for the `Size` attribute of objects is:
A `Size` clause should be supported for an object if the specified `Size` is at least as large as its subtype's `Size`, and corresponds to a size in storage elements that is a multiple of the object's `Alignment` (if the `Alignment` is nonzero).

Followed.

13.3(50-56): Size Clauses

If the `Size` of a subtype is specified, and allows for efficient independent addressability (see 9.10) on the target architecture, then the `Size` of the following objects of the subtype should equal the `Size` of the subtype:
Aliased objects (including components).

Followed.

`Size` clause on a composite subtype should not affect the internal layout of components.

Followed.

The recommended level of support for the `Size` attribute of subtypes is:

The `Size` (if not specified) of a static discrete or fixed point subtype should be the number of bits needed to represent each value belonging to the subtype using an unbiased representation, leaving space for a sign bit only if the subtype contains negative values. If such a subtype is a first subtype, then an implementation should support a specified `Size` for it that reflects this representation.

Followed.

For a subtype implemented with levels of indirection, the `Size` should include the size of the pointers, but not the size of what they point at.

Followed.

13.3(71-73): Component Size Clauses

The recommended level of support for the `Component_Size` attribute is:

An implementation need not support specified `Component_Sizes` that are less than the `Size` of the component subtype.

Followed.

An implementation should support specified `Component_Sizes` that are factors and multiples of the word size. For such `Component_Sizes`, the array should contain no gaps between components. For other `Component_Sizes` (if supported), the array should contain no gaps between components when packing is also specified; the implementation should forbid this combination in cases where it cannot support a no-gaps representation.

Followed.

13.4(9-10): Enumeration Representation Clauses

The recommended level of support for enumeration representation clauses is:
An implementation need not support enumeration representation clauses for boolean types, but should at minimum support the internal codes in the range `System.Min_Int.System.Max_Int`.

Followed.

13.5.1(17-22): Record Representation Clauses

The recommended level of support for
`record_representation_clauses` is:
An implementation should support storage places that can be extracted with a load, mask, shift sequence of machine code, and set with a load, shift, mask, store sequence, given the available machine instructions and run-time model.

Followed.

A storage place should be supported if its size is equal to the `Size` of the component subtype, and it starts and ends on a boundary that obeys the `Alignment` of the component subtype.

Followed.

If the default bit ordering applies to the declaration of a given type, then for a component whose subtype's `Size` is less than the word size, any storage place that does not cross an aligned word boundary should be supported.

Followed.

An implementation may reserve a storage place for the tag field of a tagged type, and disallow other components from overlapping that place.

Followed. The storage place for the tag field is the beginning of the tagged record, and its size is Address'Size. GNAT will reject an explicit component clause for the tag field.

An implementation need not support a `component_clause` for a component of an extension part if the storage place is not after the storage places of all components of the parent type, whether or not those storage places had been specified.

Followed. The above advice on record representation clauses is followed, and all mentioned features are implemented.

13.5.2(5): Storage Place Attributes

If a component is represented using some form of pointer (such as an offset) to the actual data of the component, and this data is contiguous with the rest of the object, then the storage place attributes should reflect the place of the actual data, not the pointer. If a component is allocated discontinuously from the rest of the object, then a warning should be generated upon reference to one of its storage place attributes.

Followed. There are no such components in GNAT.

13.5.3(7-8): Bit Ordering

The recommended level of support for the non-default bit ordering is:

If `Word_Size` = `Storage_Unit`, then the implementation should support the non-default bit ordering in addition to the default bit ordering.

Followed. Word size does not equal storage size in this implementation. Thus non-default bit ordering is not supported.

13.7(37): Address as Private

`Address` should be of a private type.

Followed.

13.7.1(16): Address Operations

> Operations in `System` and its children should reflect the target environment semantics as closely as is reasonable. For example, on most machines, it makes sense for address arithmetic to "wrap around". Operations that do not make sense should raise `Program_Error`.

Followed. Address arithmetic is modular arithmetic that wraps around. No operation raises `Program_Error`, since all operations make sense.

13.9(14-17): Unchecked Conversion

> The `Size` of an array object should not include its bounds; hence, the bounds should not be part of the converted data.

Followed.

> The implementation should not generate unnecessary run-time checks to ensure that the representation of S is a representation of the target type. It should take advantage of the permission to return by reference when possible. Restrictions on unchecked conversions should be avoided unless required by the target environment.

Followed. There are no restrictions on unchecked conversion. A warning is generated if the source and target types do not have the same size since the semantics in this case may be target dependent.

> The recommended level of support for unchecked conversions is:

> Unchecked conversions should be supported and should be reversible in the cases where this clause defines the result. To enable meaningful use of unchecked conversion, a contiguous representation should be used for elementary subtypes, for statically constrained array subtypes whose component subtype is one of the subtypes described in this paragraph, and for record subtypes without discriminants whose component subtypes are described in this paragraph.

Followed.

13.11(23-25): Implicit Heap Usage

> An implementation should document any cases in which it dynamically allocates heap storage for a purpose other than the evaluation of an allocator.

Followed, the only other points at which heap storage is dynamically allocated are as follows:

- At initial elaboration time, to allocate dynamically sized global objects.

- To allocate space for a task when a task is created.
- To extend the secondary stack dynamically when needed. The secondary stack is used for returning variable length results.

A default (implementation-provided) storage pool for an access-to-constant type should not have overhead to support deallocation of individual objects.

Followed.

A storage pool for an anonymous access type should be created at the point of an allocator for the type, and be reclaimed when the designated object becomes inaccessible.

Followed.

13.11.2(17): Unchecked De-allocation

For a standard storage pool, `Free` should actually reclaim the storage.

Followed.

13.13.2(17): Stream Oriented Attributes

If a stream element is the same size as a storage element, then the normal in-memory representation should be used by `Read` and `Write` for scalar objects. Otherwise, `Read` and `Write` should use the smallest number of stream elements needed to represent all values in the base range of the scalar type.

Followed. In particular, the interpretation chosen is that of AI-195, which specifies that the size to be used is that of the first subtype.

A.1(52): Implementation Advice

If an implementation provides additional named predefined integer types, then the names should end with '`Integer`' as in '`Long_Integer`'. If an implementation provides additional named predefined floating point types, then the names should end with '`Float`' as in '`Long_Float`'.

Followed.

A.3.2(49): `Ada.Characters.Handling`

If an implementation provides a localized definition of `Character` or `Wide_Character`, then the effects of the subprograms in `Characters.Handling` should reflect the localizations. See also 3.5.2.

Followed. GNAT provides no such localized definitions.

A.4.4(106): Bounded-Length String Handling

> Bounded string objects should not be implemented by implicit pointers and dynamic allocation.

Followed. No implicit pointers or dynamic allocation are used.

A.5.2(46-47): Random Number Generation

> Any storage associated with an object of type `Generator` should be reclaimed on exit from the scope of the object.

Followed.

> If the generator period is sufficiently long in relation to the number of distinct initiator values, then each possible value of `Initiator` passed to `Reset` should initiate a sequence of random numbers that does not, in a practical sense, overlap the sequence initiated by any other value. If this is not possible, then the mapping between initiator values and generator states should be a rapidly varying function of the initiator value.

Followed. The generator period is sufficiently long for the first condition here to hold true.

A.10.7(23): `Get_Immediate`

> The `Get_Immediate` procedures should be implemented with unbuffered input. For a device such as a keyboard, input should be *available* if a key has already been typed, whereas for a disk file, input should always be available except at end of file. For a file associated with a keyboard-like device, any line-editing features of the underlying operating system should be disabled during the execution of `Get_Immediate`.

Followed.

B.1(39-41): Pragma `Export`

> If an implementation supports pragma `Export` to a given language, then it should also allow the main subprogram to be written in that language. It should support some mechanism for invoking the elaboration of the Ada library units included in the system, and for invoking the finalization of the environment task. On typical systems, the recommended mechanism is to provide two subprograms whose link names are `adainit` and `adafinal`. `adainit` should contain the elaboration code for library units. `adafinal` should contain the finalization code. These subprograms should have no effect the second and subsequent time they are called.

Followed.

Automatic elaboration of pre-elaborated packages should be provided when pragma `Export` is supported.

Followed when the main program is in Ada. If the main program is in a foreign language, then `adainit` must be called to elaborate pre-elaborated packages.

For each supported convention L other than `Intrinsic`, an implementation should support `Import` and `Export` pragmas for objects of L-compatible types and for subprograms, and pragma `Convention` for L-eligible types and for subprograms, presuming the other language has corresponding features. Pragma `Convention` need not be supported for scalar types.

Followed.

B.2(12-13): Package `Interfaces`

For each implementation-defined convention identifier, there should be a child package of package Interfaces with the corresponding name. This package should contain any declarations that would be useful for interfacing to the language (implementation) represented by the convention. Any declarations useful for interfacing to any language on the given hardware architecture should be provided directly in `Interfaces`.

Followed. An additional package not defined in the Ada 95 Reference Manual is `Interfaces.CPP`, used for interfacing to C++.

An implementation supporting an interface to C, COBOL, or Fortran should provide the corresponding package or packages described in the following clauses.

Followed. GNAT provides all the packages described in this section.

B.3(63-71): Interfacing with C

An implementation should support the following interface correspondences between Ada and C.

Followed.

An Ada procedure corresponds to a void-returning C function.

Followed.

An Ada function corresponds to a non-void C function.

Followed.

An Ada `in` scalar parameter is passed as a scalar argument to a C function.

Followed.

An Ada `in` parameter of an access-to-object type with designated type T is passed as a `t*` argument to a C function, where t is the C type corresponding to the Ada type T.

Followed.

An Ada access T parameter, or an Ada `out` or `in out` parameter of an elementary type T, is passed as a `t*` argument to a C function, where t is the C type corresponding to the Ada type T. In the case of an elementary `out` or `in out` parameter, a pointer to a temporary copy is used to preserve by-copy semantics.

Followed.

An Ada parameter of a record type T, of any mode, is passed as a `t*` argument to a C function, where t is the C structure corresponding to the Ada type T.

Followed. This convention may be overridden by the use of the C_Pass_By_Copy pragma, or Convention, or by explicitly specifying the mechanism for a given call using an extended import or export pragma.

An Ada parameter of an array type with component type T, of any mode, is passed as a `t*` argument to a C function, where t is the C type corresponding to the Ada type T.

Followed.

An Ada parameter of an access-to-subprogram type is passed as a pointer to a C function whose prototype corresponds to the designated subprogram's specification.

Followed.

B.4(95-98): Interfacing with COBOL

An Ada implementation should support the following interface correspondences between Ada and COBOL.

Followed.

An Ada access T parameter is passed as a '`BY REFERENCE`' data item of the COBOL type corresponding to T.

Followed.

An Ada in scalar parameter is passed as a '`BY CONTENT`' data item of the corresponding COBOL type.

Followed.

Any other Ada parameter is passed as a '`BY REFERENCE`' data item of the COBOL type corresponding to the Ada parameter type; for scalars, a local copy is used if necessary to ensure by-copy semantics.

Followed.

B.5(22-26): Interfacing with Fortran

An Ada implementation should support the following interface correspondences between Ada and Fortran:

Followed.

An Ada procedure corresponds to a Fortran subroutine.

Followed.

An Ada function corresponds to a Fortran function.

Followed.

An Ada parameter of an elementary, array, or record type T is passed as a T argument to a Fortran procedure, where T is the Fortran type corresponding to the Ada type T, and where the INTENT attribute of the corresponding dummy argument matches the Ada formal parameter mode; the Fortran implementation's parameter passing conventions are used. For elementary types, a local copy is used if necessary to ensure by-copy semantics.

Followed.

An Ada parameter of an access-to-subprogram type is passed as a reference to a Fortran procedure whose interface corresponds to the designated subprogram's specification.

Followed.

C.1(3-5): Access to Machine Operations

> The machine code or intrinsic support should allow access to all operations normally available to assembly language programmers for the target environment, including privileged instructions, if any.

Followed.

> The interfacing pragmas (see Annex B) should support interface to assembler; the default assembler should be associated with the convention identifier `Assembler`.

Followed.

> If an entity is exported to assembly language, then the implementation should allocate it at an addressable location, and should ensure that it is retained by the linking process, even if not otherwise referenced from the Ada code. The implementation should assume that any call to a machine code or assembler subprogram is allowed to read or update every object that is specified as exported.

Followed.

C.1(10-16): Access to Machine Operations

> The implementation should ensure that little or no overhead is associated with calling intrinsic and machine-code subprograms.

Followed for both intrinsics and machine-code subprograms.

> It is recommended that intrinsic subprograms be provided for convenient access to any machine operations that provide special capabilities or efficiency and that are not otherwise available through the language constructs.

Followed. A full set of machine operation intrinsic subprograms is provided.

> Atomic read-modify-write operations—e.g., test and set, compare and swap, decrement and test, enqueue/dequeue.

Followed on any target supporting such operations.

> Standard numeric functions—e.g., sin, log.

Followed on any target supporting such operations.

String manipulation operations—e.g., translate and test.

Followed on any target supporting such operations.

Vector operations—e.g., compare vector against thresholds.

Followed on any target supporting such operations.

Direct operations on I/O ports.

Followed on any target supporting such operations.

C.3(28): Interrupt Support

If the `Ceiling_Locking` policy is not in effect, the implementation should provide means for the application to specify which interrupts are to be blocked during protected actions, if the underlying system allows for a finer-grain control of interrupt blocking.

Followed. The underlying system does not allow for finer-grain control of interrupt blocking.

C.3.1(20-21): Protected Procedure Handlers

Whenever possible, the implementation should allow interrupt handlers to be called directly by the hardware.

Followed on any target where the underlying operating system permits such direct calls.

Whenever practical, violations of any implementation-defined restrictions should be detected before run time.

Followed. Compile time warnings are given when possible.

C.3.2(25): Package `Interrupts`

If implementation-defined forms of interrupt handler procedures are supported, such as protected procedures with parameters, then for each such form of a handler, a type analogous to `Parameterless_Handler` should be specified in a child package of `Interrupts`, with the same operations as in the predefined package Interrupts.

Followed.

C.4(14): Pre-elaboration Requirements

> It is recommended that pre-elaborated packages be implemented in such a way that there should be little or no code executed at run time for the elaboration of entities not already covered by the Implementation Requirements.

Followed. Executable code is generated in some cases, e.g. loops to initialize large arrays.

C.5(8): Pragma `Discard_Names`

> If the pragma applies to an entity, then the implementation should reduce the amount of storage used for storing names associated with that entity.

Followed.

C.7.2(30): The Package Task_Attributes

> Some implementations are targeted to domains in which memory use at run time must be completely deterministic. For such implementations, it is recommended that the storage for task attributes will be pre-allocated statically and not from the heap. This can be accomplished by either placing restrictions on the number and the size of the task's attributes, or by using the pre-allocated storage for the first N attribute objects, and the heap for the others. In the latter case, N should be documented.

Not followed. This implementation is not targeted to such a domain.

D.3(17): Locking Policies

> The implementation should use names that end with '`_Locking`' for locking policies defined by the implementation.

Followed. A single implementation-defined locking policy is defined, whose name (`Inheritance_Locking`) follows this suggestion.

D.4(16): Entry Queuing Policies

> Names that end with '`_Queuing`' should be used for all implementation-defined queuing policies.

Followed. No such implementation-defined queueing policies exist.

D.6(9-10): Preemptive Abort

> Even though the `abort_statement` is included in the list of potentially blocking operations (see 9.5.1), it is recommended that this statement be implemented in a way that never requires the task executing the `abort_statement` to block.

Followed.

> On a multi-processor, the delay associated with aborting a task on another processor should be bounded; the implementation should use periodic polling, if necessary, to achieve this.

Followed.

D.7(21): Tasking Restrictions

> When feasible, the implementation should take advantage of the specified restrictions to produce a more efficient implementation.

GNAT currently takes advantage of these restrictions by providing an optimized run time when the Ravenscar profile and the GNAT restricted run time set of restrictions are specified. See pragma `Ravenscar` and pragma `Restricted_Run_Time` for more details.

D.8(47-49): Monotonic Time

> When appropriate, implementations should provide configuration mechanisms to change the value of `Tick`.

Such configuration mechanisms are not appropriate to this implementation and are thus not supported.

> It is recommended that `Calendar.Clock` and `Real_Time.Clock` be implemented as transformations of the same time base.

Followed.

> It is recommended that the *best* time base which exists in the underlying system be available to the application through `Clock`. *Best* may mean highest accuracy or largest range.

Followed.

E.5(28-29): Partition Communication Subsystem

> Whenever possible, the PCS on the called partition should allow for multiple tasks to call the RPC-receiver with different messages and should allow them to block until the corresponding subprogram body returns.

Followed by GLADE, a separately supplied PCS that can be used with GNAT.

> The `Write` operation on a stream of type `Params_Stream_Type` should raise `Storage_Error` if it runs out of space trying to write the `Item` into the stream.

Followed by GLADE, a separately supplied PCS that can be used with GNAT.

F(7): COBOL Support

> If COBOL (respectively, C) is widely supported in the target environment, implementations supporting the Information Systems Annex should provide the child package `Interfaces.COBOL` (respectively, `Interfaces.C`) specified in Annex B and should support a `convention_identifier` of COBOL (respectively, C) in the interfacing pragmas (see Annex B), thus allowing Ada programs to interface with programs written in that language.

Followed.

F.1(2): Decimal Radix Support

> Packed decimal should be used as the internal representation for objects of subtype S when S'Machine_Radix = 10.

Not followed. GNAT ignores S'Machine_Radix and always uses binary representations.

G: Numerics

> If Fortran (respectively, C) is widely supported in the target environment, implementations supporting the Numerics Annex should provide the child package `Interfaces.Fortran` (respectively, `Interfaces.C`) specified in Annex B and should support a `convention_identifier` of Fortran (respectively, C) in the interfacing pragmas (see Annex B), thus allowing Ada programs to interface with programs written in that language.

Followed.

G.1.1(56-58): Complex Types

> Because the usual mathematical meaning of multiplication of a complex operand and a real operand is that of the scaling of both components of the former by the latter, an implementation should not perform this operation by first promoting the real operand to complex type and then performing a full complex multiplication. In systems that, in the future, support an Ada binding to IEC 559:1989, the latter technique will not generate the required result when one of the components of the complex operand is infinite. (Explicit multiplication of the infinite component by the zero component obtained during promotion yields a NaN that propagates into the final result.) Analogous advice applies in the case of multiplication of a complex operand and a pure-imaginary operand, and in the case of division of a complex operand by a real or pure-imaginary operand.

Not followed.

Similarly, because the usual mathematical meaning of addition of a complex operand and a real operand is that the imaginary operand remains unchanged, an implementation should not perform this operation by first promoting the real operand to complex type and then performing a full complex addition. In implementations in which the `Signed_Zeros` attribute of the component type is `True` (and which therefore conform to IEC 559:1989 in regard to the handling of the sign of zero in predefined arithmetic operations), the latter technique will not generate the required result when the imaginary component of the complex operand is a negatively signed zero. (Explicit addition of the negative zero to the zero obtained during promotion yields a positive zero.) Analogous advice applies in the case of addition of a complex operand and a pure-imaginary operand, and in the case of subtraction of a complex operand and a real or pure-imaginary operand.

Not followed.

Implementations in which `Real'Signed_Zeros` is `True` should attempt to provide a rational treatment of the signs of zero results and result components. As one example, the result of the `Argument` function should have the sign of the imaginary component of the parameter `X` when the point represented by that parameter lies on the positive real axis; as another, the sign of the imaginary component of the `Compose_From_Polar` function should be the same as (respectively, the opposite of) that of the `Argument` parameter when that parameter has a value of zero and the `Modulus` parameter has a nonnegative (respectively, negative) value.

Followed.

G.1.2(49): Complex Elementary Functions

Implementations in which `Complex_Types.Real'Signed_Zeros` is `True` should attempt to provide a rational treatment of the signs of zero results and result components. For example, many of the complex elementary functions have components that are odd functions of one of the parameter components; in these cases, the result component should have the sign of the parameter component at the origin. Other complex elementary functions have zero components whose sign is opposite that of a parameter component at the origin, or is always positive or always negative.

Followed.

G.2.4(19): Accuracy Requirements

The versions of the forward trigonometric functions without a `Cycle` parameter should not be implemented by calling the corresponding version with a `Cycle` parameter of `2.0*Numerics.Pi`, since this will not provide the required accuracy in some portions of the domain. For the same reason, the version of `Log` without a `Base` parameter should not be implemented by calling the corresponding version with a `Base` parameter of `Numerics.e`.

Followed.

G.2.6(15): Complex Arithmetic Accuracy

> The version of the `Compose_From_Polar` function without a `Cycle` parameter should not be implemented by calling the corresponding version with a `Cycle` parameter of `2.0*Numerics.Pi`, since this will not provide the required accuracy in some portions of the domain.

Followed.

4 Implementation Defined Characteristics

In addition to the implementation dependent pragmas and attributes, and the implementation advice, there are a number of other features of Ada 95 that are potentially implementation dependent. These are mentioned throughout the Ada 95 Reference Manual, and are summarized in annex M.

A requirement for conforming Ada compilers is that they provide documentation describing how the implementation deals with each of these issues. In this chapter, you will find each point in annex M listed followed by a description in italic font of how GNAT handles the implementation dependence.

You can use this chapter as a guide to minimizing implementation dependent features in your programs if portability to other compilers and other operating systems is an important consideration. The numbers in each section below correspond to the paragraph number in the Ada 95 Reference Manual.

2. Whether or not each recommendation given in Implementation Advice is followed. See 1.1.2(37).

See Chapter 3 [Implementation Advice], page 51.

3. Capacity limitations of the implementation. See 1.1.3(3).

The complexity of programs that can be processed is limited only by the total amount of available virtual memory, and disk space for the generated object files.

4. Variations from the standard that are impractical to avoid given the implementation's execution environment. See 1.1.3(6).

There are no variations from the standard.

5. Which `code_statement`s cause external interactions. See 1.1.3(10).

Any `code_statement` can potentially cause external interactions.

6. The coded representation for the text of an Ada program. See 2.1(4).

See separate section on source representation.

7. The control functions allowed in comments. See 2.1(14).

See separate section on source representation.

8. The representation for an end of line. See 2.2(2).

See separate section on source representation.

9. Maximum supported line length and lexical element length. See 2.2(15).

The maximum line length is 255 characters an the maximum length of a lexical element is also 255 characters.

10. Implementation defined pragmas. See 2.8(14).

See Chapter 1 [Implementation Defined Pragmas], page 3.

11. Effect of pragma `Optimize`. See 2.8(27).

Pragma `Optimize`, if given with a `Time` or `Space` parameter, checks that the optimization flag is set, and aborts if it is not.

12. The sequence of characters of the value returned by *S*`'Image` when some of the graphic characters of *S*`'Wide_Image` are not defined in `Character`. See 3.5(37).

The sequence of characters is as defined by the wide character encoding method used for the source. See section on source representation for further details.

13. The predefined integer types declared in `Standard`. See 3.5.4(25).

`Short_Short_Integer`
: 8 bit signed

`Short_Integer`
: (Short) 16 bit signed

`Integer` 32 bit signed

`Long_Integer`
: 64 bit signed (Alpha OpenVMS only) 32 bit signed (all other targets)

`Long_Long_Integer`
: 64 bit signed

14. Any nonstandard integer types and the operators defined for them. See 3.5.4(26).

There are no nonstandard integer types.

15. Any nonstandard real types and the operators defined for them. See 3.5.6(8).

There are no nonstandard real types.

16. What combinations of requested decimal precision and range are supported for floating point types. See 3.5.7(7).

The precision and range is as defined by the IEEE standard.

17. The predefined floating point types declared in `Standard`. See 3.5.7(16).

`Short_Float`
: 32 bit IEEE short

`Float`
: (Short) 32 bit IEEE short

`Long_Float`
: 64 bit IEEE long

`Long_Long_Float`
: 64 bit IEEE long (80 bit IEEE long on x86 processors)

18. The small of an ordinary fixed point type. See 3.5.9(8).

`Fine_Delta` is 2**(−63)

19. What combinations of small, range, and digits are supported for fixed point types. See 3.5.9(10).

Any combinations are permitted that do not result in a small less than `Fine_Delta` and do not result in a mantissa larger than 63 bits. If the mantissa is larger than 53 bits on machines where Long_Long_Float is 64 bits (true of all architectures except ia32), then the output from Text_IO is accurate to only 53 bits, rather than the full mantissa. This is because floating-point conversions are used to convert fixed point.

20. The result of `Tags.Expanded_Name` for types declared within an unnamed `block_statement`. See 3.9(10).

Block numbers of the form `Bnnn`, where *nnn* is a decimal integer are allocated.

21. Implementation-defined attributes. See 4.1.4(12).

See Chapter 2 [Implementation Defined Attributes], page 41.

22. Any implementation-defined time types. See 9.6(6).

There are no implementation-defined time types.

23. The time base associated with relative delays.

See 9.6(20). The time base used is that provided by the C library function `gettimeofday`.

24. The time base of the type `Calendar.Time`. See 9.6(23).

The time base used is that provided by the C library function `gettimeofday`.

25. The time zone used for package `Calendar` operations. See 9.6(24).

The time zone used by package `Calendar` is the current system time zone setting for local time, as accessed by the C library function `localtime`.

26. Any limit on `delay_until_statements` of `select_statements`. See 9.6(29).

There are no such limits.

27. Whether or not two non overlapping parts of a composite object are independently addressable, in the case where packing, record layout, or `Component_Size` is specified for the object. See 9.10(1).

Separate components are independently addressable if they do not share overlapping storage units.

28. The representation for a compilation. See 10.1(2).

A compilation is represented by a sequence of files presented to the compiler in a single invocation of the `gcc` command.

29. Any restrictions on compilations that contain multiple compilation_units. See 10.1(4).

No single file can contain more than one compilation unit, but any sequence of files can be presented to the compiler as a single compilation.

30. The mechanisms for creating an environment and for adding and replacing compilation units. See 10.1.4(3).

See separate section on compilation model.

31. The manner of explicitly assigning library units to a partition. See 10.2(2).

If a unit contains an Ada main program, then the Ada units for the partition are determined by recursive application of the rules in the Ada Reference Manual section 10.2(2-6). In other words, the Ada units will be those that are needed by the main program, and then this definition of need is applied recursively to those units, and the partition contains the transitive closure determined by this relationship. In short, all the necessary units are included, with no need to explicitly specify the list. If additional units are required, e.g. by foreign language units, then all units must be mentioned in the context clause of one of the needed Ada units.

If the partition contains no main program, or if the main program is in a language other than Ada, then GNAT provides the binder options `-z` and `-n` respectively, and in this case a list of units can be explicitly supplied to the binder for inclusion in the partition (all units needed by these units will also be included automatically). For full details on the use of these options, refer to the *GNAT User's Guide* sections on Binding and Linking.

32. The implementation-defined means, if any, of specifying which compilation units are needed by a given compilation unit. See 10.2(2).

The units needed by a given compilation unit are as defined in the Ada Reference Manual section 10.2(2-6). There are no implementation-defined pragmas or other implementation-defined means for specifying needed units.

33. The manner of designating the main subprogram of a partition. See 10.2(7).

The main program is designated by providing the name of the corresponding '`ALI`' file as the input parameter to the binder.

34. The order of elaboration of `library_items`. See 10.2(18).

The first constraint on ordering is that it meets the requirements of chapter 10 of the Ada 95 Reference Manual. This still leaves some implementation dependent choices, which are resolved by first elaborating bodies as early as possible (i.e. in preference to specs where there is a choice), and second by evaluating the immediate with clauses of a unit to determine the probably best choice, and third by elaborating in alphabetical order of unit names where a choice still remains.

35. Parameter passing and function return for the main subprogram. See 10.2(21).

The main program has no parameters. It may be a procedure, or a function returning an integer type. In the latter case, the returned integer value is the return code of the program.

36. The mechanisms for building and running partitions. See 10.2(24).

GNAT itself supports programs with only a single partition. The GNATDIST tool provided with the GLADE package (which also includes an implementation of the PCS) provides a completely flexible method for building and running programs consisting of multiple partitions. See the separate GLADE manual for details.

37. The details of program execution, including program termination. See 10.2(25).

See separate section on compilation model.

38. The semantics of any non-active partitions supported by the implementation. See 10.2(28).

Passive partitions are supported on targets where shared memory is provided by the operating system. See the GLADE reference manual for further details.

39. The information returned by `Exception_Message`. See 11.4.1(10).

Exception message returns the null string unless a specific message has been passed by the program.

40. The result of `Exceptions.Exception_Name` for types declared within an unnamed `block_statement`. See 11.4.1(12).

Blocks have implementation defined names of the form `Bnnn` where *nnn* is an integer.

41. The information returned by `Exception_Information`. See 11.4.1(13).

`Exception_Information` returns a string in the following format:

```
Exception_Name: nnnnn
Message: mmmmm
PID: ppp
Call stack traceback locations:
0xhhhh 0xhhhh 0xhhhh ... 0xhhh
```

where

- `nnnn` is the fully qualified name of the exception in all upper case letters. This line is always present.
- `mmmm` is the message (this line present only if message is non-null)

- ppp is the Process Id value as a decimal integer (this line is present only if the Process Id is non-zero). Currently we are not making use of this field.
- The Call stack traceback locations line and the following values are present only if at least one traceback location was recorded. The values are given in C style format, with lower case letters for a-f, and only as many digits present as are necessary.

The line terminator sequence at the end of each line, including the last line is a single LF character (16#0A#).

42. Implementation-defined check names. See 11.5(27).

No implementation-defined check names are supported.

43. The interpretation of each aspect of representation. See 13.1(20).

See separate section on data representations.

44. Any restrictions placed upon representation items. See 13.1(20).

See separate section on data representations.

45. The meaning of Size for indefinite subtypes. See 13.3(48).

Size for an indefinite subtype is the maximum possible size, except that for the case of a subprogram parameter, the size of the parameter object is the actual size.

46. The default external representation for a type tag. See 13.3(75).

The default external representation for a type tag is the fully expanded name of the type in upper case letters.

47. What determines whether a compilation unit is the same in two different partitions. See 13.3(76).

A compilation unit is the same in two different partitions if and only if it derives from the same source file.

48. Implementation-defined components. See 13.5.1(15).

The only implementation defined component is the tag for a tagged type, which contains a pointer to the dispatching table.

49. If `Word_Size` = `Storage_Unit`, the default bit ordering. See 13.5.3(5).

`Word_Size` (32) is not the same as `Storage_Unit` (8) for this implementation, so no non-default bit ordering is supported. The default bit ordering corresponds to the natural endianness of the target architecture.

50. The contents of the visible part of package `System` and its language-defined children. See 13.7(2).

See the definition of these packages in files '`system.ads`' and '`s-stoele.ads`'.

51. The contents of the visible part of package `System.Machine_Code`, and the meaning of `code_statements`. See 13.8(7).

See the definition and documentation in file '`s-maccod.ads`'.

52. The effect of unchecked conversion. See 13.9(11).

Unchecked conversion between types of the same size and results in an uninterpreted transmission of the bits from one type to the other. If the types are of unequal sizes, then in the case of discrete types, a shorter source is first zero or sign extended as necessary, and a shorter target is simply truncated on the left. For all non-discrete types, the source is first copied if necessary to ensure that the alignment requirements of the target are met, then a pointer is constructed to the source value, and the result is obtained by dereferencing this pointer after converting it to be a pointer to the target type.

53. The manner of choosing a storage pool for an access type when `Storage_Pool` is not specified for the type. See 13.11(17).

There are 3 different standard pools used by the compiler when `Storage_Pool` is not specified depending whether the type is local to a subprogram or defined at the library level and whether `Storage_Size`is specified or not. See documentation in the runtime library units `System.Pool_Global`, `System.Pool_Size` and `System.Pool_Local` in files '`s-poosiz.ads`', '`s-pooglo.ads`' and '`s-pooloc.ads`' for full details on the default pools used.

54. Whether or not the implementation provides user-accessible names for the standard pool type(s). See 13.11(17).

See documentation in the sources of the run time mentioned in paragraph **53** . All these pools are accessible by means of `with`'ing these units.

55. The meaning of `Storage_Size`. See 13.11(18).

`Storage_Size` is measured in storage units, and refers to the total space available for an access type collection, or to the primary stack space for a task.

56. Implementation-defined aspects of storage pools. See 13.11(22).

See documentation in the sources of the run time mentioned in paragraph **53** for details on GNAT-defined aspects of storage pools.

57. The set of restrictions allowed in a pragma `Restrictions`. See 13.12(7).

All RM defined Restriction identifiers are implemented. The following additional restriction identifiers are provided. There are two separate lists of implementation dependent restriction identifiers. The first set requires consistency throughout a partition (in other words, if the restriction identifier is used for any compilation unit in the partition, then all compilation units in the partition must obey the restriction.

`Boolean_Entry_Barriers`
: This restriction ensures at compile time that barriers in entry declarations for protected types are restricted to references to simple boolean variables defined in the private part of the protected type. No other form of entry barriers is permitted. This is one of the restrictions of the Ravenscar profile for limited tasking (see also pragma `Ravenscar`).

`Max_Entry_Queue_Depth => Expr`
: This restriction is a declaration that any protected entry compiled in the scope of the restriction has at most the specified number of tasks waiting on the entry at any one time, and so no queue is required. This restriction is not checked at compile time. A program execution is erroneous if an attempt is made to queue more than the specified number of tasks on such an entry.

`No_Calendar`
: This restriction ensures at compile time that there is no implicit or explicit dependence on the package `Ada.Calendar`.

`No_Dynamic_Interrupts`
: This restriction ensures at compile time that there is no attempt to dynamically associate interrupts. Only static association is allowed.

`No_Enumeration_Maps`
: This restriction ensures at compile time that no operations requiring enumeration maps are used (that is Image and Value attributes applied to enumeration types).

`No_Entry_Calls_In_Elaboration_Code`
: This restriction ensures at compile time that no task or protected entry calls are made during elaboration code. As a result of the use of this restriction, the

compiler can assume that no code past an accept statement in a task can be executed at elaboration time.

`No_Exception_Handlers`
: This restriction ensures at compile time that there are no explicit exception handlers.

`No_Implicit_Conditionals`
: This restriction ensures that the generated code does not contain any implicit conditionals, either by modifying the generated code where possible, or by rejecting any construct that would otherwise generate an implicit conditional. The details and use of this restriction are described in more detail in the High Integrity product documentation.

`No_Implicit_Loops`
: This restriction ensures that the generated code does not contain any implicit `for` loops, either by modifying the generated code where possible, or by rejecting any construct that would otherwise generate an implicit `for` loop. The details and use of this restriction are described in more detail in the High Integrity product documentation.

`No_Local_Protected_Objects`
: This restriction ensures at compile time that protected objects are only declared at the library level.

`No_Protected_Type_Allocators`
: This restriction ensures at compile time that there are no allocator expressions that attempt to allocate protected objects.

`No_Secondary_Stack`
: This restriction ensures at compile time that the generated code does not contain any reference to the secondary stack. The secondary stack is used to implement functions returning unconstrained objects (arrays or records) on some targets. The details and use of this restriction are described in more detail in the High Integrity product documentation.

`No_Select_Statements`
: This restriction ensures at compile time no select statements of any kind are permitted, that is the keyword `select` may not appear. This is one of the restrictions of the Ravenscar profile for limited tasking (see also pragma `Ravenscar`).

`No_Standard_Storage_Pools`
: This restriction ensures at compile time that no access types use the standard default storage pool. Any access type declared must have an explicit Storage_Pool attribute defined specifying a user-defined storage pool.

`No_Streams`
: This restriction ensures at compile time that there are no implicit or explicit dependencies on the package `Ada.Streams`.

`No_Task_Attributes`
: This restriction ensures at compile time that there are no implicit or explicit dependencies on the package `Ada.Task_Attributes`.

`No_Task_Termination`
: This restriction ensures at compile time that no terminate alternatives appear in any task body.

`No_Tasking`
: This restriction prevents the declaration of tasks or task types throughout the partition. It is similar in effect to the use of `Max_Tasks => 0` except that violations are caught at compile time and cause an error message to be output either by the compiler or binder.

`No_Wide_Characters`
: This restriction ensures at compile time that no uses of the types `Wide_Character` or `Wide_String` appear, and that no wide character literals appear in the program (that is literals representing characters not in type `Character`.

`Static_Priorities`
: This restriction ensures at compile time that all priority expressions are static, and that there are no dependencies on the package `Ada.Dynamic_Priorities`.

`Static_Storage_Size`
: This restriction ensures at compile time that any expression appearing in a Storage_Size pragma or attribute definition clause is static.

The second set of implementation dependent restriction identifiers does not require partition-wide consistency. The restriction may be enforced for a single compilation unit without any effect on any of the other compilation units in the partition.

`No_Elaboration_Code`
: This restriction ensures at compile time that no elaboration code is generated. Note that this is not the same condition as is enforced by pragma `Preelaborate`. There are cases in which pragma `Preelaborate` still permits code to be generated (e.g. code to initialize a large array to all zeroes), and there are cases of units which do not meet the requirements for pragma `Preelaborate`, but for which no elaboration code is generated. Generally, it is the case that preelaborable units will meet the restrictions, with the exception of large aggregates initialized with an others_clause, and exception declarations (which generate calls to a run-time registry procedure). Note that this restriction is enforced on a unit by unit basis, it need not be obeyed consistently throughout a partition.

`No_Entry_Queue`
: This restriction is a declaration that any protected entry compiled in the scope of the restriction has at most one task waiting on the entry at any one time, and so no queue is required. This restriction is not checked at compile time. A program execution is erroneous if an attempt is made to queue a second task on such an entry.

`No_Implementation_Attributes`
: This restriction checks at compile time that no GNAT-defined attributes are present. With this restriction, the only attributes that can be used are those defined in the Ada 95 Reference Manual.

`No_Implementation_Pragmas`
: This restriction checks at compile time that no GNAT-defined pragmas are present. With this restriction, the only pragmas that can be used are those defined in the Ada 95 Reference Manual.

`No_Implementation_Restrictions`
: This restriction checks at compile time that no GNAT-defined restriction identifiers (other than `No_Implementation_Restrictions` itself) are present. With this restriction, the only other restriction identifiers that can be used are those defined in the Ada 95 Reference Manual.

58. The consequences of violating limitations on `Restrictions` pragmas. See 13.12(9).

Restrictions that can be checked at compile time result in illegalities if violated. Currently there are no other consequences of violating restrictions.

59. The representation used by the `Read` and `Write` attributes of elementary types in terms of stream elements. See 13.13.2(9).

The representation is the in-memory representation of the base type of the type, using the number of bits corresponding to the *type*`'Size` value, and the natural ordering of the machine.

60. The names and characteristics of the numeric subtypes declared in the visible part of package `Standard`. See A.1(3).

See items describing the integer and floating-point types supported.

61. The accuracy actually achieved by the elementary functions. See A.5.1(1).

The elementary functions correspond to the functions available in the C library. Only fast math mode is implemented.

62. The sign of a zero result from some of the operators or functions in `Numerics.Generic_Elementary_Functions`, when `Float_Type'Signed_Zeros` is `True`. See A.5.1(46).

The sign of zeroes follows the requirements of the IEEE 754 standard on floating-point.

63. The value of `Numerics.Float_Random.Max_Image_Width`. See A.5.2(27).

Maximum image width is 649, see library file '`a-numran.ads`'.

64. The value of `Numerics.Discrete_Random.Max_Image_Width`. See A.5.2(27).

Maximum image width is 80, see library file '`a-nudira.ads`'.

65. The algorithms for random number generation. See A.5.2(32).

The algorithm is documented in the source files '`a-numran.ads`' and '`a-numran.adb`'.

66. The string representation of a random number generator's state. See A.5.2(38).

See the documentation contained in the file '`a-numran.adb`'.

67. The minimum time interval between calls to the time-dependent Reset procedure that are guaranteed to initiate different random number sequences. See A.5.2(45).

The minimum period between reset calls to guarantee distinct series of random numbers is one microsecond.

68. The values of the `Model_Mantissa`, `Model_Emin`, `Model_Epsilon`, `Model`, `Safe_First`, and `Safe_Last` attributes, if the Numerics Annex is not supported. See A.5.3(72).

See the source file '`ttypef.ads`' for the values of all numeric attributes.

69. Any implementation-defined characteristics of the input-output packages. See A.7(14).

There are no special implementation defined characteristics for these packages.

70. The value of `Buffer_Size` in `Storage_IO`. See A.9(10).

All type representations are contiguous, and the `Buffer_Size` is the value of *`type`*`'Size` rounded up to the next storage unit boundary.

71. External files for standard input, standard output, and standard error See A.10(5).

These files are mapped onto the files provided by the C streams libraries. See source file '`i-cstrea.ads`' for further details.

72. The accuracy of the value produced by `Put`. See A.10.9(36).

If more digits are requested in the output than are represented by the precision of the value, zeroes are output in the corresponding least significant digit positions.

73. The meaning of `Argument_Count`, `Argument`, and `Command_Name`. See A.15(1).

These are mapped onto the `argv` and `argc` parameters of the main program in the natural manner.

74. Implementation-defined convention names. See B.1(11).

The following convention names are supported

`Ada` Ada

`Assembler`
Assembly language

`Asm` Synonym for Assembler

`Assembly` Synonym for Assembler

`C` C

`C_Pass_By_Copy`
Allowed only for record types, like C, but also notes that record is to be passed by copy rather than reference.

`COBOL` COBOL

`CPP` C++

`Default` Treated the same as C

`External` Treated the same as C

`Fortran` Fortran

`Intrinsic`
For support of pragma `Import` with convention Intrinsic, see separate section on Intrinsic Subprograms.

`Stdcall` Stdcall (used for Windows implementations only). This convention correspond to the WINAPI (previously called Pascal convention) C/C++ convention under Windows. A function with this convention cleans the stack before exit.

`DLL` Synonym for Stdcall

`Win32` Synonym for Stdcall

`Stubbed` Stubbed is a special convention used to indicate that the body of the subprogram will be entirely ignored. Any call to the subprogram is converted into a raise of the `Program_Error` exception. If a pragma `Import` specifies convention `stubbed` then no body need be present at all. This convention is useful during development for the inclusion of subprograms whose body has not yet been written.

In addition, all otherwise unrecognized convention names are also treated as being synonymous with convention C. In all implementations except for VMS, use of such other names results in a warning. In VMS implementations, these names are accepted silently.

75. The meaning of link names. See B.1(36).

Link names are the actual names used by the linker.

76. The manner of choosing link names when neither the link name nor the address of an imported or exported entity is specified. See B.1(36).

The default linker name is that which would be assigned by the relevant external language, interpreting the Ada name as being in all lower case letters.

77. The effect of pragma `Linker_Options`. See B.1(37).

The string passed to `Linker_Options` is presented uninterpreted as an argument to the link command, unless it contains Ascii.NUL characters. NUL characters if they appear act as argument separators, so for example

```
pragma Linker_Options ("-labc" & ASCII.Nul & "-ldef");
```

causes two separate arguments `-labc` and `-ldef` to be passed to the linker. The order of linker options is preserved for a given unit. The final list of options passed to the linker is in reverse order of the elaboration order. For example, linker options fo a body always appear before the options from the corresponding package spec.

78. The contents of the visible part of package `Interfaces` and its language-defined descendants. See B.2(1).

See files with prefix '`i-`' in the distributed library.

79. Implementation-defined children of package `Interfaces`. The contents of the visible part of package `Interfaces`. See B.2(11).

See files with prefix '`i-`' in the distributed library.

80. The types `Floating`, `Long_Floating`, `Binary`, `Long_Binary`, `Decimal_ Element`, and `COBOL_Character`; and the initialization of the variables `Ada_To_COBOL` and `COBOL_To_Ada`, in `Interfaces.COBOL`. See B.4(50).

`Floating` Float

`Long_Floating`
(Floating) Long_Float

`Binary` Integer

`Long_Binary`
Long_Long_Integer

`Decimal_Element`
: Character

`COBOL_Character`
: Character

For initialization, see the file '`i-cobol.ads`' in the distributed library.

81. Support for access to machine instructions. See C.1(1).

See documentation in file '`s-maccod.ads`' in the distributed library.

82. Implementation-defined aspects of access to machine operations. See C.1(9).

See documentation in file '`s-maccod.ads`' in the distributed library.

83. Implementation-defined aspects of interrupts. See C.3(2).

Interrupts are mapped to signals or conditions as appropriate. See definition of unit `Ada.Interrupt_Names` in source file '`a-intnam.ads`' for details on the interrupts supported on a particular target.

84. Implementation-defined aspects of pre-elaboration. See C.4(13).

GNAT does not permit a partition to be restarted without reloading, except under control of the debugger.

85. The semantics of pragma `Discard_Names`. See C.5(7).

Pragma `Discard_Names` causes names of enumeration literals to be suppressed. In the presence of this pragma, the Image attribute provides the image of the Pos of the literal, and Value accepts Pos values.

86. The result of the `Task_Identification.Image` attribute. See C.7.1(7).

The result of this attribute is an 8-digit hexadecimal string representing the virtual address of the task control block.

87. The value of `Current_Task` when in a protected entry or interrupt handler. See C.7.1(17).

Protected entries or interrupt handlers can be executed by any convenient thread, so the value of `Current_Task` is undefined.

88. The effect of calling `Current_Task` from an entry body or interrupt handler. See C.7.1(19).

The effect of calling `Current_Task` from an entry body or interrupt handler is to return the identification of the task currently executing the code.

89. Implementation-defined aspects of `Task_Attributes`. See C.7.2(19).

There are no implementation-defined aspects of `Task_Attributes`.

90. Values of all `Metrics`. See D(2).

The metrics information for GNAT depends on the performance of the underlying operating system. The sources of the run-time for tasking implementation, together with the output from `-gnatG` can be used to determine the exact sequence of operating systems calls made to implement various tasking constructs. Together with appropriate information on the performance of the underlying operating system, on the exact target in use, this information can be used to determine the required metrics.

91. The declarations of `Any_Priority` and `Priority`. See D.1(11).

See declarations in file '`system.ads`'.

92. Implementation-defined execution resources. See D.1(15).

There are no implementation-defined execution resources.

93. Whether, on a multiprocessor, a task that is waiting for access to a protected object keeps its processor busy. See D.2.1(3).

On a multi-processor, a task that is waiting for access to a protected object does not keep its processor busy.

94. The affect of implementation defined execution resources on task dispatching. See D.2.1(9).

Tasks map to threads in the threads package used by GNAT. Where possible and appropriate, these threads correspond to native threads of the underlying operating system.

95. Implementation-defined `policy_identifiers` allowed in a pragma `Task_Dispatching_Policy`. See D.2.2(3).

There are no implementation-defined policy-identifiers allowed in this pragma.

96. Implementation-defined aspects of priority inversion. See D.2.2(16).

Execution of a task cannot be preempted by the implementation processing of delay expirations for lower priority tasks.

97. Implementation defined task dispatching. See D.2.2(18).

The policy is the same as that of the underlying threads implementation.

98. Implementation-defined `policy_identifiers` allowed in a pragma `Locking_Policy`. See D.3(4).

The only implementation defined policy permitted in GNAT is `Inheritance_Locking`. On targets that support this policy, locking is implemented by inheritance, i.e. the task owning the lock operates at a priority equal to the highest priority of any task currently requesting the lock.

99. Default ceiling priorities. See D.3(10).

The ceiling priority of protected objects of the type `System.Interrupt_Priority'Last` as described in the Ada 95 Reference Manual D.3(10),

100. The ceiling of any protected object used internally by the implementation. See D.3(16).

The ceiling priority of internal protected objects is `System.Priority'Last`.

101. Implementation-defined queuing policies. See D.4(1).

There are no implementation-defined queueing policies.

102. On a multiprocessor, any conditions that cause the completion of an aborted construct to be delayed later than what is specified for a single processor. See D.6(3).

The semantics for abort on a multi-processor is the same as on a single processor, there are no further delays.

103. Any operations that implicitly require heap storage allocation. See D.7(8).

The only operation that implicitly requires heap storage allocation is task creation.

104. Implementation-defined aspects of pragma `Restrictions`. See D.7(20).

There are no such implementation-defined aspects.

105. Implementation-defined aspects of package `Real_Time`. See D.8(17).

There are no implementation defined aspects of package `Real_Time`.

106. Implementation-defined aspects of `delay_statements`. See D.9(8).

Any difference greater than one microsecond will cause the task to be delayed (see D.9(7)).

107. The upper bound on the duration of interrupt blocking caused by the implementation. See D.12(5).

The upper bound is determined by the underlying operating system. In no cases is it more than 10 milliseconds.

108. The means for creating and executing distributed programs. See E(5).

The GLADE package provides a utility GNATDIST for creating and executing distributed programs. See the GLADE reference manual for further details.

109. Any events that can result in a partition becoming inaccessible. See E.1(7).

See the GLADE reference manual for full details on such events.

110. The scheduling policies, treatment of priorities, and management of shared resources between partitions in certain cases. See E.1(11).

See the GLADE reference manual for full details on these aspects of multi-partition execution.

111. Events that cause the version of a compilation unit to change. See E.3(5).

Editing the source file of a compilation unit, or the source files of any units on which it is dependent in a significant way cause the version to change. No other actions cause the

version number to change. All changes are significant except those which affect only layout, capitalization or comments.

112. Whether the execution of the remote subprogram is immediately aborted as a result of cancellation. See E.4(13).

See the GLADE reference manual for details on the effect of abort in a distributed application.

113. Implementation-defined aspects of the PCS. See E.5(25).

See the GLADE reference manual for a full description of all implementation defined aspects of the PCS.

114. Implementation-defined interfaces in the PCS. See E.5(26).

See the GLADE reference manual for a full description of all implementation defined interfaces.

115. The values of named numbers in the package `Decimal`. See F.2(7).

`Max_Scale`
+18

`Min_Scale`
-18

`Min_Delta`
1.0E-18

`Max_Delta`
1.0E+18

`Max_Decimal_Digits`
18

116. The value of `Max_Picture_Length` in the package `Text_IO.Editing`. See F.3.3(16).

64

117. The value of `Max_Picture_Length` in the package `Wide_Text_IO.Editing`. See F.3.4(5).

64

118. The accuracy actually achieved by the complex elementary functions and by other complex arithmetic operations. See G.1(1).

Standard library functions are used for the complex arithmetic operations. Only fast math mode is currently supported.

119. The sign of a zero result (or a component thereof) from any operator or function in `Numerics.Generic_Complex_Types`, when `Real'Signed_Zeros` is True. See G.1.1(53).

The signs of zero values are as recommended by the relevant implementation advice.

120. The sign of a zero result (or a component thereof) from any operator or function in `Numerics.Generic_Complex_Elementary_Functions`, when `Real'Signed_Zeros` is `True`. See G.1.2(45).

The signs of zero values are as recommended by the relevant implementation advice.

121. Whether the strict mode or the relaxed mode is the default. See G.2(2).

The strict mode is the default. There is no separate relaxed mode. GNAT provides a highly efficient implementation of strict mode.

122. The result interval in certain cases of fixed-to-float conversion. See G.2.1(10).

For cases where the result interval is implementation dependent, the accuracy is that provided by performing all operations in 64-bit IEEE floating-point format.

123. The result of a floating point arithmetic operation in overflow situations, when the `Machine_Overflows` attribute of the result type is `False`. See G.2.1(13).

Infinite and Nan values are produced as dictated by the IEEE floating-point standard.

124. The result interval for division (or exponentiation by a negative exponent), when the floating point hardware implements division as multiplication by a reciprocal. See G.2.1(16).

Not relevant, division is IEEE exact.

125. The definition of close result set, which determines the accuracy of certain fixed point multiplications and divisions. See G.2.3(5).

Operations in the close result set are performed using IEEE long format floating-point arithmetic. The input operands are converted to floating-point, the operation is done in floating-point, and the result is converted to the target type.

126. Conditions on a `universal_real` operand of a fixed point multiplication or division for which the result shall be in the perfect result set. See G.2.3(22).

The result is only defined to be in the perfect result set if the result can be computed by a single scaling operation involving a scale factor representable in 64-bits.

127. The result of a fixed point arithmetic operation in overflow situations, when the `Machine_Overflows` attribute of the result type is `False`. See G.2.3(27).

Not relevant, `Machine_Overflows` is `True` for fixed-point types.

128. The result of an elementary function reference in overflow situations, when the `Machine_Overflows` attribute of the result type is `False`. See G.2.4(4).

IEEE infinite and Nan values are produced as appropriate.

129. The value of the angle threshold, within which certain elementary functions, complex arithmetic operations, and complex elementary functions yield results conforming to a maximum relative error bound. See G.2.4(10).

Information on this subject is not yet available.

130. The accuracy of certain elementary functions for parameters beyond the angle threshold. See G.2.4(10).

Information on this subject is not yet available.

131. The result of a complex arithmetic operation or complex elementary function reference in overflow situations, when the `Machine_Overflows` attribute of the corresponding real type is `False`. See G.2.6(5).

IEEE infinite and Nan values are produced as appropriate.

132. The accuracy of certain complex arithmetic operations and certain complex elementary functions for parameters (or components thereof) beyond the angle threshold. See G.2.6(8).

Information on those subjects is not yet available.

133. Information regarding bounded errors and erroneous execution. See H.2(1).

Information on this subject is not yet available.

134. Implementation-defined aspects of pragma `Inspection_Point`. See H.3.2(8).

Pragma `Inspection_Point` ensures that the variable is live and can be examined by the debugger at the inspection point.

135. Implementation-defined aspects of pragma `Restrictions`. See H.4(25).

There are no implementation-defined aspects of pragma `Restrictions`. The use of pragma `Restrictions [No_Exceptions]` has no effect on the generated code. Checks must suppressed by use of pragma `Suppress`.

136. Any restrictions on pragma `Restrictions`. See H.4(27).

There are no restrictions on pragma `Restrictions`.

5 Intrinsic Subprograms

GNAT allows a user application program to write the declaration:

```
pragma Import (Intrinsic, name);
```

providing that the name corresponds to one of the implemented intrinsic subprograms in GNAT, and that the parameter profile of the referenced subprogram meets the requirements. This chapter describes the set of implemented intrinsic subprograms, and the requirements on parameter profiles. Note that no body is supplied; as with other uses of pragma Import, the body is supplied elsewhere (in this case by the compiler itself). Note that any use of this feature is potentially non-portable, since the Ada standard does not require Ada compilers to implement this feature.

5.1 Intrinsic Operators

All the predefined numeric operators in package Standard in `pragma Import (Intrinsic,..)` declarations. In the binary operator case, the operands must have the same size. The operand or operands must also be appropriate for the operator. For example, for addition, the operands must both be floating-point or both be fixed-point, and the right operand for `"**"` must have a root type of `Standard.Integer'Base`. You can use an intrinsic operator declaration as in the following example:

```
type Int1 is new Integer;
type Int2 is new Integer;

function "+" (X1 : Int1; X2 : Int2) return Int1;
function "+" (X1 : Int1; X2 : Int2) return Int2;
pragma Import (Intrinsic, "+");
```

This declaration would permit "mixed mode" arithmetic on items of the differing types `Int1` and `Int2`. It is also possible to specify such operators for private types, if the full views are appropriate arithmetic types.

5.2 Enclosing_Entity

This intrinsic subprogram is used in the implementation of the library routine `GNAT.Source_Info`. The only useful use of the intrinsic import in this case is the one in this unit, so an application program should simply call the function `GNAT.Source_Info.Enclosing_Entity` to obtain the name of the current subprogram, package, task, entry, or protected subprogram.

5.3 Exception_Information

This intrinsic subprogram is used in the implementation of the library routine `GNAT.Current_Exception`. The only useful use of the intrinsic import in this case is the one in this unit, so an application program should simply call the function `GNAT.Current_Exception.Exception_Information` to obtain the exception information associated with the current exception.

5.4 Exception_Message

This intrinsic subprogram is used in the implementation of the library routine `GNAT.Current_Exception`. The only useful use of the intrinsic import in this case is the one in this unit, so an application program should simply call the function `GNAT.Current_Exception.Exception_Message` to obtain the message associated with the current exception.

5.5 Exception_Name

This intrinsic subprogram is used in the implementation of the library routine `GNAT.Current_Exception`. The only useful use of the intrinsic import in this case is the one in this unit, so an application program should simply call the function `GNAT.Current_Exception.Exception_Name` to obtain the name of the current exception.

5.6 File

This intrinsic subprogram is used in the implementation of the library routine `GNAT.Source_Info`. The only useful use of the intrinsic import in this case is the one in this unit, so an application program should simply call the function `GNAT.Source_Info.File` to obtain the name of the current file.

5.7 Line

This intrinsic subprogram is used in the implementation of the library routine `GNAT.Source_Info`. The only useful use of the intrinsic import in this case is the one in this unit, so an application program should simply call the function `GNAT.Source_Info.Line` to obtain the number of the current source line.

5.8 Rotate_Left

In standard Ada 95, the `Rotate_Left` function is available only for the predefined modular types in package `Interfaces`. However, in GNAT it is possible to define a Rotate_Left function for a user defined modular type or any signed integer type as in this example:

```
function Shift_Left
  (Value  : My_Modular_Type;
   Amount : Natural)
   return   My_Modular_Type;
```

The requirements are that the profile be exactly as in the example above. The only modifications allowed are in the formal parameter names, and in the type of `Value` and the return type, which must be the same, and must be either a signed integer type, or a modular integer type with a binary modulus, and the size must be 8. 16, 32 or 64 bits.

5.9 Rotate_Right

A `Rotate_Right` function can be defined for any user defined binary modular integer type, or signed integer type, as described above for `Rotate_Left`.

5.10 Shift_Left

A `Shift_Left` function can be defined for any user defined binary modular integer type, or signed integer type, as described above for `Rotate_Left`.

5.11 Shift_Right

A `Shift_Right` function can be defined for any user defined binary modular integer type, or signed integer type, as described above for `Rotate_Left`.

5.12 Shift_Right_Arithmetic

A `Shift_Right_Arithmetic` function can be defined for any user defined binary modular integer type, or signed integer type, as described above for `Rotate_Left`.

5.13 Source_Location

This intrinsic subprogram is used in the implementation of the library routine `GNAT.Source_Info`. The only useful use of the intrinsic import in this case is the one in this unit, so an application program should simply call the function `GNAT.Source_Info.Source_Location` to obtain the current source file location.

6 Representation Clauses and Pragmas

This section describes the representation clauses accepted by GNAT, and their effect on the representation of corresponding data objects.

GNAT fully implements Annex C (Systems Programming). This means that all the implementation advice sections in chapter 13 are fully implemented. However, these sections only require a minimal level of support for representation clauses. GNAT provides much more extensive capabilities, and this section describes the additional capabilities provided.

6.1 Alignment Clauses

GNAT requires that all alignment clauses specify a power of 2, and all default alignments are always a power of 2. The default alignment values are as follows:

- Primitive Types For primitive types, the alignment is the maximum of the actual size of objects of the type, and the maximum alignment supported by the target. For example, for type Long_Float, the object size is 8 bytes, and the default alignment will be 8 on any target that supports alignments this large, but on some targets, the maximum alignment may be smaller than 8, in which case objects of type Long_Float will be maximally aligned.
- Arrays For arrays, the alignment is equal to the alignment of the component type for the normal case where no packing or component size is given. If the array is packed, and the packing is effective (see separate section on packed arrays), then the alignment will be one for long packed arrays, or arrays whose length is not known at compile time. For short packed arrays, which are handled internally as modular types, the alignment will be as described for primitive types, e.g. a packed array of length 31 bits will have an object size of four bytes, and an alignment of 4.
- Records For the normal non-packed case, the alignment of a record is equal to the maximum alignment of any of its components. For tagged records, this includes the implicit access type used for the tag. If a pragma `Pack` is used and all fields are packable (see separate section on pragma `Pack`), then the resulting alignment is 1.

  A special case is when the size of the record is given explicitly, or a full record representation clause is given, and the size of the record is 2, 4, or 8 bytes. In this case, an alignment is chosen to match the size of the record. For example, if we have:

  ```
  type Small is record
     A, B : Character;
  end record;
  ```

  then the default alignment of the record type `Small` is 2, not 1. This leads to more efficient code when the record is treated as a unit, and also allows the type to specified as `Atomic` on architectures requiring strict alignment.

An alignment clause may always specify a larger alignment than the default value, up to some maximum value dependent on the target (obtainable by using the attribute reference System'Maximum_Alignment). The only case in which it is permissible to specify a smaller alignment than the default value is in the case of a record for which a record representation clause is given. In this case, packable fields for which a component clause is given still result in a default alignment corresponding to the original type, but this may be overridden, since these components in fact only require an alignment of one byte. For example, given

```
type v is record
   a : integer;
end record;

for v use record
   a at 0  range 0 .. 31;
end record;

for v'alignment use 1;
```

The default alignment for the type `v` is 4, as a result of the integer field in the record, but since this field is placed with a component clause, it is permissible, as shown, to override the default alignment of the record to a smaller value.

6.2 Size Clauses

The default size of types is as specified in the reference manual. For objects, GNAT will generally increase the type size so that the object size is a multiple of storage units, and also a multiple of the alignment. For example

```
type Smallint is range 1 .. 6;

type Rec is record
   y1 : integer;
   y2 : boolean;
end record;
```

In this example, `Smallint` has a size of 3, as specified by the RM rules, but objects of this type will have a size of 8, since objects by default occupy an integral number of storage units. On some targets, notably older versions of the Digital Alpha, the size of stand alone objects of this type may be 32, reflecting the inability of the hardware to do byte load/stores.

Similarly, the size of type `Rec` is 40 bits, but the alignment is 4, so objects of this type will have their size increased to 64 bits so that it is a multiple of the alignment. The reason for this decision, which is in accordance with the specific note in RM 13.3(43):

```
A Size clause should be supported for an object if the specified
Size is at least as large as its subtype's Size, and corresponds
to a size in storage elements that is a multiple of the object's
Alignment (if the Alignment is nonzero).
```

An explicit size clause may be used to override the default size by increasing it. For example, if we have:

```
type My_Boolean is new Boolean;
for My_Boolean'Size use 32;
```

then objects of this type will always be 32 bits long. In the case of discrete types, the size can be increased up to 64 bits, with the effect that the entire specified field is used to hold the value, sign- or zero-extended as appropriate. If more than 64 bits is specified, then padding space is allocated after the value, and a warning is issued that there are unused bits.

Similarly the size of records and arrays may be increased, and the effect is to add padding bits after the value. This also causes a warning message to be generated.

The largest Size value permitted in GNAT is 2**32−1. Since this is a Size in bits, this corresponds to an object of size 256 megabytes (minus one). This limitation is true on all

targets. The reason for this limitation is that it improves the quality of the code in many cases if it is known that a Size value can be accommodated in an object of type Integer.

6.3 Storage_Size Clauses

For tasks, the `Storage_Size` clause specifies the amount of space to be allocated for the task stack. This cannot be extended, and if the stack is exhausted, then `Storage_Error` will be raised if stack checking is enabled. If the default size of 20K bytes is insufficient, then you need to use a `Storage_Size` attribute definition clause, or a `Storage_Size` pragma in the task definition to set the appropriate required size. A useful technique is to include in every task definition a pragma of the form:

```
pragma Storage_Size (Default_Stack_Size);
```

Then Default_Stack_Size can be defined in a global package, and modified as required. Any tasks requiring different task stack sizes from the default can have an appropriate alternative reference in the pragma.

For access types, the `Storage_Size` clause specifies the maximum space available for allocation of objects of the type. If this space is exceeded then `Storage_Error` will be raised by an allocation attempt. In the case where the access type is declared local to a subprogram, the use of a `Storage_Size` clause triggers automatic use of a special predefined storage pool (`System.Pool_Size`) that ensures that all space for the pool is automatically reclaimed on exit from the scope in which the type is declared.

A special case recognized by the compiler is the specification of a `Storage_Size` of zero for an access type. This means that no items can be allocated from the pool, and this is recognized at compile time, and all the overhead normally associated with maintaining a fixed size storage pool is eliminated. Consider the following example:

```
procedure p is
   type R is array (Natural) of Character;
   type P is access all R;
   for P'Storage_Size use 0;
   --  Above access type intended only for interfacing purposes

   y : P;

   procedure g (m : P);
   pragma Import (C, g);

   --  ...

begin
   --  ...
   y := new R;
end;
```

As indicated in this example, these dummy storage pools are often useful in connection with interfacing where no object will ever be allocated. If you compile the above example, you get the warning:

```
p.adb:16:09: warning: allocation from empty storage pool
p.adb:16:09: warning: Storage_Error will be raised at run time
```

Of course in practice, there will not be any explicit allocators in the case of such an access declaration.

6.4 Size of Variant Record Objects

An issue arises in the case of variant record objects of whether Size gives information about a particular variant, or the maximum size required for any variant. Consider the following program

```
with Text_IO; use Text_IO;
procedure q is
   type R1 (A : Boolean := False) is record
     case A is
       when True  => X : Character;
       when False => null;
     end case;
   end record;

   V1 : R1 (False);
   V2 : R1;

begin
   Put_Line (Integer'Image (V1'Size));
   Put_Line (Integer'Image (V2'Size));
end q;
```

Here we are dealing with a variant record, where the True variant requires 16 bits, and the False variant requires 8 bits. In the above example, both V1 and V2 contain the False variant, which is only 8 bits long. However, the result of running the program is:

```
8
16
```

The reason for the difference here is that the discriminant value of V1 is fixed, and will always be False. It is not possible to assign a True variant value to V1, therefore 8 bits is sufficient. On the other hand, in the case of V2, the initial discriminant value is False (from the default), but it is possible to assign a True variant value to V2, therefore 16 bits must be allocated for V2 in the general case, even fewer bits may be needed at any particular point during the program execution.

As can be seen from the output of this program, the `'Size` attribute applied to such an object in GNAT gives the actual allocated size of the variable, which is the largest size of any of the variants. The Ada Reference Manual is not completely clear on what choice should be made here, but the GNAT behavior seems most consistent with the language in the RM.

In some cases, it may be desirable to obtain the size of the current variant, rather than the size of the largest variant. This can be achieved in GNAT by making use of the fact that in the case of a subprogram parameter, GNAT does indeed return the size of the current variant (because a subprogram has no way of knowing how much space is actually allocated for the actual).

Consider the following modified version of the above program:

```
with Text_IO; use Text_IO;
procedure q is
   type R1 (A : Boolean := False) is record
     case A is
       when True  => X : Character;
       when False => null;
     end case;
```

```
      end record;

      V2 : R1;

      function Size (V : R1) return Integer is
      begin
         return V'Size;
      end Size;

   begin
      Put_Line (Integer'Image (V2'Size));
      Put_Line (Integer'IMage (Size (V2)));
      V2 := (True, 'x');
      Put_Line (Integer'Image (V2'Size));
      Put_Line (Integer'IMage (Size (V2)));
   end q;
```

The output from this program is

```
16
8
16
16
```

Here we see that while the `'Size` attribute always returns the maximum size, regardless of the current variant value, the `Size` function does indeed return the size of the current variant value.

6.5 Biased Representation

In the case of scalars with a range starting at other than zero, it is possible in some cases to specify a size smaller than the default minimum value, and in such cases, GNAT uses an unsigned biased representation, in which zero is used to represent the lower bound, and successive values represent successive values of the type.

For example, suppose we have the declaration:

```
type Small is range -7 .. -4;
for Small'Size use 2;
```

Although the default size of type `Small` is 4, the `Size` clause is accepted by GNAT and results in the following representation scheme:

```
-7 is represented as 2#00#
-6 is represented as 2#01#
-5 is represented as 2#10#
-4 is represented as 2#11#
```

Biased representation is only used if the specified `Size` clause cannot be accepted in any other manner. These reduced sizes that force biased representation can be used for all discrete types except for enumeration types for which a representation clause is given.

6.6 Value_Size and Object_Size Clauses

In Ada 95, the `Size` of a discrete type is the minimum number of bits required to hold values of the type. Although this interpretation was allowed in Ada 83, it was not required, and this requirement in practice can cause some significant difficulties. For example, in most

Ada 83 compilers, `Natural'Size` was 32. However, in Ada-95, `Natural'Size` is typically 31. This means that code may change in behavior when moving from Ada 83 to Ada 95. For example, consider:

```
type Rec is record;
   A : Natural;
   B : Natural;
end record;

for Rec use record
   for A use at 0  range 0 .. Natural'Size - 1;
   for B use at 0  range Natural'Size .. 2 * Natural'Size - 1;
end record;
```

In the above code, since the typical size of `Natural` objects is 32 bits and `Natural'Size` is 31, the above code can cause unexpected inefficient packing in Ada 95, and in general there are surprising cases where the fact that the object size can exceed the size of the type causes surprises.

To help get around this problem GNAT provides two implementation dependent attributes `Value_Size` and `Object_Size`. When applied to a type, these attributes yield the size of the type (corresponding to the RM defined size attribute), and the size of objects of the type respectively.

The `Object_Size` is used for determining the default size of objects and components. This size value can be referred to using the `Object_Size` attribute. The phrase "is used" here means that it is the basis of the determination of the size. The backend is free to pad this up if necessary for efficiency, e.g. an 8-bit stand-alone character might be stored in 32 bits on a machine with no efficient byte access instructions such as the Alpha.

The default rules for the value of `Object_Size` for fixed-point and discrete types are as follows:

- The `Object_Size` for base subtypes reflect the natural hardware size in bits (run the utility `gnatpsta` to find those values for numeric types). Enumeration types and fixed-point base subtypes have 8, 16, 32 or 64 bits for this size, depending on the range of values to be stored.
- The `Object_Size` of a subtype is the same as the `Object_Size` of the type from which it is obtained.
- The `Object_Size` of a derived base type is copied from the parent base type, and the `Object_Size` of a derived first subtype is copied from the parent first subtype.

The `Value_Size` attribute is the number of bits required to store a value of the type. This size can be referred to using the `Value_Size` attribute. This value is used to determine how tightly to pack records or arrays with components of this type, and also affects the semantics of unchecked conversion (unchecked conversions where the `Value_Size` values differ generate a warning, and are potentially target dependent).

The default rules for the value of `Value_Size` are as follows:

- The `Value_Size` for a base subtype is the minimum number of bits required to store all values of the type (including the sign bit only if negative values are possible).
- If a subtype statically matches the first subtype of a given type, then it has by default the same `Value_Size` as the first subtype. This is a consequence of RM 13.1(14) ("if two subtypes statically match, then their subtype-specific aspects are the same".)

- All other subtypes have a `Value_Size` corresponding to the minimum number of bits required to store all values of the subtype. For dynamic bounds, it is assumed that the value can range down or up to the corresponding bound of the ancestor

The RM defined attribute `Size` corresponds to the `Value_Size` attribute.

The `Size` attribute may be defined for a first-named subtype. This sets the `Value_Size` of the first-named subtype to the given value, and the `Object_Size` of this first-named subtype to the given value padded up to an appropriate boundary. It is a consequence of the default rules above that this `Object_Size` will apply to all further subtypes. On the other hand, `Value_Size` is affected only for the first subtype, any dynamic subtypes obtained from it directly, and any statically matching subtypes. The `Value_Size` of any other static subtypes is not affected.

`Value_Size` and `Object_Size` may be explicitly set for any subtype using an attribute definition clause. Note that the use of these attributes can cause the RM 13.1(14) rule to be violated. If two access types reference aliased objects whose subtypes have differing `Object_Size` values as a result of explicit attribute definition clauses, then it is erroneous to convert from one access subtype to the other.

At the implementation level, Esize stores the Object_SIze and the RM_Size field stores the `Value_Size` (and hence the value of the `Size` attribute, which, as noted above, is equivalent to `Value_Size`).

To get a feel for the difference, consider the following examples (note that in each case the base is short_short_integer with a size of 8):

```
                                        Object_Size      Value_Size

type x1 is range 0 .. 5;                     8               3

type x2 is range 0 .. 5;
for x2'size use 12;                         12              12

subtype x3 is x2 range 0 .. 3;              12               2

subtype x4 is x2'base range 0 .. 10;         8               4

subtype x5 is x2 range 0 .. dynamic;        12              (7)

subtype x6 is x2'base range 0 .. dynamic;    8              (7)
```

Note: the entries marked (7) are not actually specified by the Ada 95 RM, but it seems in the spirit of the RM rules to allocate the minimum number of bits known to be large enough to hold the given range of values.

So far, so good, but GNAT has to obey the RM rules, so the question is under what conditions must the RM `Size` be used. The following is a list of the occasions on which the RM `Size` must be used:

- Component size for packed arrays or records
- Value of the attribute `Size` for a type
- Warning about sizes not matching for unchecked conversion

For types other than discrete and fixed-point types, the `Object_Size` and Value_Size are the same (and equivalent to the RM attribute `Size`). Only `Size` may be specified for such types.

6.7 Component_Size Clauses

Normally, the value specified in a component clause must be consistent with the subtype of the array component with regard to size and alignment. In other words, the value specified must be at least equal to the size of this subtype, and must be a multiple of the alignment value.

In addition, component size clauses are allowed which cause the array to be packed, by specifying a smaller value. The cases in which this is allowed are for component size values in the range 1 through 63. The value specified must not be smaller than the Size of the subtype. GNAT will accurately honor all packing requests in this range. For example, if we have:

```
type r is array (1 .. 8) of Natural;
for r'Size use 31;
```

then the resulting array has a length of 31 bytes (248 bits = 8 * 31). Of course access to the components of such an array is considerably less efficient than if the natural component size of 32 is used.

6.8 Bit_Order Clauses

For record subtypes, GNAT permits the specification of the `Bit_Order` attribute. The specification may either correspond to the default bit order for the target, in which case the specification has no effect and places no additional restrictions, or it may be for the non-standard setting (that is the opposite of the default).

In the case where the non-standard value is specified, the effect is to renumber bits within each byte, but the ordering of bytes is not affected. There are certain restrictions placed on component clauses as follows:

- Components fitting within a single storage unit. These are unrestricted, and the effect is merely to renumber bits. For example if we are on a little-endian machine with `Low_Order_First` being the default, then the following two declarations have exactly the same effect:

```
type R1 is record
   A : Boolean;
   B : Integer range 1 .. 120;
end record;

for R1 use record
   A at 0 range 0 .. 0;
   B at 0 range 1 .. 7;
end record;

type R2 is record
   A : Boolean;
   B : Integer range 1 .. 120;
end record;
```

```
for R2'Bit_Order use High_Order_First;

for R2 use record
   A at 0 range 7 .. 7;
   B at 0 range 0 .. 6;
end record;
```

The useful application here is to write the second declaration with the `Bit_Order` attribute definition clause, and know that it will be treated the same, regardless of whether the target is little-endian or big-endian.

- Components occupying an integral number of bytes. These are components that exactly fit in two or more bytes. Such component declarations are allowed, but have no effect, since it is important to realize that the `Bit_Order` specification does not affect the ordering of bytes. In particular, the following attempt at getting an endian-independent integer does not work:

```
type R2 is record
   A : Integer;
end record;

for R2'Bit_Order use High_Order_First;

for R2 use record
   A at 0 range 0 .. 31;
end record;
```

  This declaration will result in a little-endian integer on a little-endian machine, and a big-endian integer on a big-endian machine. If byte flipping is required for interoperability between big- and little-endian machines, this must be explicitly programmed. This capability is not provided by `Bit_Order`.

- Components that are positioned across byte boundaries but do not occupy an integral number of bytes. Given that bytes are not reordered, such fields would occupy a non-contiguous sequence of bits in memory, requiring non-trivial code to reassemble. They are for this reason not permitted, and any component clause specifying such a layout will be flagged as illegal by GNAT.

Since the misconception that Bit_Order automatically deals with all endian-related incompatibilities is a common one, the specification of a component field that is an integral number of bytes will always generate a warning. This warning may be suppressed using `pragma Suppress` if desired. The following section contains additional details regarding the issue of byte ordering.

6.9 Effect of Bit_Order on Byte Ordering

In this section we will review the effect of the `Bit_Order` attribute definition clause on byte ordering. Briefly, it has no effect at all, but a detailed example will be helpful. Before giving this example, let us review the precise definition of the effect of defining `Bit_Order`. The effect of a non-standard bit order is described in section 15.5.3 of the Ada Reference Manual:

```
2   A bit ordering is a method of interpreting the meaning of
the storage place attributes.
```

To understand the precise definition of storage place attributes in this context, we visit section 13.5.1 of the manual:

```
13   A record_representation_clause (without the mod_clause)
specifies the layout.  The storage place attributes (see 13.5.2)
are taken from the values of the position, first_bit, and last_bit
expressions after normalizing those values so that first_bit is
less than Storage_Unit.
```

The critical point here is that storage places are taken from the values after normalization, not before. So the `Bit_Order` interpretation applies to normalized values. The interpretation is described in the later part of the 15.5.3 paragraph:

```
2   A bit ordering is a method of interpreting the meaning of
the storage place attributes.  High_Order_First (known in the
vernacular as ``big endian'') means that the first bit of a
storage element (bit 0) is the most significant bit (interpreting
the sequence of bits that represent a component as an unsigned
integer value).  Low_Order_First (known in the vernacular as
``little endian'') means the opposite: the first bit is the
least significant.
```

Note that the numbering is with respect to the bits of a storage unit. In other words, the specification affects only the numbering of bits within a single storage unit.

We can make the effect clearer by giving an example.

Suppose that we have an external device which presents two bytes, the first byte presented, which is the first (low addressed byte) of the two byte record is called Master, and the second byte is called Slave.

The left most (most significant bit is called Control for each byte, and the remaining 7 bits are called V1, V2, ... V7, where V7 is the rightmost (least significant) bit.

On a big-endian machine, we can write the following representation clause

```
type Data is record
   Master_Control : Bit;
   Master_V1      : Bit;
   Master_V2      : Bit;
   Master_V3      : Bit;
   Master_V4      : Bit;
   Master_V5      : Bit;
   Master_V6      : Bit;
   Master_V7      : Bit;
   Slave_Control  : Bit;
   Slave_V1       : Bit;
   Slave_V2       : Bit;
   Slave_V3       : Bit;
   Slave_V4       : Bit;
   Slave_V5       : Bit;
   Slave_V6       : Bit;
   Slave_V7       : Bit;
end record;

for Data use record
   Master_Control at 0 range 0 .. 0;
   Master_V1      at 0 range 1 .. 1;
   Master_V2      at 0 range 2 .. 2;
   Master_V3      at 0 range 3 .. 3;
   Master_V4      at 0 range 4 .. 4;
   Master_V5      at 0 range 5 .. 5;
```

```
   Master_V6      at 0 range 6 .. 6;
   Master_V7      at 0 range 7 .. 7;
   Slave_Control  at 1 range 0 .. 0;
   Slave_V1       at 1 range 1 .. 1;
   Slave_V2       at 1 range 2 .. 2;
   Slave_V3       at 1 range 3 .. 3;
   Slave_V4       at 1 range 4 .. 4;
   Slave_V5       at 1 range 5 .. 5;
   Slave_V6       at 1 range 6 .. 6;
   Slave_V7       at 1 range 7 .. 7;
end record;
```

Now if we move this to a little endian machine, then the bit ordering within the byte is backwards, so we have to rewrite the record rep clause as:

```
for Data use record
   Master_Control at 0 range 7 .. 7;
   Master_V1      at 0 range 6 .. 6;
   Master_V2      at 0 range 5 .. 5;
   Master_V3      at 0 range 4 .. 4;
   Master_V4      at 0 range 3 .. 3;
   Master_V5      at 0 range 2 .. 2;
   Master_V6      at 0 range 1 .. 1;
   Master_V7      at 0 range 0 .. 0;
   Slave_Control  at 1 range 7 .. 7;
   Slave_V1       at 1 range 6 .. 6;
   Slave_V2       at 1 range 5 .. 5;
   Slave_V3       at 1 range 4 .. 4;
   Slave_V4       at 1 range 3 .. 3;
   Slave_V5       at 1 range 2 .. 2;
   Slave_V6       at 1 range 1 .. 1;
   Slave_V7       at 1 range 0 .. 0;
end record;
```

It is a nuisance to have to rewrite the clause, especially if the code has to be maintained on both machines. However, this is a case that we can handle with the `Bit_Order` attribute if it is implemented. Note that the implementation is not required on byte addressed machines, but it is indeed implemented in GNAT. This means that we can simply use the first record clause, together with the declaration

```
for Data'Bit_Order use High_Order_First;
```

and the effect is what is desired, namely the layout is exactly the same, independent of whether the code is compiled on a big-endian or little-endian machine.

The important point to understand is that byte ordering is not affected. A `Bit_Order` attribute definition never affects which byte a field ends up in, only where it ends up in that byte. To make this clear, let us rewrite the record rep clause of the previous example as:

```
for Data'Bit_Order use High_Order_First;
for Data use record
   Master_Control at 0 range  0 .. 0;
   Master_V1      at 0 range  1 .. 1;
   Master_V2      at 0 range  2 .. 2;
   Master_V3      at 0 range  3 .. 3;
   Master_V4      at 0 range  4 .. 4;
   Master_V5      at 0 range  5 .. 5;
   Master_V6      at 0 range  6 .. 6;
   Master_V7      at 0 range  7 .. 7;
   Slave_Control  at 0 range  8 .. 8;
   Slave_V1       at 0 range  9 .. 9;
```

```
   Slave_V2       at 0 range 10 .. 10;
   Slave_V3       at 0 range 11 .. 11;
   Slave_V4       at 0 range 12 .. 12;
   Slave_V5       at 0 range 13 .. 13;
   Slave_V6       at 0 range 14 .. 14;
   Slave_V7       at 0 range 15 .. 15;
end record;
```

This is exactly equivalent to saying (a repeat of the first example):

```
for Data'Bit_Order use High_Order_First;
for Data use record
   Master_Control at 0 range 0 .. 0;
   Master_V1      at 0 range 1 .. 1;
   Master_V2      at 0 range 2 .. 2;
   Master_V3      at 0 range 3 .. 3;
   Master_V4      at 0 range 4 .. 4;
   Master_V5      at 0 range 5 .. 5;
   Master_V6      at 0 range 6 .. 6;
   Master_V7      at 0 range 7 .. 7;
   Slave_Control  at 1 range 0 .. 0;
   Slave_V1       at 1 range 1 .. 1;
   Slave_V2       at 1 range 2 .. 2;
   Slave_V3       at 1 range 3 .. 3;
   Slave_V4       at 1 range 4 .. 4;
   Slave_V5       at 1 range 5 .. 5;
   Slave_V6       at 1 range 6 .. 6;
   Slave_V7       at 1 range 7 .. 7;
end record;
```

Why are they equivalent? Well take a specific field, the `Slave_V2` field. The storage place attributes are obtained by normalizing the values given so that the `First_Bit` value is less than 8. After nromalizing the values (0,10,10) we get (1,2,2) which is exactly what we specified in the other case.

Now one might expect that the `Bit_Order` attribute might affect bit numbering within the entire record component (two bytes in this case, thus affecting which byte fields end up in), but that is not the way this feature is defined, it only affects numbering of bits, not which byte they end up in.

Consequently it never makes sense to specify a starting bit number greater than 7 (for a byte addressable field) if an attribute definition for `Bit_Order` has been given, and indeed it may be actively confusing to specify such a value, so the compiler generates a warning for such usage.

If you do need to control byte ordering then appropriate conditional values must be used. If in our example, the slave byte came first on some machines we might write:

```
Master_Byte_First constant Boolean := ...;

Master_Byte : constant Natural :=
                1 - Boolean'Pos (Master_Byte_First);
Slave_Byte  : constant Natural :=
                Boolean'Pos (Master_Byte_First);

for Data'Bit_Order use High_Order_First;
for Data use record
   Master_Control at Master_Byte range 0 .. 0;
   Master_V1      at Master_Byte range 1 .. 1;
   Master_V2      at Master_Byte range 2 .. 2;
```

```
   Master_V3      at Master_Byte range 3 .. 3;
   Master_V4      at Master_Byte range 4 .. 4;
   Master_V5      at Master_Byte range 5 .. 5;
   Master_V6      at Master_Byte range 6 .. 6;
   Master_V7      at Master_Byte range 7 .. 7;
   Slave_Control  at Slave_Byte  range 0 .. 0;
   Slave_V1       at Slave_Byte  range 1 .. 1;
   Slave_V2       at Slave_Byte  range 2 .. 2;
   Slave_V3       at Slave_Byte  range 3 .. 3;
   Slave_V4       at Slave_Byte  range 4 .. 4;
   Slave_V5       at Slave_Byte  range 5 .. 5;
   Slave_V6       at Slave_Byte  range 6 .. 6;
   Slave_V7       at Slave_Byte  range 7 .. 7;
end record;
```

Now to switch between machines, all that is necessary is to set the boolean constant `Master_Byte_First` in an appropriate manner.

6.10 Pragma Pack for Arrays

Pragma `Pack` applied to an array has no effect unless the component type is packable. For a component type to be packable, it must be one of the following cases:

- Any scalar type
- Any fixed-point type
- Any type whose size is specified with a size clause
- Any packed array type with a static size

For all these cases, if the component subtype size is in the range 1 through 63, then the effect of the pragma `Pack` is exactly as though a component size were specified giving the component subtype size. For example if we have:

```
type r is range 0 .. 17;

type ar is array (1 .. 8) of r;
pragma Pack (ar);
```

Then the component size of `ar` will be set to 5 (i.e. to `r'size`, and the size of the array `ar` will be exactly 40 bits.

Note that in some cases this rather fierce approach to packing can produce unexpected effects. For example, in Ada 95, type Natural typically has a size of 31, meaning that if you pack an array of Natural, you get 31-bit close packing, which saves a few bits, but results in far less efficient access. Since many other Ada compilers will ignore such a packing request, GNAT will generate a warning on some uses of pragma `Pack` that it guesses might not be what is intended. You can easily remove this warning by using an explicit `Component_Size` setting instead, which never generates a warning, since the intention of the programmer is clear in this case.

GNAT treats packed arrays in one of two ways. If the size of the array is known at compile time and is less than 64 bits, then internally the array is represented as a single modular type, of exactly the appropriate number of bits. If the length is greater than 63 bits, or is not known at compile time, then the packed array is represented as an array of bytes, and the length is always a multiple of 8 bits.

6.11 Pragma Pack for Records

Pragma `Pack` applied to a record will pack the components to reduce wasted space from alignment gaps and by reducing the amount of space taken by components. We distinguish between package components and non-packable components. Components of the following types are considered packable:

- All scalar types are packable.
- All fixed-point types are represented internally as integers, and are packable.
- Small packed arrays, whose size does not exceed 64 bits, and where the size is statically known at compile time, are represented internally as modular integers, and so they are also packable.

All packable components occupy the exact number of bits corresponding to their `Size` value, and are packed with no padding bits, i.e. they can start on an arbitrary bit boundary.

All other types are non-packable, they occupy an integral number of storage units, and are placed at a boundary corresponding to their alignment requirements.

For example, consider the record

```
type Rb1 is array (1 .. 13) of Boolean;
pragma Pack (rb1);

type Rb2 is array (1 .. 65) of Boolean;
pragma Pack (rb2);

type x2 is record
   l1 : Boolean;
   l2 : Duration;
   l3 : Float;
   l4 : Boolean;
   l5 : Rb1;
   l6 : Rb2;
end record;
pragma Pack (x2);
```

The representation for the record x2 is as follows:

```
for x2'Size use 224;
for x2 use record
   l1 at  0 range  0 .. 0;
   l2 at  0 range  1 .. 64;
   l3 at 12 range  0 .. 31;
   l4 at 16 range  0 .. 0;
   l5 at 16 range  1 .. 13;
   l6 at 18 range  0 .. 71;
end record;
```

Studying this example, we see that the packable fields `l1` and `l2` are of length equal to their sizes, and placed at specific bit boundaries (and not byte boundaries) to eliminate padding. But `l3` is of a non-packable float type, so it is on the next appropriate alignment boundary.

The next two fields are fully packable, so `l4` and `l5` are minimally packed with no gaps. However, type `Rb2` is a packed array that is longer than 64 bits, so it is itself non-packable. Thus the `l6` field is aligned to the next byte boundary, and takes an integral number of bytes, i.e. 72 bits.

6.12 Record Representation Clauses

Record representation clauses may be given for all record types, including types obtained by record extension. Component clauses are allowed for any static component. The restrictions on component clauses depend on the type of the component.

For all components of an elementary type, the only restriction on component clauses is that the size must be at least the 'Size value of the type (actually the Value_Size). There are no restrictions due to alignment, and such components may freely cross storage boundaries.

Packed arrays with a size up to and including 64 bits are represented internally using a modular type with the appropriate number of bits, and thus the same lack of restriction applies. For example, if you declare:

```
type R is array (1 .. 49) of Boolean;
pragma Pack (R);
for R'Size use 49;
```

then a component clause for a component of type R may start on any specified bit boundary, and may specify a value of 49 bits or greater.

For non-primitive types, including packed arrays with a size greater than 64 bits, component clauses must respect the alignment requirement of the type, in particular, always starting on a byte boundary, and the length must be a multiple of the storage unit.

The tag field of a tagged type always occupies an address sized field at the start of the record. No component clause may attempt to overlay this tag.

In the case of a record extension T1, of a type T, no component clause applied to the type T1 can specify a storage location that would overlap the first T'Size bytes of the record.

6.13 Enumeration Clauses

The only restriction on enumeration clauses is that the range of values must be representable. For the signed case, if one or more of the representation values are negative, all values must be in the range:

```
System.Min_Int .. System.Max_Int
```

For the unsigned case, where all values are non negative, the values must be in the range:

```
0 .. System.Max_Binary_Modulus;
```

A *confirming* representation clause is one in which the values range from 0 in sequence, i.e. a clause that confirms the default representation for an enumeration type. Such a confirming representation is permitted by these rules, and is specially recognized by the compiler so that no extra overhead results from the use of such a clause.

If an array has an index type which is an enumeration type to which an enumeration clause has been applied, then the array is stored in a compact manner. Consider the declarations:

```
type r is (A, B, C);
for r use (A => 1, B => 5, C => 10);
type t is array (r) of Character;
```

The array type t corresponds to a vector with exactly three elements and has a default size equal to `3*Character'Size`. This ensures efficient use of space, but means that accesses to elements of the array will incur the overhead of converting representation values to the corresponding positional values, (i.e. the value delivered by the `Pos` attribute).

6.14 Address Clauses

The reference manual allows a general restriction on representation clauses, as found in RM 13.1(22):

```
An implementation need not support representation
items containing nonstatic expressions, except that
an implementation should support a representation item
for a given entity if each nonstatic expression in the
representation item is a name that statically denotes
a constant declared before the entity.
```

In practice this is applicable only to address clauses, since this is the only case in which a non-static expression is permitted by the syntax. As the AARM notes in sections 13.1 (22.a-22.h):

```
22.a   Reason: This is to avoid the following sort
       of thing:

22.b        X : Integer := F(...);
            Y : Address := G(...);
            for X'Address use Y;

22.c   In the above, we have to evaluate the
       initialization expression for X before we
       know where to put the result.  This seems
       like an unreasonable implementation burden.

22.d   The above code should instead be written
       like this:

22.e        Y : constant Address := G(...);
            X : Integer := F(...);
            for X'Address use Y;

22.f   This allows the expression ``Y'' to be safely
       evaluated before X is created.

22.g   The constant could be a formal parameter of mode in.

22.h   An implementation can support other nonstatic
       expressions if it wants to.  Expressions of type
       Address are hardly ever static, but their value
       might be known at compile time anyway in many
       cases.
```

GNAT does indeed permit many additional cases of non-static expressions. In particular, if the type involved is elementary there are no restrictions (since in this case, holding a temporary copy of the initialization value, if one is present, is inexpensive). In addition, if there is no implicit or explicit initialization, then there are no restrictions. GNAT will reject only the case where all three of these conditions hold:

- The type of the item is non-elementary (e.g. a record or array).
- There is explicit or implicit initialization required for the object.
- The address value is non-static. Here GNAT is more permissive than the RM, and allows the address value to be the address of a previously declared stand-alone variable, as long as it does not itself have an address clause.

```
Anchor : Some_Initialized_Type;
Overlay : Some_Initialized_Type;
for Overlay'Address use Anchor'Address;
```

However, the prefix of the address clause cannot be an array component, or a component of a discriminated record.

As noted above in section 22.h, address values are typically non-static. In particular the To_Address function, even if applied to a literal value, is a non-static function call. To avoid this minor annoyance, GNAT provides the implementation defined attribute 'To_Address. The following two expressions have identical values:

Another issue with address clauses is the interaction with alignment requirements. When an address clause is given for an object, the address value must be consistent with the alignment of the object (which is usually the same as the alignment of the type of the object). If an address clause is given that specifies an inappropriately aligned address value, then the program execution is erroneous.

Since this source of erroneous behavior can have unfortunate effects, GNAT checks (at compile time if possible, generating a warning, or at execution time with a run-time check) that the alignment is appropriate. If the run-time check fails, then `Program_Error` is raised. This run-time check is suppressed if range checks are suppressed, or if `pragma Restrictions (No_Elaboration_Code)` is in effect.

```
To_Address (16#1234_0000#)
System'To_Address (16#1234_0000#);
```

except that the second form is considered to be a static expression, and thus when used as an address clause value is always permitted.

Additionally, GNAT treats as static an address clause that is an unchecked_conversion of a static integer value. This simplifies the porting of legacy code, and provides a portable equivalent to the GNAT attribute To_Address.

An address clause cannot be given for an exported object. More understandably the real restriction is that objects with an address clause cannot be exported. This is because such variables are not defined by the Ada program, so there is no external object so export.

It is permissible to give an address clause and a pragma Import for the same object. In this case, the variable is not really defined by the Ada program, so there is no external symbol to be linked. The link name and the external name are ignored in this case. The reason that we allow this combination is that it provides a useful idiom to avoid unwanted initializations on objects with address clauses.

When an address clause is given for an object that has implicit or explicit initialization, then by default initialization takes place. This means that the effect of the object declaration is to overwrite the memory at the specified address. This is almost always not what the programmer wants, so GNAT will output a warning:

```
with System;
package G is
   type R is record
      M : Integer := 0;
   end record;

   Ext : R;
   for Ext'Address use System'To_Address (16#1234_1234#);
       |
```

```
>>> warning: implicit initialization of "Ext" may
    modify overlaid storage
>>> warning: use pragma Import for "Ext" to suppress
    initialization (RM B(24))

end G;
```

As indicated by the warning message, the solution is to use a (dummy) pragma Import to suppress this initialization. The pragma tell the compiler that the object is declared and initialized elsewhere. The following package compiles without warnings (and the initialization is suppressed):

```
with System;
package G is
   type R is record
      M : Integer := 0;
   end record;

   Ext : R;
   for Ext'Address use System'To_Address (16#1234_1234#);
   pragma Import (Ada, Ext);
end G;
```

6.15 Effect of Convention on Representation

Normally the specification of a foreign language convention for a type or an object has no effect on the chosen representation. In particular, the representation chosen for data in GNAT generally meets the standard system conventions, and for example records are laid out in a manner that is consistent with C. This means that specifying convention C (for example) has no effect.

There are three exceptions to this general rule:

- Convention Fortran and array subtypes If pragma Convention Fortran is specified for an array subtype, then in accordance with the implementation advice in section 3.6.2(11) of the Ada Reference Manual, the array will be stored in a Fortran-compatible column-major manner, instead of the normal default row-major order.
- Convention C and enumeration types GNAT normally stores enumeration types in 8, 16, or 32 bits as required to accommodate all values of the type. For example, for the enumeration type declared by:

  ```
  type Color is (Red, Green, Blue);
  ```

  8 bits is sufficient to store all values of the type, so by default, objects of type `Color` will be represented using 8 bits. However, normal C convention is to use 32 bits for all enum values in C, since enum values are essentially of type int. If pragma `Convention C` is specified for an Ada enumeration type, then the size is modified as necessary (usually to 32 bits) to be consistent with the C convention for enum values.
- Convention C/Fortran and Boolean types In C, the usual convention for boolean values, that is values used for conditions, is that zero represents false, and nonzero values represent true. In Ada, the normal convention is that two specific values, typically 0/1, are used to represent false/true respectively.

  Fortran has a similar convention for `LOGICAL` values (any nonzero value represents true).

To accommodate the Fortran and C conventions, if a pragma Convention specifies C or Fortran convention for a derived Boolean, as in the following example:

```
type C_Switch is new Boolean;
pragma Convention (C, C_Switch);
```

then the GNAT generated code will treat any nonzero value as true. For truth values generated by GNAT, the conventional value 1 will be used for True, but when one of these values is read, any nonzero value is treated as True.

6.16 Determining the Representations chosen by GNAT

Although the descriptions in this section are intended to be complete, it is often easier to simply experiment to see what GNAT accepts and what the effect is on the layout of types and objects.

As required by the Ada RM, if a representation clause is not accepted, then it must be rejected as illegal by the compiler. However, when a representation clause or pragma is accepted, there can still be questions of what the compiler actually does. For example, if a partial record representation clause specifies the location of some components and not others, then where are the non-specified components placed? Or if pragma `Pack` is used on a record, then exactly where are the resulting fields placed? The section on pragma `Pack` in this chapter can be used to answer the second question, but it is often easier to just see what the compiler does.

For this purpose, GNAT provides the option `-gnatR`. If you compile with this option, then the compiler will output information on the actual representations chosen, in a format similar to source representation clauses. For example, if we compile the package:

```
package q is
   type r (x : boolean) is tagged record
      case x is
         when True => S : String (1 .. 100);
         when False => null;
      end case;
   end record;

   type r2 is new r (false) with record
      y2 : integer;
   end record;

   for r2 use record
      y2 at 16 range 0 .. 31;
   end record;

   type x is record
      y : character;
   end record;

   type x1 is array (1 .. 10) of x;
   for x1'component_size use 11;

   type ia is access integer;

   type Rb1 is array (1 .. 13) of Boolean;
   pragma Pack (rb1);
```

```
   type Rb2 is array (1 .. 65) of Boolean;
   pragma Pack (rb2);

   type x2 is record
      l1 : Boolean;
      l2 : Duration;
      l3 : Float;
      l4 : Boolean;
      l5 : Rb1;
      l6 : Rb2;
   end record;
   pragma Pack (x2);
end q;
```

using the switch `-gnatR` we obtain the following output:

```
Representation information for unit q
-------------------------------------

for r'Size use ??;
for r'Alignment use 4;
for r use record
   x    at 4 range  0 .. 7;
   _tag at 0 range  0 .. 31;
   s    at 5 range  0 .. 799;
end record;

for r2'Size use 160;
for r2'Alignment use 4;
for r2 use record
   x       at  4 range  0 .. 7;
   _tag    at  0 range  0 .. 31;
   _parent at  0 range  0 .. 63;
   y2      at 16 range  0 .. 31;
end record;

for x'Size use 8;
for x'Alignment use 1;
for x use record
   y at 0 range  0 .. 7;
end record;

for x1'Size use 112;
for x1'Alignment use 1;
for x1'Component_Size use 11;

for rb1'Size use 13;
for rb1'Alignment use 2;
for rb1'Component_Size use 1;

for rb2'Size use 72;
for rb2'Alignment use 1;
for rb2'Component_Size use 1;

for x2'Size use 224;
for x2'Alignment use 4;
for x2 use record
   l1 at  0 range  0 .. 0;
   l2 at  0 range  1 .. 64;
```

```
      l3 at 12 range  0 .. 31;
      l4 at 16 range  0 .. 0;
      l5 at 16 range  1 .. 13;
      l6 at 18 range  0 .. 71;
   end record;
```

The Size values are actually the Object_Size, i.e. the default size that will be allocated for objects of the type. The ?? size for type r indicates that we have a variant record, and the actual size of objects will depend on the discriminant value.

The Alignment values show the actual alignment chosen by the compiler for each record or array type.

The record representation clause for type r shows where all fields are placed, including the compiler generated tag field (whose location cannot be controlled by the programmer).

The record representation clause for the type extension r2 shows all the fields present, including the parent field, which is a copy of the fields of the parent type of r2, i.e. r1.

The component size and size clauses for types rb1 and rb2 show the exact effect of pragma `Pack` on these arrays, and the record representation clause for type x2 shows how pragma `Pack` affects this record type.

In some cases, it may be useful to cut and paste the representation clauses generated by the compiler into the original source to fix and guarantee the actual representation to be used.

7 Standard Library Routines

The Ada 95 Reference Manual contains in Annex A a full description of an extensive set of standard library routines that can be used in any Ada program, and which must be provided by all Ada compilers. They are analogous to the standard C library used by C programs.

GNAT implements all of the facilities described in annex A, and for most purposes the description in the Ada 95 reference manual, or appropriate Ada text book, will be sufficient for making use of these facilities.

In the case of the input-output facilities, See Chapter 8 [The Implementation of Standard I/O], page 133, gives details on exactly how GNAT interfaces to the file system. For the remaining packages, the Ada 95 reference manual should be sufficient. The following is a list of the packages included, together with a brief description of the functionality that is provided.

For completeness, references are included to other predefined library routines defined in other sections of the Ada 95 reference manual (these are cross-indexed from annex A).

`Ada (A.2)` This is a parent package for all the standard library packages. It is usually included implicitly in your program, and itself contains no useful data or routines.

`Ada.Calendar (9.6)`
`Calendar` provides time of day access, and routines for manipulating times and durations.

`Ada.Characters (A.3.1)`
This is a dummy parent package that contains no useful entities

`Ada.Characters.Handling (A.3.2)`
This package provides some basic character handling capabilities, including classification functions for classes of characters (e.g. test for letters, or digits).

`Ada.Characters.Latin_1 (A.3.3)`
This package includes a complete set of definitions of the characters that appear in type CHARACTER. It is useful for writing programs that will run in international environments. For example, if you want an upper case E with an acute accent in a string, it is often better to use the definition of `UC_E_Acute` in this package. Then your program will print in an understandable manner even if your environment does not support these extended characters.

`Ada.Command_Line (A.15)`
This package provides access to the command line parameters and the name of the current program (analogous to the use of `argc` and `argv` in C), and also allows the exit status for the program to be set in a system-independent manner.

`Ada.Decimal (F.2)`
This package provides constants describing the range of decimal numbers implemented, and also a decimal divide routine (analogous to the COBOL verb DIVIDE .. GIVING .. REMAINDER ..)

`Ada.Direct_IO (A.8.4)`
: This package provides input-output using a model of a set of records of fixed-length, containing an arbitrary definite Ada type, indexed by an integer record number.

`Ada.Dynamic_Priorities (D.5)`
: This package allows the priorities of a task to be adjusted dynamically as the task is running.

`Ada.Exceptions (11.4.1)`
: This package provides additional information on exceptions, and also contains facilities for treating exceptions as data objects, and raising exceptions with associated messages.

`Ada.Finalization (7.6)`
: This package contains the declarations and subprograms to support the use of controlled types, providing for automatic initialization and finalization (analogous to the constructors and destructors of C++)

`Ada.Interrupts (C.3.2)`
: This package provides facilities for interfacing to interrupts, which includes the set of signals or conditions that can be raised and recognized as interrupts.

`Ada.Interrupts.Names (C.3.2)`
: This package provides the set of interrupt names (actually signal or condition names) that can be handled by GNAT.

`Ada.IO_Exceptions (A.13)`
: This package defines the set of exceptions that can be raised by use of the standard IO packages.

`Ada.Numerics`
: This package contains some standard constants and exceptions used throughout the numerics packages. Note that the constants pi and e are defined here, and it is better to use these definitions than rolling your own.

`Ada.Numerics.Complex_Elementary_Functions`
: Provides the implementation of standard elementary functions (such as log and trigonometric functions) operating on complex numbers using the standard `Float` and the `Complex` and `Imaginary` types created by the package `Numerics.Complex_Types`.

`Ada.Numerics.Complex_Types`
: This is a predefined instantiation of `Numerics.Generic_Complex_Types` using `Standard.Float` to build the type `Complex` and `Imaginary`.

`Ada.Numerics.Discrete_Random`
: This package provides a random number generator suitable for generating random integer values from a specified range.

`Ada.Numerics.Float_Random`
: This package provides a random number generator suitable for generating uniformly distributed floating point values.

`Ada.Numerics.Generic_Complex_Elementary_Functions`
: This is a generic version of the package that provides the implementation of standard elementary functions (such as log and trigonometric functions) for an arbitrary complex type.

  The following predefined instantiations of this package are provided:

  `Short_Float`
  : `Ada.Numerics.Short_Complex_Elementary_Functions`

  `Float`
  : `Ada.Numerics.Complex_Elementary_Functions`

  `Long_Float`
  : `Ada.Numerics. Long_Complex_Elementary_Functions`

`Ada.Numerics.Generic_Complex_Types`
: This is a generic package that allows the creation of complex types, with associated complex arithmetic operations.

  The following predefined instantiations of this package exist

  `Short_Float`
  : `Ada.Numerics.Short_Complex_Complex_Types`

  `Float`
  : `Ada.Numerics.Complex_Complex_Types`

  `Long_Float`
  : `Ada.Numerics.Long_Complex_Complex_Types`

`Ada.Numerics.Generic_Elementary_Functions`
: This is a generic package that provides the implementation of standard elementary functions (such as log an trigonometric functions) for an arbitrary float type.

  The following predefined instantiations of this package exist

  `Short_Float`
  : `Ada.Numerics.Short_Elementary_Functions`

  `Float`
  : `Ada.Numerics.Elementary_Functions`

  `Long_Float`
  : `Ada.Numerics.Long_Elementary_Functions`

`Ada.Real_Time (D.8)`
: This package provides facilities similar to those of `Calendar`, but operating with a finer clock suitable for real time control. Note that annex D requires that there be no backward clock jumps, and GNAT generally guarantees this behavior, but of course if the external clock on which the GNAT runtime depends is deliberately reset by some external event, then such a backward jump may occur.

`Ada.Sequential_IO (A.8.1)`
: This package provides input-output facilities for sequential files, which can contain a sequence of values of a single type, which can be any Ada type, including indefinite (unconstrained) types.

`Ada.Storage_IO (A.9)`
: This package provides a facility for mapping arbitrary Ada types to and from a storage buffer. It is primarily intended for the creation of new IO packages.

`Ada.Streams (13.13.1)`
: This is a generic package that provides the basic support for the concept of streams as used by the stream attributes (`Input`, `Output`, `Read` and `Write`).

`Ada.Streams.Stream_IO (A.12.1)`
: This package is a specialization of the type `Streams` defined in package `Streams` together with a set of operations providing Stream_IO capability. The Stream_IO model permits both random and sequential access to a file which can contain an arbitrary set of values of one or more Ada types.

`Ada.Strings (A.4.1)`
: This package provides some basic constants used by the string handling packages.

`Ada.Strings.Bounded (A.4.4)`
: This package provides facilities for handling variable length strings. The bounded model requires a maximum length. It is thus somewhat more limited than the unbounded model, but avoids the use of dynamic allocation or finalization.

`Ada.Strings.Fixed (A.4.3)`
: This package provides facilities for handling fixed length strings.

`Ada.Strings.Maps (A.4.2)`
: This package provides facilities for handling character mappings and arbitrarily defined subsets of characters. For instance it is useful in defining specialized translation tables.

`Ada.Strings.Maps.Constants (A.4.6)`
: This package provides a standard set of predefined mappings and predefined character sets. For example, the standard upper to lower case conversion table is found in this package. Note that upper to lower case conversion is non-trivial if you want to take the entire set of characters, including extended characters like E with an acute accent, into account. You should use the mappings in this package (rather than adding 32 yourself) to do case mappings.

`Ada.Strings.Unbounded (A.4.5)`
: This package provides facilities for handling variable length strings. The unbounded model allows arbitrary length strings, but requires the use of dynamic allocation and finalization.

`Ada.Strings.Wide_Bounded (A.4.7)`
`Ada.Strings.Wide_Fixed (A.4.7)`
`Ada.Strings.Wide_Maps (A.4.7)`
`Ada.Strings.Wide_Maps.Constants (A.4.7)`
`Ada.Strings.Wide_Unbounded (A.4.7)`
: These package provide analogous capabilities to the corresponding packages without '`Wide_`' in the name, but operate with the types `Wide_String` and `Wide_Character` instead of `String` and `Character`.

`Ada.Synchronous_Task_Control (D.10)`
: This package provides some standard facilities for controlling task communication in a synchronous manner.

`Ada.Tags` This package contains definitions for manipulation of the tags of tagged values.

`Ada.Task_Attributes`
: This package provides the capability of associating arbitrary task-specific data with separate tasks.

`Ada.Text_IO`
: This package provides basic text input-output capabilities for character, string and numeric data. The subpackages of this package are listed next.

`Ada.Text_IO.Decimal_IO`
: Provides input-output facilities for decimal fixed-point types

`Ada.Text_IO.Enumeration_IO`
: Provides input-output facilities for enumeration types.

`Ada.Text_IO.Fixed_IO`
: Provides input-output facilities for ordinary fixed-point types.

`Ada.Text_IO.Float_IO`
: Provides input-output facilities for float types. The following predefined instantiations of this generic package are available:

  `Short_Float`
  : `Short_Float_Text_IO`

  `Float` `Float_Text_IO`

  `Long_Float`
  : `Long_Float_Text_IO`

`Ada.Text_IO.Integer_IO`
: Provides input-output facilities for integer types. The following predefined instantiations of this generic package are available:

  `Short_Short_Integer`
  : `Ada.Short_Short_Integer_Text_IO`

  `Short_Integer`
  : `Ada.Short_Integer_Text_IO`

  `Integer` `Ada.Integer_Text_IO`

  `Long_Integer`
  : `Ada.Long_Integer_Text_IO`

  `Long_Long_Integer`
  : `Ada.Long_Long_Integer_Text_IO`

`Ada.Text_IO.Modular_IO`
: Provides input-output facilities for modular (unsigned) types

`Ada.Text_IO.Complex_IO (G.1.3)`
: This package provides basic text input-output capabilities for complex data.

`Ada.Text_IO.Editing (F.3.3)`
: This package contains routines for edited output, analogous to the use of pictures in COBOL. The picture formats used by this package are a close copy of the facility in COBOL.

`Ada.Text_IO.Text_Streams (A.12.2)`
: This package provides a facility that allows Text_IO files to be treated as streams, so that the stream attributes can be used for writing arbitrary data, including binary data, to Text_IO files.

`Ada.Unchecked_Conversion (13.9)`
: This generic package allows arbitrary conversion from one type to another of the same size, providing for breaking the type safety in special circumstances.

  If the types have the same Size (more accurately the same Value_Size), then the effect is simply to transfer the bits from the source to the target type without any modification. This usage is well defined, and for simple types whose representation is typically the same across all implementations, gives a portable method of performing such conversions.

  If the types do not have the same size, then the result is implementation defined, and thus may be non-portable. The following describes how GNAT handles such unchecked conversion cases.

  If the types are of different sizes, and are both discrete types, then the effect is of a normal type conversion without any constraint checking. In particular if the result type has a larger size, the result will be zero or sign extended. If the result type has a smaller size, the result will be truncated by ignoring high order bits.

  If the types are of different sizes, and are not both discrete types, then the conversion works as though pointers were created to the source and target, and the pointer value is converted. The effect is that bits are copied from successive low order storage units and bits of the source up to the length of the target type.

  A warning is issued if the lengths differ, since the effect in this case is implementation dependent, and the above behavior may not match that of some other compiler.

  A pointer to one type may be converted to a pointer to another type using unchecked conversion. The only case in which the effect is undefined is when one or both pointers are pointers to unconstrained array types. In this case, the bounds information may get incorrectly transferred, and in particular, GNAT uses double size pointers for such types, and it is meaningless to convert between such pointer types. GNAT will issue a warning if the alignment of the target designated type is more strict than the alignment of the source designated type (since the result may be unaligned in this case).

  A pointer other than a pointer to an unconstrained array type may be converted to and from System.Address. Such usage is common in Ada 83 programs, but note that Ada.Address_To_Access_Conversions is the preferred method of performing such conversions in Ada 95. Neither unchecked conversion nor

Ada.Address_To_Access_Conversions should be used in conjunction with pointers to unconstrained objects, since the bounds information cannot be handled correctly in this case.

`Ada.Unchecked_Deallocation (13.11.2)`
This generic package allows explicit freeing of storage previously allocated by use of an allocator.

`Ada.Wide_Text_IO (A.11)`
This package is similar to `Ada.Text_IO`, except that the external file supports wide character representations, and the internal types are `Wide_Character` and `Wide_String` instead of `Character` and `String`. It contains generic subpackages listed next.

`Ada.Wide_Text_IO.Decimal_IO`
Provides input-output facilities for decimal fixed-point types

`Ada.Wide_Text_IO.Enumeration_IO`
Provides input-output facilities for enumeration types.

`Ada.Wide_Text_IO.Fixed_IO`
Provides input-output facilities for ordinary fixed-point types.

`Ada.Wide_Text_IO.Float_IO`
Provides input-output facilities for float types. The following predefined instantiations of this generic package are available:

```
Short_Float
          Short_Float_Wide_Text_IO

Float     Float_Wide_Text_IO

Long_Float
          Long_Float_Wide_Text_IO
```

`Ada.Wide_Text_IO.Integer_IO`
Provides input-output facilities for integer types. The following predefined instantiations of this generic package are available:

```
Short_Short_Integer
          Ada.Short_Short_Integer_Wide_Text_IO

Short_Integer
          Ada.Short_Integer_Wide_Text_IO

Integer   Ada.Integer_Wide_Text_IO

Long_Integer
          Ada.Long_Integer_Wide_Text_IO

Long_Long_Integer
          Ada.Long_Long_Integer_Wide_Text_IO
```

`Ada.Wide_Text_IO.Modular_IO`
Provides input-output facilities for modular (unsigned) types

`Ada.Wide_Text_IO.Complex_IO (G.1.3)`
: This package is similar to `Ada.Text_IO.Complex_IO`, except that the external file supports wide character representations.

`Ada.Wide_Text_IO.Editing (F.3.4)`
: This package is similar to `Ada.Text_IO.Editing`, except that the types are `Wide_Character` and `Wide_String` instead of `Character` and `String`.

`Ada.Wide_Text_IO.Streams (A.12.3)`
: This package is similar to `Ada.Text_IO.Streams`, except that the types are `Wide_Character` and `Wide_String` instead of `Character` and `String`.

8 The Implementation of Standard I/O

GNAT implements all the required input-output facilities described in A.6 through A.14. These sections of the Ada 95 reference manual describe the required behavior of these packages from the Ada point of view, and if you are writing a portable Ada program that does not need to know the exact manner in which Ada maps to the outside world when it comes to reading or writing external files, then you do not need to read this chapter. As long as your files are all regular files (not pipes or devices), and as long as you write and read the files only from Ada, the description in the Ada 95 reference manual is sufficient.

However, if you want to do input-output to pipes or other devices, such as the keyboard or screen, or if the files you are dealing with are either generated by some other language, or to be read by some other language, then you need to know more about the details of how the GNAT implementation of these input-output facilities behaves.

In this chapter we give a detailed description of exactly how GNAT interfaces to the file system. As always, the sources of the system are available to you for answering questions at an even more detailed level, but for most purposes the information in this chapter will suffice.

Another reason that you may need to know more about how input-output is implemented arises when you have a program written in mixed languages where, for example, files are shared between the C and Ada sections of the same program. GNAT provides some additional facilities, in the form of additional child library packages, that facilitate this sharing, and these additional facilities are also described in this chapter.

8.1 Standard I/O Packages

The Standard I/O packages described in Annex A for

- Ada.Text_IO
- Ada.Text_IO.Complex_IO
- Ada.Text_IO.Text_Streams,
- Ada.Wide_Text_IO
- Ada.Wide_Text_IO.Complex_IO,
- Ada.Wide_Text_IO.Text_Streams
- Ada.Stream_IO
- Ada.Sequential_IO
- Ada.Direct_IO

are implemented using the C library streams facility; where

- All files are opened using `fopen`.
- All input/output operations use `fread`/`fwrite`.

There is no internal buffering of any kind at the Ada library level. The only buffering is that provided at the system level in the implementation of the C library routines that support streams. This facilitates shared use of these streams by mixed language programs.

8.2 FORM Strings

The format of a FORM string in GNAT is:

```
"keyword=value,keyword=value,...,keyword=value"
```

where letters may be in upper or lower case, and there are no spaces between values. The order of the entries is not important. Currently there are two keywords defined.

```
SHARED=[YES|NO]
WCEM=[n|h|u|s\e]
```

The use of these parameters is described later in this section.

8.3 Direct_IO

Direct_IO can only be instantiated for definite types. This is a restriction of the Ada language, which means that the records are fixed length (the length being determined by *type*'Size, rounded up to the next storage unit boundary if necessary).

The records of a Direct_IO file are simply written to the file in index sequence, with the first record starting at offset zero, and subsequent records following. There is no control information of any kind. For example, if 32-bit integers are being written, each record takes 4-bytes, so the record at index K starts at offset $(K-1)*4$.

There is no limit on the size of Direct_IO files, they are expanded as necessary to accommodate whatever records are written to the file.

8.4 Sequential_IO

Sequential_IO may be instantiated with either a definite (constrained) or indefinite (unconstrained) type.

For the definite type case, the elements written to the file are simply the memory images of the data values with no control information of any kind. The resulting file should be read using the same type, no validity checking is performed on input.

For the indefinite type case, the elements written consist of two parts. First is the size of the data item, written as the memory image of a `Interfaces.C.size_t` value, followed by the memory image of the data value. The resulting file can only be read using the same (unconstrained) type. Normal assignment checks are performed on these read operations, and if these checks fail, `Data_Error` is raised. In particular, in the array case, the lengths must match, and in the variant record case, if the variable for a particular read operation is constrained, the discriminants must match.

Note that it is not possible to use Sequential_IO to write variable length array items, and then read the data back into different length arrays. For example, the following will raise `Data_Error`:

```
package IO is new Sequential_IO (String);
F : IO.File_Type;
S : String (1..4);
...
IO.Create (F)
IO.Write (F, "hello!")
```

```
IO.Reset (F, Mode=>In_File);
IO.Read (F, S);
Put_Line (S);
```

On some Ada implementations, this will print '`hell`', but the program is clearly incorrect, since there is only one element in the file, and that element is the string '`hello!`'.

In Ada 95, this kind of behavior can be legitimately achieved using Stream_IO, and this is the preferred mechanism. In particular, the above program fragment rewritten to use Stream_IO will work correctly.

8.5 Text_IO

Text_IO files consist of a stream of characters containing the following special control characters:

```
LF (line feed, 16#0A#) Line Mark
FF (form feed, 16#0C#) Page Mark
```

A canonical Text_IO file is defined as one in which the following conditions are met:

- The character `LF` is used only as a line mark, i.e. to mark the end of the line.
- The character `FF` is used only as a page mark, i.e. to mark the end of a page and consequently can appear only immediately following a `LF` (line mark) character.
- The file ends with either `LF` (line mark) or `LF-FF` (line mark, page mark). In the former case, the page mark is implicitly assumed to be present.

A file written using Text_IO will be in canonical form provided that no explicit `LF` or `FF` characters are written using `Put` or `Put_Line`. There will be no `FF` character at the end of the file unless an explicit `New_Page` operation was performed before closing the file.

A canonical Text_IO file that is a regular file, i.e. not a device or a pipe, can be read using any of the routines in Text_IO. The semantics in this case will be exactly as defined in the Ada 95 reference manual and all the routines in Text_IO are fully implemented.

A text file that does not meet the requirements for a canonical Text_IO file has one of the following:

- The file contains `FF` characters not immediately following a `LF` character.
- The file contains `LF` or `FF` characters written by `Put` or `Put_Line`, which are not logically considered to be line marks or page marks.
- The file ends in a character other than `LF` or `FF`, i.e. there is no explicit line mark or page mark at the end of the file.

Text_IO can be used to read such non-standard text files but subprograms to do with line or page numbers do not have defined meanings. In particular, a `FF` character that does not follow a `LF` character may or may not be treated as a page mark from the point of view of page and line numbering. Every `LF` character is considered to end a line, and there is an implied `LF` character at the end of the file.

8.5.1 Stream Pointer Positioning

`Ada.Text_IO` has a definition of current position for a file that is being read. No internal buffering occurs in Text_IO, and usually the physical position in the stream used to implement the file corresponds to this logical position defined by Text_IO. There are two exceptions:

- After a call to `End_Of_Page` that returns `True`, the stream is positioned past the `LF` (line mark) that precedes the page mark. Text_IO maintains an internal flag so that subsequent read operations properly handle the logical position which is unchanged by the `End_Of_Page` call.
- After a call to `End_Of_File` that returns `True`, if the Text_IO file was positioned before the line mark at the end of file before the call, then the logical position is unchanged, but the stream is physically positioned right at the end of file (past the line mark, and past a possible page mark following the line mark. Again Text_IO maintains internal flags so that subsequent read operations properly handle the logical position.

These discrepancies have no effect on the observable behavior of Text_IO, but if a single Ada stream is shared between a C program and Ada program, or shared (using '`shared=yes`' in the form string) between two Ada files, then the difference may be observable in some situations.

8.5.2 Reading and Writing Non-Regular Files

A non-regular file is a device (such as a keyboard), or a pipe. Text_IO can be used for reading and writing. Writing is not affected and the sequence of characters output is identical to the normal file case, but for reading, the behavior of Text_IO is modified to avoid undesirable look-ahead as follows:

An input file that is not a regular file is considered to have no page marks. Any `Ascii.FF` characters (the character normally used for a page mark) appearing in the file are considered to be data characters. In particular:

- `Get_Line` and `Skip_Line` do not test for a page mark following a line mark. If a page mark appears, it will be treated as a data character.
- This avoids the need to wait for an extra character to be typed or entered from the pipe to complete one of these operations.
- `End_Of_Page` always returns `False`
- `End_Of_File` will return `False` if there is a page mark at the end of the file.

Output to non-regular files is the same as for regular files. Page marks may be written to non-regular files using `New_Page`, but as noted above they will not be treated as page marks on input if the output is piped to another Ada program.

Another important discrepancy when reading non-regular files is that the end of file indication is not "sticky". If an end of file is entered, e.g. by pressing the ⟨EOT⟩ key, then end of file is signalled once (i.e. the test `End_Of_File` will yield `True`, or a read will raise `End_Error`), but then reading can resume to read data past that end of file indication, until another end of file indication is entered.

8.5.3 Get_Immediate

Get_Immediate returns the next character (including control characters) from the input file. In particular, Get_Immediate will return LF or FF characters used as line marks or page marks. Such operations leave the file positioned past the control character, and it is thus not treated as having its normal function. This means that page, line and column counts after this kind of Get_Immediate call are set as though the mark did not occur. In the case where a Get_Immediate leaves the file positioned between the line mark and page mark (which is not normally possible), it is undefined whether the FF character will be treated as a page mark.

8.5.4 Treating Text_IO Files as Streams

The package `Text_IO.Streams` allows a Text_IO file to be treated as a stream. Data written to a Text_IO file in this stream mode is binary data. If this binary data contains bytes 16#0A# (`LF`) or 16#0C# (`FF`), the resulting file may have non-standard format. Similarly if read operations are used to read from a Text_IO file treated as a stream, then `LF` and `FF` characters may be skipped and the effect is similar to that described above for `Get_Immediate`.

8.5.5 Text_IO Extensions

A package GNAT.IO_Aux in the GNAT library provides some useful extensions to the standard `Text_IO` package:

- function File_Exists (Name : String) return Boolean; Determines if a file of the given name exists and can be successfully opened (without actually performing the open operation).
- function Get_Line return String; Reads a string from the standard input file. The value returned is exactly the length of the line that was read.
- function Get_Line (File : Ada.Text_IO.File_Type) return String; Similar, except that the parameter File specifies the file from which the string is to be read.

8.5.6 Text_IO Facilities for Unbounded Strings

The package `Ada.Strings.Unbounded.Text_IO` in library files `a-suteio.ads/adb` contains some GNAT-specific subprograms useful for Text_IO operations on unbounded strings:

- function Get_Line (File : File_Type) return Unbounded_String; Reads a line from the specified file and returns the result as an unbounded string.
- procedure Put (File : File_Type; U : Unbounded_String); Writes the value of the given unbounded string to the specified file Similar to the effect of `Put (To_String (U))` except that an extra copy is avoided.
- procedure Put_Line (File : File_Type; U : Unbounded_String); Writes the value of the given unbounded string to the specified file, followed by a `New_Line`. Similar to the effect of `Put_Line (To_String (U))` except that an extra copy is avoided.

In the above procedures, `File` is of type `Ada.Text_IO.File_Type` and is optional. If the parameter is omitted, then the standard input or output file is referenced as appropriate.

The package `Ada.Strings.Wide_Unbounded.Wide_Text_IO` in library files '`a-swuwti.ads`' and '`a-swuwti.adb`' provides similar extended `Wide_Text_IO` functionality for unbounded wide strings.

8.6 Wide_Text_IO

`Wide_Text_IO` is similar in most respects to Text_IO, except that both input and output files may contain special sequences that represent wide character values. The encoding scheme for a given file may be specified using a FORM parameter:

```
WCEM=x
```

as part of the FORM string (WCEM = wide character encoding method), where *x* is one of the following characters

'`h`' Hex ESC encoding

'`u`' Upper half encoding

'`s`' Shift-JIS encoding

'`e`' EUC Encoding

'`8`' UTF-8 encoding

'`b`' Brackets encoding

The encoding methods match those that can be used in a source program, but there is no requirement that the encoding method used for the source program be the same as the encoding method used for files, and different files may use different encoding methods.

The default encoding method for the standard files, and for opened files for which no WCEM parameter is given in the FORM string matches the wide character encoding specified for the main program (the default being brackets encoding if no coding method was specified with -gnatW).

Hex Coding

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

```
ESC a b c d
```

where *a*, *b*, *c*, *d* are the four hexadecimal characters (using upper case 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 (i.e. a is in the range 8-F) is represented as two bytes 16#ab# and 16#cd#. The second byte may never be a format control character, but is not required to be in the upper half. This method can be also used for shift-JIS or EUC where the internal coding matches the external coding.

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

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

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

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

  where the xxx bits correspond to the left-padded bits of the 16-bit character value. Note that all lower half ASCII characters are represented as ASCII bytes and all upper half characters and other wide characters are represented as sequences of upper-half (The full UTF-8 scheme allows for encoding 31-bit characters as 6-byte sequences, but in this implementation, all UTF-8 sequences of four or more bytes length will raise a Constraint_Error, as will all invalid UTF-8 sequences.)

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#`. This scheme is compatible with use of the full Wide_Character set. On input, brackets coding can also be used for upper half characters, e.g. `["C1"]` for lower case a. However, on output, brackets notation is only used for wide characters with a code greater than `16#FF#`.

For the coding schemes other than Hex and Brackets encoding, not all wide character values can be represented. An attempt to output a character that cannot be represented using the encoding scheme for the file causes Constraint_Error to be raised. An invalid wide character sequence on input also causes Constraint_Error to be raised.

8.6.1 Stream Pointer Positioning

`Ada.Wide_Text_IO` is similar to `Ada.Text_IO` in its handling of stream pointer positioning (see Section 8.5 [Text_IO], page 135). There is one additional case:

If `Ada.Wide_Text_IO.Look_Ahead` reads a character outside the normal lower ASCII set (i.e. a character in the range:

```
Wide_Character'Val (16#0080#) .. Wide_Character'Val (16#FFFF#)
```

then although the logical position of the file pointer is unchanged by the `Look_Ahead` call, the stream is physically positioned past the wide character sequence. Again this is to avoid the need for buffering or backup, and all `Wide_Text_IO` routines check the internal indication that this situation has occurred so that this is not visible to a normal program using `Wide_Text_IO`. However, this discrepancy can be observed if the wide text file shares a stream with another file.

8.6.2 Reading and Writing Non-Regular Files

As in the case of Text_IO, when a non-regular file is read, it is assumed that the file contains no page marks (any form characters are treated as data characters), and `End_Of_Page` always returns `False`. Similarly, the end of file indication is not sticky, so it is possible to read beyond an end of file.

8.7 Stream_IO

A stream file is a sequence of bytes, where individual elements are written to the file as described in the Ada 95 reference manual. The type `Stream_Element` is simply a byte. There are two ways to read or write a stream file.

- The operations `Read` and `Write` directly read or write a sequence of stream elements with no control information.
- The stream attributes applied to a stream file transfer data in the manner described for stream attributes.

8.8 Shared Files

Section A.14 of the Ada 95 Reference Manual allows implementations to provide a wide variety of behavior if an attempt is made to access the same external file with two or more internal files.

To provide a full range of functionality, while at the same time minimizing the problems of portability caused by this implementation dependence, GNAT handles file sharing as follows:

- In the absence of a '`shared=xxx`' form parameter, an attempt to open two or more files with the same full name is considered an error and is not supported. The exception `Use_Error` will be raised. Note that a file that is not explicitly closed by the program remains open until the program terminates.
- If the form parameter '`shared=no`' appears in the form string, the file can be opened or created with its own separate stream identifier, regardless of whether other files sharing the same external file are opened. The exact effect depends on how the C stream routines handle multiple accesses to the same external files using separate streams.

- If the form parameter '`shared=yes`' appears in the form string for each of two or more files opened using the same full name, the same stream is shared between these files, and the semantics are as described in Ada 95 Reference Manual, Section A.14.

When a program that opens multiple files with the same name is ported from another Ada compiler to GNAT, the effect will be that `Use_Error` is raised.

The documentation of the original compiler and the documentation of the program should then be examined to determine if file sharing was expected, and '`shared=xxx`' parameters added to `Open` and `Create` calls as required.

When a program is ported from GNAT to some other Ada compiler, no special attention is required unless the '`shared=xxx`' form parameter is used in the program. In this case, you must examine the documentation of the new compiler to see if it supports the required file sharing semantics, and form strings modified appropriately. Of course it may be the case that the program cannot be ported if the target compiler does not support the required functionality. The best approach in writing portable code is to avoid file sharing (and hence the use of the '`shared=xxx`' parameter in the form string) completely.

One common use of file sharing in Ada 83 is the use of instantiations of Sequential_IO on the same file with different types, to achieve heterogeneous input-output. Although this approach will work in GNAT if '`shared=yes`' is specified, it is preferable in Ada 95 to use Stream_IO for this purpose (using the stream attributes)

8.9 Open Modes

`Open` and `Create` calls result in a call to `fopen` using the mode shown in Table 6.1

Table 6-1 `Open` and `Create` Call Modes

	OPEN	**CREATE**
`Append_File`	`"r+"`	`"w+"`
`In_File`	`"r"`	`"w+"`
`Out_File (Direct_IO)`	`"r+"`	`"w"`
`Out_File (all other cases)`	`"w"`	`"w"`
`Inout_File`	`"r+"`	`"w+"`

If text file translation is required, then either '`b`' or '`t`' is added to the mode, depending on the setting of Text. Text file translation refers to the mapping of CR/LF sequences in an external file to LF characters internally. This mapping only occurs in DOS and DOS-like systems, and is not relevant to other systems.

A special case occurs with Stream_IO. As shown in the above table, the file is initially opened in '`r`' or '`w`' mode for the `In_File` and `Out_File` cases. If a `Set_Mode` operation subsequently requires switching from reading to writing or vice-versa, then the file is reopened in '`r+`' mode to permit the required operation.

8.10 Operations on C Streams

The package `Interfaces.C_Streams` provides an Ada program with direct access to the C library functions for operations on C streams:

```
package Interfaces.C_Streams is
  -- Note: the reason we do not use the types that are in
  -- Interfaces.C is that we want to avoid dragging in the
  -- code in this unit if possible.
  subtype chars is System.Address;
  -- Pointer to null-terminated array of characters
  subtype FILEs is System.Address;
  -- Corresponds to the C type FILE*
  subtype voids is System.Address;
  -- Corresponds to the C type void*
  subtype int is Integer;
  subtype long is Long_Integer;
  -- Note: the above types are subtypes deliberately, and it
  -- is part of this spec that the above correspondences are
  -- guaranteed.  This means that it is legitimate to, for
  -- example, use Integer instead of int.  We provide these
  -- synonyms for clarity, but in some cases it may be
  -- convenient to use the underlying types (for example to
  -- avoid an unnecessary dependency of a spec on the spec
  -- of this unit).
  type size_t is mod 2 ** Standard'Address_Size;
  NULL_Stream : constant FILEs;
  -- Value returned (NULL in C) to indicate an
  -- fdopen/fopen/tmpfile error
  ---------------------------------
  -- Constants Defined in stdio.h --
  ---------------------------------
  EOF : constant int;
  -- Used by a number of routines to indicate error or
  -- end of file
  IOFBF : constant int;
  IOLBF : constant int;
  IONBF : constant int;
  -- Used to indicate buffering mode for setvbuf call
  SEEK_CUR : constant int;
  SEEK_END : constant int;
  SEEK_SET : constant int;
  -- Used to indicate origin for fseek call
  function stdin return FILEs;
  function stdout return FILEs;
  function stderr return FILEs;
  -- Streams associated with standard files
  --------------------------
  -- Standard C functions --
  --------------------------
  -- The functions selected below are ones that are
  -- available in DOS, OS/2, UNIX and Xenix (but not
  -- necessarily in ANSI C).  These are very thin interfaces
  -- which copy exactly the C headers.  For more
  -- documentation on these functions, see the Microsoft C
  -- "Run-Time Library Reference" (Microsoft Press, 1990,
  -- ISBN 1-55615-225-6), which includes useful information
  -- on system compatibility.
  procedure clearerr (stream : FILEs);
  function fclose (stream : FILEs) return int;
  function fdopen (handle : int; mode : chars) return FILEs;
  function feof (stream : FILEs) return int;
  function ferror (stream : FILEs) return int;
```

```
function fflush (stream : FILEs) return int;
function fgetc (stream : FILEs) return int;
function fgets (strng : chars; n : int; stream : FILEs)
    return chars;
function fileno (stream : FILEs) return int;
function fopen (filename : chars; Mode : chars)
    return FILEs;
-- Note: to maintain target independence, use
-- text_translation_required, a boolean variable defined in
-- a-sysdep.c to deal with the target dependent text
-- translation requirement.  If this variable is set,
-- then  b/t should be appended to the standard mode
-- argument to set the text translation mode off or on
-- as required.
function fputc (C : int; stream : FILEs) return int;
function fputs (Strng : chars; Stream : FILEs) return int;
function fread
   (buffer : voids;
    size : size_t;
    count : size_t;
    stream : FILEs)
    return size_t;
function freopen
   (filename : chars;
    mode : chars;
    stream : FILEs)
    return FILEs;
function fseek
   (stream : FILEs;
    offset : long;
    origin : int)
    return int;
function ftell (stream : FILEs) return long;
function fwrite
   (buffer : voids;
    size : size_t;
    count : size_t;
    stream : FILEs)
    return size_t;
function isatty (handle : int) return int;
procedure mktemp (template : chars);
-- The return value (which is just a pointer to template)
-- is discarded
procedure rewind (stream : FILEs);
function rmtmp return int;
function setvbuf
   (stream : FILEs;
    buffer : chars;
    mode : int;
    size : size_t)
    return int;

function tmpfile return FILEs;
function ungetc (c : int; stream : FILEs) return int;
function unlink (filename : chars) return int;
---------------------
-- Extra functions --
---------------------
```

```
   -- These functions supply slightly thicker bindings than
   -- those above.  They are derived from functions in the
   -- C Run-Time Library, but may do a bit more work than
   -- just directly calling one of the Library functions.
   function is_regular_file (handle : int) return int;
   -- Tests if given handle is for a regular file (result 1)
   -- or for a non-regular file (pipe or device, result 0).
   --------------------------------
   -- Control of Text/Binary Mode --
   --------------------------------
   -- If text_translation_required is true, then the following
   -- functions may be used to dynamically switch a file from
   -- binary to text mode or vice versa.  These functions have
   -- no effect if text_translation_required is false (i.e.  in
   -- normal UNIX mode).  Use fileno to get a stream handle.
   procedure set_binary_mode (handle : int);
   procedure set_text_mode (handle : int);
   ----------------------------
   -- Full Path Name support --
   ----------------------------
   procedure full_name (nam : chars; buffer : chars);
   -- Given a NUL terminated string representing a file
   -- name, returns in buffer a NUL terminated string
   -- representing the full path name for the file name.
   -- On systems where it is relevant the   drive is also
   -- part of the full path name.  It is the responsibility
   -- of the caller to pass an actual parameter for buffer
   -- that is big enough for any full path name.  Use
   -- max_path_len given below as the size of buffer.
   max_path_len : integer;
   -- Maximum length of an allowable full path name on the
   -- system, including a terminating NUL character.
end Interfaces.C_Streams;
```

8.11 Interfacing to C Streams

The packages in this section permit interfacing Ada files to C Stream operations.

```
with Interfaces.C_Streams;
package Ada.Sequential_IO.C_Streams is
   function C_Stream (F : File_Type)
      return Interfaces.C_Streams.FILEs;
   procedure Open
     (File : in out File_Type;
      Mode : in File_Mode;
      C_Stream : in Interfaces.C_Streams.FILEs;
      Form : in String := "");
end Ada.Sequential_IO.C_Streams;

with Interfaces.C_Streams;
package Ada.Direct_IO.C_Streams is
   function C_Stream (F : File_Type)
      return Interfaces.C_Streams.FILEs;
   procedure Open
     (File : in out File_Type;
      Mode : in File_Mode;
      C_Stream : in Interfaces.C_Streams.FILEs;
      Form : in String := "");
```

```
end Ada.Direct_IO.C_Streams;

with Interfaces.C_Streams;
package Ada.Text_IO.C_Streams is
   function C_Stream (F : File_Type)
      return Interfaces.C_Streams.FILEs;
   procedure Open
     (File : in out File_Type;
      Mode : in File_Mode;
      C_Stream : in Interfaces.C_Streams.FILEs;
      Form : in String := "");
end Ada.Text_IO.C_Streams;

with Interfaces.C_Streams;
package Ada.Wide_Text_IO.C_Streams is
   function C_Stream (F : File_Type)
      return Interfaces.C_Streams.FILEs;
   procedure Open
     (File : in out File_Type;
      Mode : in File_Mode;
      C_Stream : in Interfaces.C_Streams.FILEs;
      Form : in String := "");
end Ada.Wide_Text_IO.C_Streams;

with Interfaces.C_Streams;
package Ada.Stream_IO.C_Streams is
   function C_Stream (F : File_Type)
      return Interfaces.C_Streams.FILEs;
   procedure Open
     (File : in out File_Type;
      Mode : in File_Mode;
      C_Stream : in Interfaces.C_Streams.FILEs;
      Form : in String := "");
end Ada.Stream_IO.C_Streams;
```

In each of these five packages, the `C_Stream` function obtains the `FILE` pointer from a currently opened Ada file. It is then possible to use the `Interfaces.C_Streams` package to operate on this stream, or the stream can be passed to a C program which can operate on it directly. Of course the program is responsible for ensuring that only appropriate sequences of operations are executed.

One particular use of relevance to an Ada program is that the `setvbuf` function can be used to control the buffering of the stream used by an Ada file. In the absence of such a call the standard default buffering is used.

The `Open` procedures in these packages open a file giving an existing C Stream instead of a file name. Typically this stream is imported from a C program, allowing an Ada file to operate on an existing C file.

9 The GNAT Library

The GNAT library contains a number of general and special purpose packages. It represents functionality that the GNAT developers have found useful, and which is made available to GNAT users. The packages described here are fully supported, and upwards compatibility will be maintained in future releases, so you can use these facilities with the confidence that the same functionality will be available in future releases.

The chapter here simply gives a brief summary of the facilities available. The full documentation is found in the spec file for the package. The full sources of these library packages, including both spec and body, are provided with all GNAT releases. For example, to find out the full specifications of the SPITBOL pattern matching capability, including a full tutorial and extensive examples, look in the '`g-spipat.ads`' file in the library.

For each entry here, the package name (as it would appear in a `with` clause) is given, followed by the name of the corresponding spec file in parentheses. The packages are children in four hierarchies, `Ada`, `Interfaces`, `System`, and `GNAT`, the latter being a GNAT-specific hierarchy.

Note that an application program should only use packages in one of these four hierarchies if the package is defined in the Ada Reference Manual, or is listed in this section of the GNAT Programmers Reference Manual. All other units should be considered internal implementation units and should not be directly `with`'ed by application code. The use of a `with` statement that references one of these internal implementation units makes an application potentially dependent on changes in versions of GNAT, and will generate a warning message.

9.1 Ada.Characters.Latin_9 ('a-chlat9.ads')

This child of `Ada.Characters` provides a set of definitions corresponding to those in the RM-defined package `Ada.Characters.Latin_1` but with the few modifications required for `Latin-9` The provision of such a package is specifically authorized by the Ada Reference Manual (RM A.3(27)).

9.2 Ada.Characters.Wide_Latin_1 ('a-cwila1.ads')

This child of `Ada.Characters` provides a set of definitions corresponding to those in the RM-defined package `Ada.Characters.Latin_1` but with the types of the constants being `Wide_Character` instead of `Character`. The provision of such a package is specifically authorized by the Ada Reference Manual (RM A.3(27)).

9.3 Ada.Characters.Wide_Latin_9 ('a-cwila1.ads')

This child of `Ada.Characters` provides a set of definitions corresponding to those in the GNAT defined package `Ada.Characters.Latin_9` but with the types of the constants being `Wide_Character` instead of `Character`. The provision of such a package is specifically authorized by the Ada Reference Manual (RM A.3(27)).

9.4 Ada.Command_Line.Remove ('a-colire.ads')

This child of `Ada.Command_Line` provides a mechanism for logically removing arguments from the argument list. Once removed, an argument is not visible to further calls on the subprograms in `Ada.Command_Line` will not see the removed argument.

9.5 Ada.Direct_IO.C_Streams ('a-diocst.ads')

This package provides subprograms that allow interfacing between C streams and `Direct_IO`. The stream identifier can be extracted from a file opened on the Ada side, and an Ada file can be constructed from a stream opened on the C side.

9.6 Ada.Exceptions.Is_Null_Occurrence ('a-einuoc.ads')

This child subprogram provides a way of testing for the null exception occurrence (`Null_Occurrence`) without raising an exception.

9.7 Ada.Sequential_IO.C_Streams ('a-siocst.ads')

This package provides subprograms that allow interfacing between C streams and `Sequential_IO`. The stream identifier can be extracted from a file opened on the Ada side, and an Ada file can be constructed from a stream opened on the C side.

9.8 Ada.Streams.Stream_IO.C_Streams ('a-ssicst.ads')

This package provides subprograms that allow interfacing between C streams and `Stream_IO`. The stream identifier can be extracted from a file opened on the Ada side, and an Ada file can be constructed from a stream opened on the C side.

9.9 Ada.Strings.Unbounded.Text_IO ('a-suteio.ads')

This package provides subprograms for Text_IO for unbounded strings, avoiding the necessity for an intermediate operation with ordinary strings.

9.10 Ada.Strings.Wide_Unbounded.Wide_Text_IO ('a-swuwti.ads')

This package provides subprograms for Text_IO for unbounded wide strings, avoiding the necessity for an intermediate operation with ordinary wide strings.

9.11 Ada.Text_IO.C_Streams ('a-tiocst.ads')

This package provides subprograms that allow interfacing between C streams and `Text_IO`. The stream identifier can be extracted from a file opened on the Ada side, and an Ada file can be constructed from a stream opened on the C side.

9.12 Ada.Wide_Text_IO.C_Streams ('a-wtcstr.ads')

This package provides subprograms that allow interfacing between C streams and `Wide_Text_IO`. The stream identifier can be extracted from a file opened on the Ada side, and an Ada file can be constructed from a stream opened on the C side.

9.13 GNAT.AWK ('g-awk.ads')

Provides AWK-like parsing functions, with an easy interface for parsing one or more files containing formatted data. The file is viewed as a database where each record is a line and a field is a data element in this line.

9.14 GNAT.Bubble_Sort_A ('g-busora.ads')

Provides a general implementation of bubble sort usable for sorting arbitrary data items. Move and comparison procedures are provided by passing access-to-procedure values.

9.15 GNAT.Bubble_Sort_G ('g-busorg.ads')

Similar to `Bubble_Sort_A` except that the move and sorting procedures are provided as generic parameters, this improves efficiency, especially if the procedures can be inlined, at the expense of duplicating code for multiple instantiations.

9.16 GNAT.Calendar ('g-calend.ads')

Extends the facilities provided by `Ada.Calendar` to include handling of days of the week, an extended `Split` and `Time_Of` capability. Also provides conversion of `Ada.Calendar.Time` values to and from the C `timeval` format.

9.17 GNAT.Calendar.Time_IO ('g-catiio.ads')

9.18 GNAT.CRC32 ('g-crc32.ads')

This package implements the CRC-32 algorithm. For a full description of this algorithm you should have a look at: "Computation of Cyclic Redundancy Checks via Table Look-Up", *Communications of the ACM*, Vol. 31 No. 8, pp. 1008-1013, Aug. 1988. Sarwate, D.V.

Provides an extended capability for formatted output of time values with full user control over the format. Modeled on the GNU Date specification.

9.19 GNAT.Case_Util ('g-casuti.ads')

A set of simple routines for handling upper and lower casing of strings without the overhead of the full casing tables in `Ada.Characters.Handling`.

9.20 GNAT.CGI ('g-cgi.ads')

This is a package for interfacing a GNAT program with a Web server via the Common Gateway Interface (CGI). Basically this package parses the CGI parameters, which are a set of key/value pairs sent by the Web server. It builds a table whose index is the key and provides some services to deal with this table.

9.21 GNAT.CGI.Cookie ('g-cgicoo.ads')

This is a package to interface a GNAT program with a Web server via the Common Gateway Interface (CGI). It exports services to deal with Web cookies (piece of information kept in the Web client software).

9.22 GNAT.CGI.Debug ('g-cgideb.ads')

This is a package to help debugging CGI (Common Gateway Interface) programs written in Ada.

9.23 GNAT.Command_Line ('g-comlin.ads')

Provides a high level interface to `Ada.Command_Line` facilities, including the ability to scan for named switches with optional parameters and expand file names using wild card notations.

9.24 GNAT.Current_Exception ('g-curexc.ads')

Provides access to information on the current exception that has been raised without the need for using the Ada-95 exception choice parameter specification syntax. This is particularly useful in simulating typical facilities for obtaining information about exceptions provided by Ada 83 compilers.

9.25 GNAT.Debug_Pools ('g-debpoo.ads')

Provide a debugging storage pools that helps tracking memory corruption problems. See section "Finding memory problems with GNAT Debug Pool" in the *GNAT User's Guide*.

9.26 GNAT.Debug_Utilities ('g-debuti.ads')

Provides a few useful utilities for debugging purposes, including conversion to and from string images of address values.

9.27 GNAT.Directory_Operations (g-dirope.ads)

Provides a set of routines for manipulating directories, including changing the current directory, making new directories, and scanning the files in a directory.

9.28 GNAT.Dynamic_Tables ('g-dyntab.ads')

A generic package providing a single dimension array abstraction where the length of the array can be dynamically modified.

This package provides a facility similar to that of GNAT.Table, except that this package declares a type that can be used to define dynamic instances of the table, while an instantiation of GNAT.Table creates a single instance of the table type.

9.29 GNAT.Exception_Traces ('g-exctra.ads')

Provides an interface allowing to control automatic output upon exception occurrences.

9.30 GNAT.Expect ('g-expect.ads')

Provides a set of subprograms similar to what is available with the standard Tcl Expect tool. It allows you to easily spawn and communicate with an external process. You can send commands or inputs to the process, and compare the output with some expected regular expression. Currently GNAT.Expect is implemented on all native GNAT ports except for OpenVMS. It is not implemented for cross ports, and in particular is not implemented for VxWorks or LynxOS.

9.31 GNAT.Float_Control ('g-flocon.ads')

Provides an interface for resetting the floating-point processor into the mode required for correct semantic operation in Ada. Some third party library calls may cause this mode to be modified, and the Reset procedure in this package can be used to reestablish the required mode.

9.32 GNAT.Heap_Sort_A ('g-hesora.ads')

Provides a general implementation of heap sort usable for sorting arbitrary data items. Move and comparison procedures are provided by passing access-to-procedure values. The algorithm used is a modified heap sort that performs approximately N*log(N) comparisons in the worst case.

9.33 GNAT.Heap_Sort_G ('g-hesorg.ads')

Similar to `Heap_Sort_A` except that the move and sorting procedures are provided as generic parameters, this improves efficiency, especially if the procedures can be inlined, at the expense of duplicating code for multiple instantiations.

9.34 GNAT.HTable ('g-htable.ads')

A generic implementation of hash tables that can be used to hash arbitrary data. Provides two approaches, one a simple static approach, and the other allowing arbitrary dynamic hash tables.

9.35 GNAT.IO ('g-io.ads')

A simple preealborable input-output package that provides a subset of simple Text_IO functions for reading characters and strings from Standard_Input, and writing characters, strings and integers to either Standard_Output or Standard_Error.

9.36 GNAT.IO_Aux ('g-io_aux.ads')

Provides some auxiliary functions for use with Text_IO, including a test for whether a file exists, and functions for reading a line of text.

9.37 GNAT.Lock_Files ('g-locfil.ads')

Provides a general interface for using files as locks. Can be used for providing program level synchronization.

9.38 GNAT.MD5 ('g-md5.ads')

Implements the MD5 Message-Digest Algorithm as described in RFC 1321.

9.39 GNAT.Most_Recent_Exception ('g-moreex.ads')

Provides access to the most recently raised exception. Can be used for various logging purposes, including duplicating functionality of some Ada 83 implementation dependent extensions.

9.40 GNAT.OS_Lib ('g-os_lib.ads')

Provides a range of target independent operating system interface functions, including time/date management, file operations, subprocess management, including a portable spawn procedure, and access to environment variables and error return codes.

9.41 GNAT.Regexp ('g-regexp.ads')

A simple implementation of regular expressions, using a subset of regular expression syntax copied from familiar Unix style utilities. This is the simples of the three pattern matching packages provided, and is particularly suitable for "file globbing" applications.

9.42 GNAT.Registry ('g-regist.ads')

This is a high level binding to the Windows registry. It is possible to do simple things like reading a key value, creating a new key. For full registry API, but at a lower level of abstraction, refer to the Win32.Winreg package provided with the Win32Ada binding

9.43 GNAT.Regpat ('g-regpat.ads')

A complete implementation of Unix-style regular expression matching, copied from the original V7 style regular expression library written in C by Henry Spencer (and binary compatible with this C library).

9.44 GNAT.Sockets ('g-socket.ads')

A high level and portable interface to develop sockets based applications. This package is based on the sockets thin binding found in GNAT.Sockets.Thin. Currently GNAT.Sockets is implemented on all native GNAT ports except for OpenVMS. It is not implemented for the LynxOS cross port.

9.45 GNAT.Source_Info ('g-souinf.ads')

Provides subprograms that give access to source code information known at compile time, such as the current file name and line number.

9.46 GNAT.Spell_Checker ('g-speche.ads')

Provides a function for determining whether one string is a plausible near misspelling of another string.

9.47 GNAT.Spitbol.Patterns ('g-spipat.ads')

A complete implementation of SNOBOL4 style pattern matching. This is the most elaborate of the pattern matching packages provided. It fully duplicates the SNOBOL4 dynamic pattern construction and matching capabilities, using the efficient algorithm developed by Robert Dewar for the SPITBOL system.

9.48 GNAT.Spitbol ('g-spitbo.ads')

The top level package of the collection of SPITBOL-style functionality, this package provides basic SNOBOL4 string manipulation functions, such as Pad, Reverse, Trim, Substr capability, as well as a generic table function useful for constructing arbitrary mappings from strings in the style of the SNOBOL4 TABLE function.

9.49 GNAT.Spitbol.Table_Boolean ('g-sptabo.ads')

A library level of instantiation of `GNAT.Spitbol.Patterns.Table` for type `Standard.Boolean`, giving an implementation of sets of string values.

9.50 GNAT.Spitbol.Table_Integer ('g-sptain.ads')

A library level of instantiation of `GNAT.Spitbol.Patterns.Table` for type `Standard.Integer`, giving an implementation of maps from string to integer values.

9.51 GNAT.Spitbol.Table_VString ('g-sptavs.ads')

A library level of instantiation of GNAT.Spitbol.Patterns.Table for a variable length string type, giving an implementation of general maps from strings to strings.

9.52 GNAT.Table ('g-table.ads')

A generic package providing a single dimension array abstraction where the length of the array can be dynamically modified.

This package provides a facility similar to that of GNAT.Dynamic_Tables, except that this package declares a single instance of the table type, while an instantiation of GNAT.Dynamic_Tables creates a type that can be used to define dynamic instances of the table.

9.53 GNAT.Task_Lock ('g-tasloc.ads')

A very simple facility for locking and unlocking sections of code using a single global task lock. Appropriate for use in situations where contention between tasks is very rarely expected.

9.54 GNAT.Threads ('g-thread.ads')

Provides facilities for creating and destroying threads with explicit calls. These threads are known to the GNAT run-time system. These subprograms are exported C-convention procedures intended to be called from foreign code. By using these primitives rather than directly calling operating systems routines, compatibility with the Ada tasking runt-time is provided.

9.55 GNAT.Traceback ('g-traceb.ads')

Provides a facility for obtaining non-symbolic traceback information, useful in various debugging situations.

9.56 GNAT.Traceback.Symbolic ('g-trasym.ads')

Provides symbolic traceback information that includes the subprogram name and line number information.

9.57 Interfaces.C.Extensions ('i-cexten.ads')

This package contains additional C-related definitions, intended for use with either manually or automatically generated bindings to C libraries.

9.58 Interfaces.C.Streams ('i-cstrea.ads')

This package is a binding for the most commonly used operations on C streams.

9.59 Interfaces.CPP ('i-cpp.ads')

This package provides facilities for use in interfacing to C++. It is primarily intended to be used in connection with automated tools for the generation of C++ interfaces.

9.60 Interfaces.Os2lib ('i-os2lib.ads')

This package provides interface definitions to the OS/2 library. It is a thin binding which is a direct translation of the various '<bse.h>' files.

9.61 Interfaces.Os2lib.Errors ('i-os2err.ads')

This package provides definitions of the OS/2 error codes.

9.62 Interfaces.Os2lib.Synchronization ('i-os2syn.ads')

This is a child package that provides definitions for interfacing to the OS/2 synchronization primitives.

9.63 Interfaces.Os2lib.Threads ('i-os2thr.ads')

This is a child package that provides definitions for interfacing to the OS/2 thread primitives.

9.64 Interfaces.Packed_Decimal ('i-pacdec.ads')

This package provides a set of routines for conversions to and from a packed decimal format compatible with that used on IBM mainframes.

9.65 Interfaces.VxWorks ('i-vxwork.ads')

This package provides a limited binding to the VxWorks API. In particular, it interfaces with the VxWorks hardware interrupt facilities.

9.66 Interfaces.VxWorks.IO ('i-vxwoio.ads')

This package provides a limited binding to the VxWorks' I/O API. In particular, it provides procedures that enable the use of Get_Immediate under VxWorks.

9.67 System.Address_Image ('s-addima.ads')

This function provides a useful debugging function that gives an (implementation dependent) string which identifies an address.

9.68 System.Assertions ('s-assert.ads')

This package provides the declaration of the exception raised by an run-time assertion failure, as well as the routine that is used internally to raise this assertion.

9.69 System.Partition_Interface ('s-parint.ads')

This package provides facilities for partition interfacing. It is used primarily in a distribution context when using Annex E with GLADE.

9.70 System.Task_Info ('s-tasinf.ads')

This package provides target dependent functionality that is used to support the Task_Info pragma

9.71 System.Wch_Cnv ('s-wchcnv.ads')

This package provides routines for converting between wide characters and a representation as a value of type `Standard.String`, using a specified wide character encoding method. It uses definitions in package `System.Wch_Con`.

9.72 System.Wch_Con ('s-wchcon.ads')

This package provides definitions and descriptions of the various methods used for encoding wide characters in ordinary strings. These definitions are used by the package `System.Wch_Cnv`.

10 Interfacing to Other Languages

The facilities in annex B of the Ada 95 Reference Manual are fully implemented in GNAT, and in addition, a full interface to C++ is provided.

10.1 Interfacing to C

Interfacing to C with GNAT can use one of two approaches:

1. The types in the package `Interfaces.C` may be used.
2. Standard Ada types may be used directly. This may be less portable to other compilers, but will work on all GNAT compilers, which guarantee correspondence between the C and Ada types.

Pragma `Convention C` maybe applied to Ada types, but mostly has no effect, since this is the default. The following table shows the correspondence between Ada scalar types and the corresponding C types.

`Integer` `int`

`Short_Integer`
: `short`

`Short_Short_Integer`
: `signed char`

`Long_Integer`
: `long`

`Long_Long_Integer`
: `long long`

`Short_Float`
: `float`

`Float` `float`

`Long_Float`
: `double`

`Long_Long_Float`
: This is the longest floating-point type supported by the hardware.

- Ada enumeration types map to C enumeration types directly if pragma `Convention C` is specified, which causes them to have int length. Without pragma `Convention C`, Ada enumeration types map to 8, 16, or 32 bits (i.e. C types `signed char`, `short`, `int`, respectively) depending on the number of values passed. This is the only case in which pragma `Convention C` affects the representation of an Ada type.
- Ada access types map to C pointers, except for the case of pointers to unconstrained types in Ada, which have no direct C equivalent.
- Ada arrays map directly to C arrays.
- Ada records map directly to C structures.
- Packed Ada records map to C structures where all members are bit fields of the length corresponding to the *type*`'Size` value in Ada.

10.2 Interfacing to C++

The interface to C++ makes use of the following pragmas, which are primarily intended to be constructed automatically using a binding generator tool, although it is possible to construct them by hand. No suitable binding generator tool is supplied with GNAT though.

Using these pragmas it is possible to achieve complete inter-operability between Ada tagged types and C class definitions. See Chapter 1 [Implementation Defined Pragmas], page 3 for more details.

`pragma CPP_Class ([Entity =>] local_name)`
: The argument denotes an entity in the current declarative region that is declared as a tagged or untagged record type. It indicates that the type corresponds to an externally declared C++ class type, and is to be laid out the same way that C++ would lay out the type.

`pragma CPP_Constructor ([Entity =>] local_name)`
: This pragma identifies an imported function (imported in the usual way with pragma `Import`) as corresponding to a C++ constructor.

`pragma CPP_Vtable ...`
: One `CPP_Vtable` pragma can be present for each component of type `CPP.Interfaces.Vtable_Ptr` in a record to which pragma `CPP_Class` applies.

10.3 Interfacing to COBOL

Interfacing to COBOL is achieved as described in section B.4 of the Ada 95 reference manual.

10.4 Interfacing to Fortran

Interfacing to Fortran is achieved as described in section B.5 of the reference manual. The pragma `Convention Fortran`, applied to a multi-dimensional array causes the array to be stored in column-major order as required for convenient interface to Fortran.

10.5 Interfacing to non-GNAT Ada code

It is possible to specify the convention `Ada` in a pragma `Import` or pragma `Export`. However this refers to the calling conventions used by GNAT, which may or may not be similar enough to those used by some other Ada 83 or Ada 95 compiler to allow interoperation.

If arguments types are kept simple, and if the foreign compiler generally follows system calling conventions, then it may be possible to integrate files compiled by other Ada compilers, provided that the elaboration issues are adequately addressed (for example by eliminating the need for any load time elaboration).

In particular, GNAT running on VMS is designed to be highly compatible with the DEC Ada 83 compiler, so this is one case in which it is possible to import foreign units of this type, provided that the data items passed are restricted to simple scalar values or simple record types without variants, or simple array types with fixed bounds.

11 Machine Code Insertions

Package `Machine_Code` provides machine code support as described in the Ada 95 Reference Manual in two separate forms:

- Machine code statements, consisting of qualified expressions that fit the requirements of RM section 13.8.
- An intrinsic callable procedure, providing an alternative mechanism of including machine instructions in a subprogram.

The two features are similar, and both closely related to the mechanism provided by the asm instruction in the GNU C compiler. Full understanding and use of the facilities in this package requires understanding the asm instruction as described in *Using and Porting the GNU Compiler Collection (GCC)* by Richard Stallman. Calls to the function `Asm` and the procedure `Asm` have identical semantic restrictions and effects as described below. Both are provided so that the procedure call can be used as a statement, and the function call can be used to form a code_statement.

The first example given in the GCC documentation is the C `asm` instruction:

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

The equivalent can be written for GNAT as:

```
Asm ("fsinx %1 %0",
     My_Float'Asm_Output ("=f", result),
     My_Float'Asm_Input  ("f",  angle));
```

The first argument to `Asm` is the assembler template, and is identical to what is used in GNU C. This string must be a static expression. The second argument is the output operand list. It is either a single `Asm_Output` attribute reference, or a list of such references enclosed in parentheses (technically an array aggregate of such references).

The `Asm_Output` attribute denotes a function that takes two parameters. The first is a string, the second is the name of a variable of the type designated by the attribute prefix. The first (string) argument is required to be a static expression and designates the constraint for the parameter (e.g. what kind of register is required). The second argument is the variable to be updated with the result. The possible values for constraint are the same as those used in the RTL, and are dependent on the configuration file used to build the GCC back end. If there are no output operands, then this argument may either be omitted, or explicitly given as `No_Output_Operands`.

The second argument of *`my_float'`*`Asm_Output` functions as though it were an `out` parameter, which is a little curious, but all names have the form of expressions, so there is no syntactic irregularity, even though normally functions would not be permitted `out` parameters. The third argument is the list of input operands. It is either a single `Asm_Input` attribute reference, or a list of such references enclosed in parentheses (technically an array aggregate of such references).

The `Asm_Input` attribute denotes a function that takes two parameters. The first is a string, the second is an expression of the type designated by the prefix. The first (string) argument is required to be a static expression, and is the constraint for the parameter, (e.g. what kind of register is required). The second argument is the value to be used as the input argument. The possible values for the constant are the same as those used in the RTL, and are dependent on the configuration file used to built the GCC back end.

If there are no input operands, this argument may either be omitted, or explicitly given as `No_Input_Operands`. The fourth argument, not present in the above example, is a list of register names, called the *clobber* argument. This argument, if given, must be a static string expression, and is a space or comma separated list of names of registers that must be considered destroyed as a result of the `Asm` call. If this argument is the null string (the default value), then the code generator assumes that no additional registers are destroyed.

The fifth argument, not present in the above example, called the *volatile* argument, is by default `False`. It can be set to the literal value `True` to indicate to the code generator that all optimizations with respect to the instruction specified should be suppressed, and that in particular, for an instruction that has outputs, the instruction will still be generated, even if none of the outputs are used. See the full description in the GCC manual for further details.

The `Asm` subprograms may be used in two ways. First the procedure forms can be used anywhere a procedure call would be valid, and correspond to what the RM calls "intrinsic" routines. Such calls can be used to intersperse machine instructions with other Ada statements. Second, the function forms, which return a dummy value of the limited private type `Asm_Insn`, can be used in code statements, and indeed this is the only context where such calls are allowed. Code statements appear as aggregates of the form:

```
Asm_Insn'(Asm (...));
Asm_Insn'(Asm_Volatile (...));
```

In accordance with RM rules, such code statements are allowed only within subprograms whose entire body consists of such statements. It is not permissible to intermix such statements with other Ada statements.

Typically the form using intrinsic procedure calls is more convenient and more flexible. The code statement form is provided to meet the RM suggestion that such a facility should be made available. The following is the exact syntax of the call to `Asm` (of course if named notation is used, the arguments may be given in arbitrary order, following the normal rules for use of positional and named arguments)

```
ASM_CALL ::= Asm (
                [Template =>] static_string_EXPRESSION
               [,[Outputs  =>] OUTPUT_OPERAND_LIST      ]
               [,[Inputs   =>] INPUT_OPERAND_LIST       ]
               [,[Clobber  =>] static_string_EXPRESSION ]
               [,[Volatile =>] static_boolean_EXPRESSION] )
OUTPUT_OPERAND_LIST ::=
  No_Output_Operands
| OUTPUT_OPERAND_ATTRIBUTE
| (OUTPUT_OPERAND_ATTRIBUTE {,OUTPUT_OPERAND_ATTRIBUTE})
OUTPUT_OPERAND_ATTRIBUTE ::=
  SUBTYPE_MARK'Asm_Output (static_string_EXPRESSION, NAME)
INPUT_OPERAND_LIST ::=
  No_Input_Operands
| INPUT_OPERAND_ATTRIBUTE
| (INPUT_OPERAND_ATTRIBUTE {,INPUT_OPERAND_ATTRIBUTE})
INPUT_OPERAND_ATTRIBUTE ::=
  SUBTYPE_MARK'Asm_Input (static_string_EXPRESSION, EXPRESSION)
```

12 GNAT Implementation of Tasking

12.1 Mapping Ada Tasks onto the Underlying Kernel Threads

GNAT run-time system comprises two layers:

- GNARL (GNAT Run-time Layer)
- GNULL (GNAT Low-level Library)

In GNAT, Ada's tasking services rely on a platform and OS independent layer known as GNARL. This code is responsible for implementing the correct semantics of Ada's task creation, rendezvous, protected operations etc.

GNARL decomposes Ada's tasking semantics into simpler lower level operations such as create a thread, set the priority of a thread, yield, create a lock, lock/unlock, etc. The spec for these low-level operations constitutes GNULLI, the GNULL Interface. This interface is directly inspired from the POSIX real-time API.

If the underlying executive or OS implements the POSIX standard faithfully, the GNULL Interface maps as is to the services offered by the underlying kernel. Otherwise, some target dependent glue code maps the services offered by the underlying kernel to the semantics expected by GNARL.

Whatever the underlying OS (VxWorks, UNIX, OS/2, Windows NT, etc.) the key point is that each Ada task is mapped on a thread in the underlying kernel. For example, in the case of VxWorks, one Ada task = one VxWorks task.

In addition Ada task priorities map onto the underlying thread priorities. Mapping Ada tasks onto the underlying kernel threads has several advantages:

1. The underlying scheduler is used to schedule the Ada tasks. This makes Ada tasks as efficient as kernel threads from a scheduling standpoint.
2. Interaction with code written in C containing threads is eased since at the lowest level Ada tasks and C threads map onto the same underlying kernel concept.
3. When an Ada task is blocked during I/O the remaining Ada tasks are able to proceed.
4. On multi-processor systems Ada Tasks can execute in parallel.

12.2 Ensuring Compliance with the Real-Time Annex

The reader will be quick to notice that while mapping Ada tasks onto the underlying threads has significant advantages, it does create some complications when it comes to respecting the scheduling semantics specified in the real-time annex (Annex D).

For instance Annex D requires that for the FIFO_Within_Priorities scheduling policy we have:

```
When the active priority of a ready task that is not running
changes, or the setting of its base priority takes effect, the
task is removed from the ready queue for its old active priority
and is added at the tail of the ready queue for its new active
priority, except in the case where the active priority is lowered
```

```
due to the loss of inherited priority, in which case the task is
added at the head of the ready queue for its new active priority.
```

While most kernels do put tasks at the end of the priority queue when a task changes its priority, (which respects the main FIFO_Within_Priorities requirement), almost none keep a thread at the beginning of its priority queue when its priority drops from the loss of inherited priority.

As a result most vendors have provided incomplete Annex D implementations.

The GNAT run-time, has a nice cooperative solution to this problem which ensures that accurate FIFO_Within_Priorities semantics are respected.

The principle is as follows. When an Ada task T is about to start running, it checks whether some other Ada task R with the same priority as T has been suspended due to the loss of priority inheritance. If this is the case, T yields and is placed at the end of its priority queue. When R arrives at the front of the queue it executes.

Note that this simple scheme preserves the relative order of the tasks that were ready to execute in the priority queue where R has been placed at the end.

13 Code generation for array aggregates

Aggregate have a rich syntax and allow the user to specify the values of complex data structures by means of a single construct. As a result, the code generated for aggregates can be quite complex and involve loops, case statements and multiple assignments. In the simplest cases, however, the compiler will recognize aggregates whose components and constraints are fully static, and in those cases the compiler will generate little or no executable code. The following is an outline of the code that GNAT generates for various aggregate constructs. For further details, the user will find it useful to examine the output produced by the -gnatG flag to see the expanded source that is input to the code generator. The user will also want to examine the assembly code generated at various levels of optimization.

The code generated for aggregates depends on the context, the component values, and the type. In the context of an object declaration the code generated is generally simpler than in the case of an assignment. As a general rule, static component values and static subtypes also lead to simpler code.

13.1 Static constant aggregates with static bounds

For the declarations:

```
type One_Dim is array (1..10) of integer;
ar0 : constant One_Dim := ( 1, 2, 3, 4, 5, 6, 7, 8, 9, 0);
```

GNAT generates no executable code: the constant ar0 is placed in static memory. The same is true for constant aggregates with named associations:

```
Cr1 : constant One_Dim := (4 => 16, 2 => 4, 3 => 9, 1=> 1);
Cr3 : constant One_Dim := (others => 7777);
```

The same is true for multidimensional constant arrays such as:

```
type two_dim is array (1..3, 1..3) of integer;
Unit : constant two_dim := ( (1,0,0), (0,1,0), (0,0,1));
```

The same is true for arrays of one-dimensional arrays: the following are static:

```
type ar1b is array (1..3) of boolean;
type ar_ar is array (1..3) of ar1b;
None : constant ar1b := (others => false);      --  fully static
None2 : constant ar_ar := (1..3 => None);       --  fully static
```

However, for multidimensional aggregates with named associations, GNAT will generate assignments and loops, even if all associations are static. The following two declarations generate a loop for the first dimension, and individual component assignments for the second dimension:

```
Zero1: constant two_dim := (1..3 => (1..3 => 0));
Zero2: constant two_dim := (others => (others => 0));
```

13.2 Constant aggregates with an unconstrained nominal types

In such cases the aggregate itself establishes the subtype, so that associations with `others` cannot be used. GNAT determines the bounds for the actual subtype of the aggregate, and allocates the aggregate statically as well. No code is generated for the following:

```
type One_Unc is array (natural range <>) of integer;
Cr_Unc : constant One_Unc := (12,24,36);
```

13.3 Aggregates with static bounds

In all previous examples the aggregate was the initial (and immutable) value of a constant. If the aggregate initializes a variable, then code is generated for it as a combination of individual assignments and loops over the target object. The declarations

```
Cr_Var1 : One_Dim := (2, 5, 7, 11);
Cr_Var2 : One_Dim := (others > -1);
```

generate the equivalent of

```
Cr_Var1 (1) := 2;
Cr_Var1 (2) := 3;
Cr_Var1 (3) := 5;
Cr_Var1 (4) := 11;

for I in Cr_Var2'range loop
   Cr_Var2 (I) := =-1;
end loop;
```

13.4 Aggregates with non-static bounds

If the bounds of the aggregate are not statically compatible with the bounds of the nominal subtype of the target, then constraint checks have to be generated on the bounds. For a multidimensional array, constraint checks may have to be applied to sub-arrays individually, if they do not have statically compatible subtypes.

13.5 Aggregates in assignments statements

In general, aggregate assignment requires the construction of a temporary, and a copy from the temporary to the target of the assignment. This is because it is not always possible to convert the assignment into a series of individual component assignments. For example, consider the simple case:

A := (A(2), A(1));

This cannot be converted into:

```
A(1) := A(2);
A(2) := A(1);
```

So the aggregate has to be built first in a separate location, and then copied into the target. GNAT recognizes simple cases where this intermediate step is not required, and the assignments can be performed in place, directly into the target. The following sufficient criteria are applied:

1. The bounds of the aggregate are static, and the associations are static.
2. The components of the aggregate are static constants, names of simple variables that are not renamings, or expressions not involving indexed components whose operands obey these rules.

If any of these conditions are violated, the aggregate will be built in a temporary (created either by the front-end or the code generator) and then that temporary will be copied onto the target.

14 Specialized Needs Annexes

Ada 95 defines a number of specialized needs annexes, which are not required in all implementations. However, as described in this chapter, GNAT implements all of these special needs annexes:

Systems Programming (Annex C)
: The Systems Programming Annex is fully implemented.

Real-Time Systems (Annex D)
: The Real-Time Systems Annex is fully implemented.

Distributed Systems (Annex E)
: Stub generation is fully implemented in the GNAT compiler. In addition, a complete compatible PCS is available as part of the GLADE system, a separate product. When the two products are used in conjunction, this annex is fully implemented.

Information Systems (Annex F)
: The Information Systems annex is fully implemented.

Numerics (Annex G)
: The Numerics Annex is fully implemented.

Safety and Security (Annex H)
: The Safety and Security annex is fully implemented.

15 Compatibility Guide

This chapter contains sections that describe compatibility issues between GNAT and other Ada 83 and Ada 95 compilation systems, to aid in porting applications developed in other Ada environments.

15.1 Compatibility with Ada 83

Ada 95 is designed to be highly upwards compatible with Ada 83. In particular, the design intention is that the difficulties associated with moving from Ada 83 to Ada 95 should be no greater than those that occur when moving from one Ada 83 system to another.

However, there are a number of points at which there are minor incompatibilities. The Ada 95 Annotated Reference Manual contains full details of these issues, and should be consulted for a complete treatment. In practice the following are the most likely issues to be encountered.

Character range
: The range of `Standard.Character` is now the full 256 characters of Latin-1, whereas in most Ada 83 implementations it was restricted to 128 characters. This may show up as compile time or runtime errors. The desirable fix is to adapt the program to accommodate the full character set, but in some cases it may be convenient to define a subtype or derived type of Character that covers only the restricted range.

New reserved words
: The identifiers `abstract`, `aliased`, `protected`, `requeue`, `tagged`, and `until` are reserved in Ada 95. Existing Ada 83 code using any of these identifiers must be edited to use some alternative name.

Freezing rules
: The rules in Ada 95 are slightly different with regard to the point at which entities are frozen, and representation pragmas and clauses are not permitted past the freeze point. This shows up most typically in the form of an error message complaining that a representation item appears too late, and the appropriate corrective action is to move the item nearer to the declaration of the entity to which it refers.

  A particular case is that representation pragmas (including the extended DEC Ada 83 compatibility pragmas such as `Export_Procedure`), cannot be applied to a subprogram body. If necessary, a separate subprogram declaration must be introduced to which the pragma can be applied.

Optional bodies for library packages
: In Ada 83, a package that did not require a package body was nevertheless allowed to have one. This lead to certain surprises in compiling large systems (situations in which the body could be unexpectedly ignored). In Ada 95, if a package does not require a body then it is not permitted to have a body. To fix this problem, simply remove a redundant body if it is empty, or, if it is non-empty, introduce a dummy declaration into the spec that makes the body

required. One approach is to add a private part to the package declaration (if necessary), and define a parameterless procedure called Requires_Body, which must then be given a dummy procedure body in the package body, which then becomes required.

`Numeric_Error` is now the same as `Constraint_Error`
: In Ada 95, the exception `Numeric_Error` is a renaming of `Constraint_Error`. This means that it is illegal to have separate exception handlers for the two exceptions. The fix is simply to remove the handler for the `Numeric_Error` case (since even in Ada 83, a compiler was free to raise `Constraint_Error` in place of `Numeric_Error` in all cases).

Indefinite subtypes in generics
: In Ada 83, it was permissible to pass an indefinite type (e.g. `String`) as the actual for a generic formal private type, but then the instantiation would be illegal if there were any instances of declarations of variables of this type in the generic body. In Ada 95, to avoid this clear violation of the contract model, the generic declaration clearly indicates whether or not such instantiations are permitted. If a generic formal parameter has explicit unknown discriminants, indicated by using `(<>)` after the type name, then it can be instantiated with indefinite types, but no variables can be declared of this type. Any attempt to declare a variable will result in an illegality at the time the generic is declared. If the `(<>)` notation is not used, then it is illegal to instantiate the generic with an indefinite type. This will show up as a compile time error, and the fix is usually simply to add the `(<>)` to the generic declaration.

All implementations of GNAT provide a switch that causes GNAT to operate in Ada 83 mode. In this mode, some but not all compatibility problems of the type described above are handled automatically. For example, the new Ada 95 protected keywords are not recognized in this mode. However, in practice, it is usually advisable to make the necessary modifications to the program to remove the need for using this switch.

15.2 Compatibility with Other Ada 95 Systems

Providing that programs avoid the use of implementation dependent and implementation defined features of Ada 95, as documented in the Ada 95 reference manual, there should be a high degree of portability between GNAT and other Ada 95 systems. The following are specific items which have proved troublesome in moving GNAT programs to other Ada 95 compilers, but do not affect porting code to GNAT.

Ada 83 Pragmas and Attributes
: Ada 95 compilers are allowed, but not required, to implement the missing Ada 83 pragmas and attributes that are no longer defined in Ada 95. GNAT implements all such pragmas and attributes, eliminating this as a compatibility concern, but some other Ada 95 compilers reject these pragmas and attributes.

Special-needs Annexes
: GNAT implements the full set of special needs annexes. At the current time, it is the only Ada 95 compiler to do so. This means that programs making use of these features may not be portable to other Ada 95 compilation systems.

Representation Clauses
: Some other Ada 95 compilers implement only the minimal set of representation clauses required by the Ada 95 reference manual. GNAT goes far beyond this minimal set, as described in the next section.

15.3 Representation Clauses

The Ada 83 reference manual was quite vague in describing both the minimal required implementation of representation clauses, and also their precise effects. The Ada 95 reference manual is much more explicit, but the minimal set of capabilities required in Ada 95 is quite limited.

GNAT implements the full required set of capabilities described in the Ada 95 reference manual, but also goes much beyond this, and in particular an effort has been made to be compatible with existing Ada 83 usage to the greatest extent possible.

A few cases exist in which Ada 83 compiler behavior is incompatible with requirements in the Ada 95 reference manual. These are instances of intentional or accidental dependence on specific implementation dependent characteristics of these Ada 83 compilers. The following is a list of the cases most likely to arise in existing legacy Ada 83 code.

Implicit Packing
: Some Ada 83 compilers allowed a Size specification to cause implicit packing of an array or record. This could cause expensive implicit conversions for change of representation in the presence of derived types, and the Ada design intends to avoid this possibility. Subsequent AI's were issued to make it clear that such implicit change of representation in response to a Size clause is inadvisable, and this recommendation is represented explicitly in the Ada 95 RM as implementation advice that is followed by GNAT. The problem will show up as an error message rejecting the size clause. The fix is simply to provide the explicit pragma `Pack`, or for more fine tuned control, provide a Component_Size clause.

Meaning of Size Attribute
: The Size attribute in Ada 95 for discrete types is defined as being the minimal number of bits required to hold values of the type. For example, on a 32-bit machine, the size of Natural will typically be 31 and not 32 (since no sign bit is required). Some Ada 83 compilers gave 31, and some 32 in this situation. This problem will usually show up as a compile time error, but not always. It is a good idea to check all uses of the 'Size attribute when porting Ada 83 code. The GNAT specific attribute Object_Size can provide a useful way of duplicating the behavior of some Ada 83 compiler systems.

Size of Access Types
: A common assumption in Ada 83 code is that an access type is in fact a pointer, and that therefore it will be the same size as a System.Address value. This assumption is true for GNAT in most cases with one exception. For the case of a pointer to an unconstrained array type (where the bounds may vary from one value of the access type to another), the default is to use a "fat pointer", which is represented as two separate pointers, one to the bounds, and one to the array. This representation has a number of advantages, including improved

efficiency. However, it may cause some difficulties in porting existing Ada 83 code which makes the assumption that, for example, pointers fit in 32 bits on a machine with 32-bit addressing.

To get around this problem, GNAT also permits the use of "thin pointers" for access types in this case (where the designated type is an unconstrained array type). These thin pointers are indeed the same size as a System.Address value. To specify a thin pointer, use a size clause for the type, for example:

```
type X is access all String;
for X'Size use Standard'Address_Size;
```

which will cause the type X to be represented using a single pointer. When using this representation, the bounds are right behind the array. This representation is slightly less efficient, and does not allow quite such flexibility in the use of foreign pointers or in using the Unrestricted_Access attribute to create pointers to non-aliased objects. But for any standard portable use of the access type it will work in a functionally correct manner and allow porting of existing code. Note that another way of forcing a thin pointer representation is to use a component size clause for the element size in an array, or a record representation clause for an access field in a record.

15.4 Compatibility with DEC Ada 83

The VMS version of GNAT fully implements all the pragmas and attributes provided by DEC Ada 83, as well as providing the standard DEC Ada 83 libraries, including Starlet. In addition, data layouts and parameter passing conventions are highly compatible. This means that porting existing DEC Ada 83 code to GNAT in VMS systems should be easier than most other porting efforts. The following are some of the most significant differences between GNAT and DEC Ada 83.

Default floating-point representation

In GNAT, the default floating-point format is IEEE, whereas in DEC Ada 83, it is VMS format. GNAT does implement the necessary pragmas (Long_Float, Float_Representation) for changing this default.

System

The package System in GNAT exactly corresponds to the definition in the Ada 95 reference manual, which means that it excludes many of the DEC Ada 83 extensions. However, a separate package Aux_DEC is provided that contains the additional definitions, and a special pragma, Extend_System allows this package to be treated transparently as an extension of package System.

To_Address

The definitions provided by Aux_DEC are exactly compatible with those in the DEC Ada 83 version of System, with one exception. DEC Ada provides the following declarations:

```
TO_ADDRESS(INTEGER)
TO_ADDRESS(UNSIGNED_LONGWORD)
TO_ADDRESS(universal_integer)
```

The version of TO_ADDRESS taking a universal integer argument is in fact an extension to Ada 83 not strictly compatible with the reference manual. In

GNAT, we are constrained to be exactly compatible with the standard, and this means we cannot provide this capability. In DEC Ada 83, the point of this definition is to deal with a call like:

```
TO_ADDRESS (16#12777#);
```

Normally, according to the Ada 83 standard, one would expect this to be ambiguous, since it matches both the INTEGER and UNSIGNED_LONGWORD forms of TO_ADDRESS. However, in DEC Ada 83, there is no ambiguity, since the definition using universal_integer takes precedence.

In GNAT, since the version with universal_integer cannot be supplied, it is not possible to be 100% compatible. Since there are many programs using numeric constants for the argument to TO_ADDRESS, the decision in GNAT was to change the name of the function in the UNSIGNED_LONGWORD case, so the declarations provided in the GNAT version of AUX_Dec are:

```
function To_Address (X : Integer) return Address;
pragma Pure_Function (To_Address);

function To_Address_Long (X : Unsigned_Longword)
 return Address;
pragma Pure_Function (To_Address_Long);
```

This means that programs using TO_ADDRESS for UNSIGNED_LONGWORD must change the name to TO_ADDRESS_LONG.

Task_Id values

The Task_Id values assigned will be different in the two systems, and GNAT does not provide a specified value for the Task_Id of the environment task, which in GNAT is treated like any other declared task.

For full details on these and other less significant compatibility issues, see appendix E of the Digital publication entitled *DEC Ada, Technical Overview and Comparison on DIGITAL Platforms*.

For GNAT running on other than VMS systems, all the DEC Ada 83 pragmas and attributes are recognized, although only a subset of them can sensibly be implemented. The description of pragmas in this reference manual indicates whether or not they are applicable to non-VMS systems.

GNU Free Documentation License

Version 1.2, November 2002

Copyright © 2000,2001,2002 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307, USA

Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document *free* in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released

under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and

that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

C. State on the Title page the name of the publisher of the Modified Version, as the publisher.

D. Preserve all the copyright notices of the Document.

E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.

F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.

G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.

H. Include an unaltered copy of this License.

I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.

L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.

M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.

N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.

O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

   You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

   The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

   In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements."

6. COLLECTIONS OF DOCUMENTS

   You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

   You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

   A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called

an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

   Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warrany Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

   If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

   You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

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

    Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

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

```
Copyright (C)  year  your name.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.
A copy of the license is included in the section entitled ``GNU
Free Documentation License''.
```

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

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

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

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

Index

-

-gnatR switch 121

A

Abort_Defer 3
Abort_Signal 41
Access, unrestricted 48
Accuracy requirements 72
Accuracy, complex arithmetic 72
Ada 83 attributes 43, 44, 45, 47
Ada 95 ISO/ANSI Standard 2
Ada.Characters.Handling 62
Ada.Characters.Latin_9 ('a-chlat9.ads') ... 147
Ada.Characters.Wide_Latin_1 ('a-cwila1.ads') 147
Ada.Characters.Wide_Latin_9 ('a-cwila1.ads') 147
Ada.Command_Line.Remove ('a-colire.ads') .. 148
Ada.Direct_IO.C_Streams ('a-diocst.ads') .. 148
Ada.Exceptions.Is_Null_Occurrence ('a-einuoc.ads') 148
Ada.Sequential_IO.C_Streams ('a-siocst.ads') 148
Ada.Streams.Stream_IO.C_Streams ('a-ssicst.ads') 148
Ada.Strings.Unbounded.Text_IO ('a-suteio.ads') 148
Ada.Strings.Wide_Unbounded.Wide_Text_IO ('a-swuwti.ads') 148
Ada.Text_IO.C_Streams ('a-tiocst.ads') 149
Ada.Wide_Text_IO.C_Streams ('a-wtcstr.ads') 149
Ada_83 3
Ada_95 3
Address Clause 118
Address clauses 56
Address image 156
Address of subprogram code 42
Address, as private type 60
Address, operations of 60
Address_Size 41
Alignment Clause 103
Alignment clauses 57
Alignment, default 104
Alignment, maximum 45
Alignments of components 6
Alternative Character Sets 53
Annotate 4
Argument passing mechanisms 12
Arrays, extendable 151, 154
Arrays, multidimensional 54
Asm_Input 41
Asm_Output 41
Assert 4
Assert_Failure, exception 156
Assertions 156
Ast_Entry 5
AST_Entry 42
Attribute 119
AWK 149

B

Biased representation 107
Big endian 43
Bit 42
bit ordering 110
Bit ordering 60
Bit_Order Clause 110
Bit_Position 42
Boolean_Entry_Barriers 83
Bounded errors 51
Bounded-length strings 62
Bubble sort 149
byte ordering 111

C

C streams, interfacing 155
C Streams, Interfacing with Direct_IO 148
C Streams, Interfacing with Sequential_IO 148
C Streams, Interfacing with Stream_IO 148
C Streams, Interfacing with Text_IO 149
C Streams, Interfacing with Wide_Text_IO 149
C++ interfacing 155
C, interfacing with 64
C_Pass_By_Copy 5
Calendar 149
Casing of External names 15
Casing utilities 150
CGI (Common Gateway Interface) 150
CGI (Common Gateway Interface) cookie support 150
CGI (Common Gateway Interface) debugging 150
Character handling (GNAT.Case_Util) 150
Character Sets 53
Checks, suppression of 55
Child Units 51
COBOL support 71
COBOL, interfacing with 65
Code_Address 42
Command line 150
Command line, argument removal 148
Comment 5
Common_Object 6
Compatibility (between Ada 83 and Ada 95) .. 169
Complex arithmetic accuracy 72

Complex elementary functions 72
Complex types 71
Complex_Representation 6
Component Clause 117
Component_Alignment 6
Component_Size 7
Component_Size Clause 110
Component_Size clauses 59
Component_Size_4 7
Convention, effect on representation 120
Convention_Identifier 7
Conventions, synonyms 7
Conventions, typographical 2
Cookie support in CGI 150
CPP_Class 8
CPP_Constructor 8
CPP_Virtual 9
CPP_Vtable 9
CRC32 150
Current exception 150
Cyclic Redundancy Check 150

D

Debug 10
Debug pools 151
Debugging 151
debugging with Initialize_Scalars 20
Dec Ada 83 15
Dec Ada 83 casing compatibility 15
Decimal radix support 71
Default_Bit_Order 43
Deferring aborts 3
Directory operations 151
Discriminants, testing for 44
Duration'Small 54

E

Elab_Body 43
Elab_Spec 43
Elaborated 43
Elaboration control 10
Elaboration_Checks 10
Eliminate 10
Elimination of unused subprograms 10
Emax 43
Enclosing_Entity 99
Entry queuing policies 69
Enum_Rep 43
Enumeration representation clauses 59
Enumeration values 53
Epsilon 44
Error detection 51
Exception information 55
Exception retrieval 150
Exception traces 151
Exception, obtaining most recent 152
Exception_Information' 99
Exception_Message 100
Exception_Name 100
Export 63, 119
Export_Exception 12
Export_Function 12
Export_Object 13
Export_Procedure 13
Export_Valued_Procedure 14
Extend_System 15
External 15
External Names, casing 15
External_Name_Casing 15

F

FDL, GNU Free Documentation License 175
File 100
File locking 152
Finalize_Storage_Only 17
Fixed_Value 44
Float types 53
Float_Representation 17
Floating-Point Processor 151
Foreign threads 155
Fortran, interfacing with 66

G

Get_Immediate 63
Get_Immediate 137
GNAT.AWK ('g-awk.ads') 149
GNAT.Bubble_Sort_A ('g-busora.ads') 149
GNAT.Bubble_Sort_G ('g-busorg.ads') 149
GNAT.Calendar ('g-calend.ads') 149
GNAT.Calendar.Time_IO ('g-catiio.ads') 149
GNAT.Case_Util ('g-casuti.ads') 150
GNAT.CGI ('g-cgi.ads') 150
GNAT.CGI.Cookie ('g-cgicoo.ads') 150
GNAT.CGI.Debug ('g-cgideb.ads') 150
GNAT.Command_Line ('g-comlin.ads') 150
GNAT.CRC32 ('g-crc32.ads') 150
GNAT.Current_Exception ('g-curexc.ads') 150
GNAT.Debug_Pools ('g-debpoo.ads') 151
GNAT.Debug_Utilities ('g-debuti.ads') 151
GNAT.Directory_Operations (g-dirope.ads) 151
GNAT.Dynamic_Tables ('g-dyntab.ads') 151
GNAT.Exception_Traces ('g-exctra.ads') 151
GNAT.Expect ('g-expect.ads') 151
GNAT.Float_Control ('g-flocon.ads') 151
GNAT.Heap_Sort_A ('g-hesora.ads') 152
GNAT.Heap_Sort_G ('g-hesorg.ads') 152
GNAT.HTable ('g-htable.ads') 152
GNAT.IO ('g-io.ads') 152
GNAT.IO_Aux ('g-io_aux.ads') 152
GNAT.Lock_Files ('g-locfil.ads') 152
GNAT.MD5 ('g-md5.ads') 152

GNAT.Most_Recent_Exception ('g-moreex.ads') ... 152
GNAT.OS_Lib ('g-os_lib.ads') ... 153
GNAT.Regexp ('g-regexp.ads') ... 153
GNAT.Registry ('g-regist.ads') ... 153
GNAT.Regpat ('g-regpat.ads') ... 153
GNAT.Sockets ('g-socket.ads') ... 153
GNAT.Source_Info ('g-souinf.ads') ... 153
GNAT.Spell_Checker ('g-speche.ads') ... 153
GNAT.Spitbol ('g-spitbo.ads') ... 154
GNAT.Spitbol.Patterns ('g-spipat.ads') ... 154
GNAT.Spitbol.Table_Boolean ('g-sptabo.ads') ... 154
GNAT.Spitbol.Table_Integer ('g-sptain.ads') ... 154
GNAT.Spitbol.Table_VString ('g-sptavs.ads') ... 154
GNAT.Table ('g-table.ads') ... 154
GNAT.Task_Lock ('g-tasloc.ads') ... 154
GNAT.Threads ('g-thread.ads') ... 155
GNAT.Traceback ('g-traceb.ads') ... 155
GNAT.Traceback.Symbolic ('g-trasym.ads') ... 155

H

Has_Discriminants ... 44
Hash tables ... 152
Heap usage, implicit ... 61

I

IBM Packed Format ... 156
Ident ... 17
Image, of an address ... 156
Img ... 44
Implementation-dependent features ... 1
Import ... 119
Import_Exception ... 17
Import_Function ... 18
Import_Object ... 19
Import_Procedure ... 19
Import_Valued_Procedure ... 20
Initialization, suppression of ... 34
Initialize_Scalars ... 20
Inline_Always ... 21
Inline_Generic ... 21
Input/Output facilities ... 152
Integer maps ... 154
Integer types ... 53
Integer_Value ... 44
Interface ... 21
Interface_Name ... 22
Interfaces ... 64
Interfaces.C.Extensions ('i-cexten.ads') ... 155
Interfaces.C.Streams ('i-cstrea.ads') ... 155
Interfaces.CPP ('i-cpp.ads') ... 155
Interfaces.Os2lib ('i-os2lib.ads') ... 155
Interfaces.Os2lib.Errors ('i-os2err.ads') ... 155
Interfaces.Os2lib.Synchronization ('i-os2syn.ads') ... 156
Interfaces.Os2lib.Threads ('i-os2thr.ads') ... 156
Interfaces.Packed_Decimal ('i-pacdec.ads') ... 156
Interfaces.VxWorks ('i-vxwork.ads') ... 156
Interfaces.VxWorks.IO ('i-vxwoio.ads') ... 156
Interfacing to C++ ... 9
Interfacing to VxWorks ... 156
Interfacing to VxWorks' I/O ... 156
Interfacing with C++ ... 8, 9
Interfacing, to C++ ... 155
Interfacing, to OS/2 ... 155, 156
Interrupt priority, maximum ... 45
Interrupt support ... 68
Interrupts ... 68
Intrinsic operator ... 99
Intrinsic Subprograms ... 99

L

Large ... 45
Latin-1 ... 169
Latin_1 constants for Wide_Character ... 147
Latin_9 constants for Character ... 147
Latin_9 constants for Wide_Character ... 147
License ... 22
License checking ... 22
Line ... 100
Link_With ... 23
Linker_Alias ... 23
Linker_Section ... 23
Little endian ... 43
Locking ... 154
Locking Policies ... 69
Locking using files ... 152
Long_Float ... 25

M

Machine operations ... 66
Machine_Attribute ... 25
Machine_Size ... 45
Main_Storage ... 25
Mantissa ... 45
Maps ... 154
Max_Entry_Queue_Depth ... 83
Max_Interrupt_Priority ... 45
Max_Priority ... 45
Maximum_Alignment ... 45
Mechanism_Code ... 45
Memory corruption debugging ... 151
Message Digest MD5 ... 152
Multidimensional arrays ... 54

N

Named numbers, representation of ... 48
No_Calendar ... 83
No_Dynamic_Interrupts ... 83
No_Elaboration_Code ... 85
No_Entry_Calls_In_Elaboration_Code ... 83
No_Entry_Queue ... 85
No_Enumeration_Maps ... 83
No_Exception_Handlers ... 84
No_Implementation_Attributes ... 85
No_Implementation_Pragmas ... 86
No_Implementation_Restrictions ... 86
No_Implicit_Conditionals ... 84
No_Implicit_Loops ... 84
No_Local_Protected_Objects ... 84
No_Protected_Type_Allocators ... 84
No_Return ... 25
No_Run_Time ... 24
No_Secondary_Stack ... 84
No_Select_Statements ... 84
No_Standard_Storage_Pools ... 84
No_Streams ... 84
No_Task_Attributes ... 84
No_Task_Termination ... 85
No_Tasking ... 85
No_Wide_Characters ... 85
Normalize_Scalars ... 24
Null_Occurrence, testing for ... 148
Null_Parameter ... 46
Numerics ... 71

O

Object_Size ... 46, 107
OpenVMS ... 5, 7, 12, 13, 17, 18, 25, 35, 42, 46
Operating System interface ... 153
Operations, on Address ... 60
ordering, of bits ... 110
ordering, of bytes ... 111
OS/2 Error codes ... 155
OS/2 interfacing ... 155
OS/2 synchronization primitives ... 156
OS/2 thread interfacing ... 156

P

Package Interfaces ... 64
Package Interrupts ... 68
Package Task_Attributes ... 69
Packed Decimal ... 156
Packed types ... 56
Parameters, passing mechanism ... 45
Parameters, when passed by reference ... 46
Parsing ... 149
Partition communication subsystem ... 70
Partition intefacing functions ... 156
Passed_By_Reference ... 46
Passing by copy ... 5
Passing by descriptor ... 13, 18
Passive ... 26
Pattern matching ... 153, 154
PCS ... 70
Polling ... 26
Portability ... 1
Pragma Pack (for arrays) ... 115
Pragma Pack (for records) ... 116
Pragma, representation ... 103
Pragmas ... 51
Pre-elaboration requirements ... 68
Preemptive abort ... 69
Priority, maximum ... 45
Propagate_Exceptions ... 26
Protected procedure handlers ... 68
Psect_Object ... 27
Pure ... 28
Pure_Function ... 27

R

Random number generation ... 63
Range_Length ... 46
Ravenscar ... 28
Record Representation Clause ... 117
Record representation clauses ... 59
Regular expressions ... 153
Removing command line arguments ... 148
Representation Clause ... 103
Representation clauses ... 55
Representation Clauses ... 103
Representation clauses, enumeration ... 59
Representation clauses, records ... 59
Representation of enums ... 43
Representation of wide characters ... 157
Representation Pragma ... 103
Representation, determination of ... 121
Restricted_Run_Time ... 30
Return values, passing mechanism ... 45
Rotate_Left ... 100
Rotate_Right ... 101

S

Safe_Emax ... 47
Safe_Large ... 47
Sets of strings ... 154
Share_Generic ... 30
Shift_Left ... 101
Shift_Right ... 101
Shift_Right_Arithmetic ... 101
Simple I/O ... 152
Size Clause ... 104
Size clauses ... 58
Size for biased representation ... 107
Size of Address ... 41
Size, of objects ... 107
Size, setting for not-first subtype ... 48

Size, used for objects ... 46
`Size`, VADS compatibility ... 38, 48
Size, variant record objects ... 106
`Small` ... 47
Sockets ... 153
Sorting ... 149, 152
Source Information ... 153
`Source_File_Name` ... 31
Source_Location ... 101
`Source_Reference` ... 31
Spawn capability ... 153
Spell checking ... 153
SPITBOL interface ... 154
SPITBOL pattern matching ... 154
SPITBOL Tables ... 154
`Static_Priorities` ... 85
`Static_Storage_Size` ... 85
Storage place attributes ... 60
Storage_Size Clause ... 105
`Storage_Unit` ... 7, 47
Stream files ... 137
Stream oriented attributes ... 62
`Stream_Convert` ... 32
String maps ... 154
`Style_Checks` ... 33
Subprogram address ... 42
`Subtitle` ... 33
`Suppress_All` ... 34
`Suppress_Initialization` ... 34
Suppressing initialization ... 34
Suppression of checks ... 55
Synchronization, OS/2 ... 156
`system`, extending ... 15
`System.Address_Image` (`s-addima.ads`) ... 156
`System.Assertions` (`s-assert.ads`) ... 156
`System.Partition_Interface` (`s-parint.ads`) ... 156
`System.Task_Info` (`s-tasinf.ads`) ... 156
`System.Wch_Cnv` (`s-wchcnv.ads`) ... 157
`System.Wch_Con` (`s-wchcon.ads`) ... 157

T

Table implementation ... 151, 154
Task locking ... 154
Task synchronization ... 154
`Task_Attributes` ... 69
`Task_Info` ... 34
Task_Info pragma ... 156
`Task_Name` ... 34
`Task_Storage` ... 35
Tasking restrictions ... 70
Text_IO ... 152
Text_IO extensions ... 137
Text_IO for unbounded strings ... 137
`Text_IO`, extensions for unbounded strings ... 148
`Text_IO`, extensions for unbounded wide strings ... 148
Thread control, OS/2 ... 156
Threads, foreign ... 155
`Tick` ... 47
Time ... 149
Time, monotonic ... 70
`Time_Slice` ... 35
`Title` ... 35
`To_Address` ... 47, 119
Trace back facilities ... 155
`Type_Class` ... 47
Typographical conventions ... 2

U

`UET_Address` ... 48
`Unbounded_String`, IO support ... 148
Unbounded_String, Text_IO operations ... 137
`Unbounded_Wide_String`, IO support ... 148
Unchecked conversion ... 61
Unchecked deallocation ... 62
`Unchecked_Union` ... 36
`Unimplemented_Unit` ... 37
Unions in C ... 36
`Universal_Literal_String` ... 48
`Unreferenced` ... 37
`Unreserve_All_Interrupts` ... 37
`Unrestricted_Access` ... 48
`Unsuppress` ... 37
`Use_VADS_Size` ... 38

V

`VADS_Size` ... 48
`Validity_Checks` ... 38
`Value_Size` ... 48, 107
Variant record objects, size ... 106
`Volatile` ... 39
VxWorks, Get_Immediate ... 156
VxWorks, I/O interfacing ... 156
VxWorks, interfacing ... 156

W

`Warnings` ... 39
Warnings, unreferenced ... 37
`Wchar_T_Size` ... 49
`Weak_External` ... 39
Wide Character, Representation ... 157
Wide String, Conversion ... 157
Windows Registry ... 153
`Word_Size` ... 49

Z

Zero address, passing ... 46
Zero Cost Exceptions ... 26

Table of Contents

About This Guide 1

What This Reference Manual Contains 1
Conventions 2
Related Information 2

1 Implementation Defined Pragmas 3

2 Implementation Defined Attributes 41

3 Implementation Advice 51

4 Implementation Defined Characteristics 75

5 Intrinsic Subprograms 99

5.1 Intrinsic Operators 99
5.2 Enclosing_Entity 99
5.3 Exception_Information 99
5.4 Exception_Message 99
5.5 Exception_Name 100
5.6 File 100
5.7 Line 100
5.8 Rotate_Left 100
5.9 Rotate_Right 100
5.10 Shift_Left 101
5.11 Shift_Right 101
5.12 Shift_Right_Arithmetic 101
5.13 Source_Location 101

6 Representation Clauses and Pragmas 103

6.1 Alignment Clauses 103
6.2 Size Clauses 104
6.3 Storage_Size Clauses 105
6.4 Size of Variant Record Objects 105
6.5 Biased Representation 107
6.6 Value_Size and Object_Size Clauses 107
6.7 Component_Size Clauses 110
6.8 Bit_Order Clauses 110
6.9 Effect of Bit_Order on Byte Ordering 111
6.10 Pragma Pack for Arrays 115
6.11 Pragma Pack for Records 115
6.12 Record Representation Clauses 116

6.13 Enumeration Clauses 117
6.14 Address Clauses 117
6.15 Effect of Convention on Representation 120
6.16 Determining the Representations chosen by GNAT 121

7 Standard Library Routines 125

8 The Implementation of Standard I/O 133

8.1 Standard I/O Packages 133
8.2 FORM Strings 133
8.3 Direct_IO 134
8.4 Sequential_IO 134
8.5 Text_IO 135
8.5.1 Stream Pointer Positioning 135
8.5.2 Reading and Writing Non-Regular Files 136
8.5.3 Get_Immediate 136
8.5.4 Treating Text_IO Files as Streams 137
8.5.5 Text_IO Extensions 137
8.5.6 Text_IO Facilities for Unbounded Strings 137
8.6 Wide_Text_IO 138
8.6.1 Stream Pointer Positioning 139
8.6.2 Reading and Writing Non-Regular Files 140
8.7 Stream_IO 140
8.8 Shared Files 140
8.9 Open Modes 141
8.10 Operations on C Streams 141
8.11 Interfacing to C Streams 144

9 The GNAT Library 147

9.1 `Ada.Characters.Latin_9` (‘`a-chlat9.ads`’) 147
9.2 `Ada.Characters.Wide_Latin_1` (‘`a-cwila1.ads`’) 147
9.3 `Ada.Characters.Wide_Latin_9` (‘`a-cwila1.ads`’) 147
9.4 `Ada.Command_Line.Remove` (‘`a-colire.ads`’) 147
9.5 `Ada.Direct_IO.C_Streams` (‘`a-diocst.ads`’) 148
9.6 `Ada.Exceptions.Is_Null_Occurrence` (‘`a-einuoc.ads`’) 148
9.7 `Ada.Sequential_IO.C_Streams` (‘`a-siocst.ads`’) 148
9.8 `Ada.Streams.Stream_IO.C_Streams` (‘`a-ssicst.ads`’) 148
9.9 `Ada.Strings.Unbounded.Text_IO` (‘`a-suteio.ads`’) 148
9.10 `Ada.Strings.Wide_Unbounded.Wide_Text_IO` (‘`a-swuwti.ads`’) 148
9.11 `Ada.Text_IO.C_Streams` (‘`a-tiocst.ads`’) 148
9.12 `Ada.Wide_Text_IO.C_Streams` (‘`a-wtcstr.ads`’) 149
9.13 `GNAT.AWK` (‘`g-awk.ads`’) 149
9.14 `GNAT.Bubble_Sort_A` (‘`g-busora.ads`’) 149
9.15 `GNAT.Bubble_Sort_G` (‘`g-busorg.ads`’) 149
9.16 `GNAT.Calendar` (‘`g-calend.ads`’) 149

9.17 GNAT.Calendar.Time_IO (‘g-catiio.ads’) 149
9.18 GNAT.CRC32 (‘g-crc32.ads’) 149
9.19 GNAT.Case_Util (‘g-casuti.ads’) 150
9.20 GNAT.CGI (‘g-cgi.ads’) 150
9.21 GNAT.CGI.Cookie (‘g-cgicoo.ads’) 150
9.22 GNAT.CGI.Debug (‘g-cgideb.ads’) 150
9.23 GNAT.Command_Line (‘g-comlin.ads’) 150
9.24 GNAT.Current_Exception (‘g-curexc.ads’) 150
9.25 GNAT.Debug_Pools (‘g-debpoo.ads’) 150
9.26 GNAT.Debug_Utilities (‘g-debuti.ads’) 151
9.27 GNAT.Directory_Operations (g-dirope.ads) 151
9.28 GNAT.Dynamic_Tables (‘g-dyntab.ads’) 151
9.29 GNAT.Exception_Traces (‘g-exctra.ads’) 151
9.30 GNAT.Expect (‘g-expect.ads’) 151
9.31 GNAT.Float_Control (‘g-flocon.ads’) 151
9.32 GNAT.Heap_Sort_A (‘g-hesora.ads’) 151
9.33 GNAT.Heap_Sort_G (‘g-hesorg.ads’) 152
9.34 GNAT.HTable (‘g-htable.ads’) 152
9.35 GNAT.IO (‘g-io.ads’) 152
9.36 GNAT.IO_Aux (‘g-io_aux.ads’) 152
9.37 GNAT.Lock_Files (‘g-locfil.ads’) 152
9.38 GNAT.MD5 (‘g-md5.ads’) 152
9.39 GNAT.Most_Recent_Exception (‘g-moreex.ads’) 152
9.40 GNAT.OS_Lib (‘g-os_lib.ads’) 152
9.41 GNAT.Regexp (‘g-regexp.ads’) 153
9.42 GNAT.Registry (‘g-regist.ads’) 153
9.43 GNAT.Regpat (‘g-regpat.ads’) 153
9.44 GNAT.Sockets (‘g-socket.ads’) 153
9.45 GNAT.Source_Info (‘g-souinf.ads’) 153
9.46 GNAT.Spell_Checker (‘g-speche.ads’) 153
9.47 GNAT.Spitbol.Patterns (‘g-spipat.ads’) 153
9.48 GNAT.Spitbol (‘g-spitbo.ads’) 154
9.49 GNAT.Spitbol.Table_Boolean (‘g-sptabo.ads’) 154
9.50 GNAT.Spitbol.Table_Integer (‘g-sptain.ads’) 154
9.51 GNAT.Spitbol.Table_VString (‘g-sptavs.ads’) 154
9.52 GNAT.Table (‘g-table.ads’) 154
9.53 GNAT.Task_Lock (‘g-tasloc.ads’) 154
9.54 GNAT.Threads (‘g-thread.ads’) 154
9.55 GNAT.Traceback (‘g-traceb.ads’) 155
9.56 GNAT.Traceback.Symbolic (‘g-trasym.ads’) 155
9.57 Interfaces.C.Extensions (‘i-cexten.ads’) 155
9.58 Interfaces.C.Streams (‘i-cstrea.ads’) 155
9.59 Interfaces.CPP (‘i-cpp.ads’) 155
9.60 Interfaces.Os2lib (‘i-os2lib.ads’) 155
9.61 Interfaces.Os2lib.Errors (‘i-os2err.ads’) 155
9.62 Interfaces.Os2lib.Synchronization (‘i-os2syn.ads’) 155
9.63 Interfaces.Os2lib.Threads (‘i-os2thr.ads’) 156

9.64 `Interfaces.Packed_Decimal` (`i-pacdec.ads`) 156
9.65 `Interfaces.VxWorks` (`i-vxwork.ads`) 156
9.66 `Interfaces.VxWorks.IO` (`i-vxwoio.ads`) 156
9.67 `System.Address_Image` (`s-addima.ads`) 156
9.68 `System.Assertions` (`s-assert.ads`) 156
9.69 `System.Partition_Interface` (`s-parint.ads`) 156
9.70 `System.Task_Info` (`s-tasinf.ads`) 156
9.71 `System.Wch_Cnv` (`s-wchcnv.ads`) 156
9.72 `System.Wch_Con` (`s-wchcon.ads`) 157

10 Interfacing to Other Languages 159
10.1 Interfacing to C 159
10.2 Interfacing to C++ 159
10.3 Interfacing to COBOL 160
10.4 Interfacing to Fortran 160
10.5 Interfacing to non-GNAT Ada code 160

11 Machine Code Insertions 161

12 GNAT Implementation of Tasking 163
12.1 Mapping Ada Tasks onto the Underlying Kernel Threads 163
12.2 Ensuring Compliance with the Real-Time Annex 163

13 Code generation for array aggregates 165
13.1 Static constant aggregates with static bounds 165
13.2 Constant aggregates with an unconstrained nominal types 165
13.3 Aggregates with static bounds 165
13.4 Aggregates with non-static bounds 166
13.5 Aggregates in assignments statements 166

14 Specialized Needs Annexes 167

15 Compatibility Guide 169
15.1 Compatibility with Ada 83 169
15.2 Compatibility with Other Ada 95 Systems 170
15.3 Representation Clauses 171
15.4 Compatibility with DEC Ada 83 172

GNU Free Documentation License 175
ADDENDUM: How to use this License for your documents 181

Index 183