

Ada

for the Embedded C Developer

Quentin Ochem

Robert Tice

Gustavo A. Hoffmann

Patrick Rogers

LEARN.
ADACORE.COM

Ada for the Embedded C Developer

Release 2021-06

**Quentin Ochem
Gustavo A. Hoffmann**

**Robert Tice
Patrick Rogers.**

Jun 11, 2021

CONTENTS

1	Introduction	3
1.1	So, what is this Ada thing anyway?	3
1.2	Ada — The Technical Details	5
2	The C Developer's Perspective on Ada	7
2.1	What we mean by Embedded Software	7
2.2	The GNAT Toolchain	7
2.3	The GNAT Toolchain for Embedded Targets	7
2.4	Hello World in Ada	8
2.5	The Ada Syntax	10
2.6	Compilation Unit Structure	11
2.7	Packages	11
2.7.1	Declaration Protection	11
2.7.2	Hierarchical Packages	12
2.7.3	Using Entities from Packages	13
2.8	Statements and Declarations	13
2.9	Conditions	19
2.10	Loops	22
2.11	Type System	29
2.11.1	Strong Typing	29
2.11.2	Language-Defined Types	32
2.11.3	Application-Defined Types	32
2.11.4	Type Ranges	35
2.11.5	Unsigned And Modular Types	38
2.11.6	Attributes	41
2.11.7	Arrays and Strings	44
2.11.8	Heterogeneous Data Structures	51
2.11.9	Pointers	53
2.12	Functions and Procedures	59
2.12.1	General Form	59
2.12.2	Overloading	62
2.12.3	Aspects	65
3	Concurrency and Real-Time	67
3.1	Understanding the various options	67
3.2	Tasks	67
3.3	Rendezvous	70
3.4	Selective Rendezvous	72
3.5	Protected Objects	74
3.6	Ravenscar	78
4	Writing Ada on Embedded Systems	81
4.1	Understanding the Ada Run-Time	81
4.2	Low Level Programming	82

4.2.1	Representation Clauses	82
4.2.2	Embedded Assembly Code	83
4.3	Interrupt Handling	84
4.4	Dealing with Absence of FPU with Fixed Point	86
4.5	Volatile and Atomic data	90
4.5.1	Volatile	91
4.5.2	Atomic	93
4.6	Interfacing with Devices	95
4.6.1	Size aspect and attribute	95
4.6.2	Register overlays	96
4.6.3	Data streams	99
4.7	ARM and svd2ada	105
5	Enhancing Verification with SPARK and Ada	107
5.1	Understanding Exceptions and Dynamic Checks	107
5.2	Understanding Dynamic Checks versus Formal Proof	113
5.3	Initialization and Correct Data Flow	116
5.4	Contract-Based Programming	118
5.5	Replacing Defensive Code	120
5.6	Proving Absence of Run-Time Errors	122
5.7	Proving Abstract Properties	123
5.8	Final Comments	123
6	C to Ada Translation Patterns	125
6.1	Naming conventions and casing considerations	125
6.2	Manually interfacing C and Ada	125
6.3	Building and Debugging mixed language code	127
6.4	Automatic interfacing	128
6.5	Using Arrays in C interfaces	128
6.6	By-value vs. by-reference types	130
6.7	Naming and prefixes	131
6.8	Pointers	132
6.9	Bitwise Operations	135
6.10	Mapping Structures to Bit-Fields	137
6.10.1	Overlays vs. Unchecked Conversions	150
7	Handling Variability and Re-usability	155
7.1	Understanding static and dynamic variability	155
7.2	Handling variability & reusability statically	155
7.2.1	Genericity	155
7.2.2	Simple derivation	158
7.2.3	Configuration pragma files	164
7.2.4	Configuration packages	165
7.3	Handling variability & reusability dynamically	170
7.3.1	Records with discriminants	170
7.3.2	Variant records	172
7.3.2.1	Variant records and unions	174
7.3.2.2	Optional components	176
7.3.2.3	Optional output information	177
7.3.3	Object orientation	180
7.3.3.1	Type extension	180
7.3.3.2	Overriding subprograms	182
7.3.3.3	Comparing untagged and tagged types	182
7.3.3.4	Dispatching calls	185
7.3.3.5	Interfaces	186
7.3.3.6	Deriving from multiple interfaces	188
7.3.3.7	Abstract tagged types	189
7.3.3.8	From simple derivation to OOP	191
7.3.3.9	Further resources	194

7.3.4	Pointer to subprograms	194
7.4	Design by components using dynamic libraries	199
8	Performance considerations	203
8.1	Overall expectations	203
8.2	Switches and optimizations	203
8.2.1	Optimizations levels	203
8.2.2	Inlining	204
8.3	Checks and assertions	205
8.3.1	Checks	205
8.3.2	Assertions	209
8.4	Dynamic vs. static structures	209
8.5	Pointers vs. data copies	210
8.5.1	Function returns	213
9	Argumentation and Business Perspectives	217
9.1	What's the expected ROI of a C to Ada transition?	217
9.2	Who is using Ada today?	218
9.3	What is the future of the Ada technology?	218
9.4	Is the Ada toolset complete?	219
9.5	Where can I find Ada or SPARK developers?	219
9.6	How to introduce Ada and SPARK in an existing code base?	220
10	Conclusion	221
11	Appendix A: Hands-On Object-Oriented Programming	225
11.1	System Overview	225
11.2	Non Object-Oriented Approach	226
11.2.1	Starting point in C	226
11.2.2	Initial translation to Ada	229
11.2.3	Improved Ada implementation	233
11.3	First Object-Oriented Approach	237
11.3.1	Interfaces	237
11.3.2	Base type	238
11.3.3	Derived types	238
11.3.4	Subprograms from parent	239
11.3.5	Type AB	240
11.3.6	Updated source-code	241
11.4	Further Improvements	245
11.4.1	Dispatching calls	245
11.4.2	Dynamic allocation	246
11.4.3	Limited controlled types	247
11.4.4	Updated source-code	248
	Bibliography	253

Copyright © 2020 – 2021, AdaCore

This book is published under a CC BY-SA license, which means that you can copy, redistribute, remix, transform, and build upon the content for any purpose, even commercially, as long as you give appropriate credit, provide a link to the license, and indicate if changes were made. If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original. You can find license details [on this page](http://creativecommons.org/licenses/by-sa/4.0)¹



This course introduces you to the Ada language by comparing it to C. It assumes that you have good knowledge of the C language. It also assumes that the choice of learning Ada is guided by considerations linked to reliability, safety or security. In that sense, it teaches you Ada paradigms that should be applied in replacement of those usually applied in C.

This course also introduces you to the SPARK subset of the Ada programming language, which removes a few features of the language with undefined behavior, so that the code is fit for sound static analysis techniques.

This course was written by Quentin Ochem, Robert Tice, Gustavo A. Hoffmann, and Patrick Rogers and reviewed by Patrick Rogers, Filip Gajowniczek, and Tucker Taft.

¹ <http://creativecommons.org/licenses/by-sa/4.0>

INTRODUCTION

1.1 So, what is this Ada thing anyway?

To answer this question let's introduce Ada as it compares to C for an embedded application. C developers are used to a certain coding semantic and style of programming. Especially in the embedded domain, developers are used to working at a very low level near the hardware to directly manipulate memory and registers. Normal operations involve mathematical operations on pointers, complex bit shifts, and logical bitwise operations. C is well designed for such operations as it is a low level language that was designed to replace assembly language for faster, more efficient programming. Because of this minimal abstraction, the programmer has to model the data that represents the problem they are trying to solve using the language of the physical hardware.

Let's look at an example of this problem in action by comparing the same program in Ada and C:

[C]

Listing 1: main.c

```
1  #include <stdio.h>
2  #include <stdlib.h>
3
4  #define DEGREES_MAX      (360)
5  typedef unsigned int degrees;
6
7  #define MOD_DEGREES(x)   (x % DEGREES_MAX)
8
9  degrees add_angles(degrees* list, int length)
10 {
11     degrees sum = 0;
12     for(int i = 0; i < length; ++i) {
13         sum += list[i];
14     }
15
16     return sum;
17 }
18
19 int main(int argc, char** argv)
20 {
21     degrees list[argc - 1];
22
23     for(int i = 1; i < argc; ++i) {
24         list[i - 1] = MOD_DEGREES(atoi(argv[i]));
25     }
26
27     printf("Sum: %d\n", add_angles(list, argc - 1));
28
29     return 0;
30 }
```

Runtime output

```
Sum: 0
```

[Ada]

Listing 2: sum_angles.adb

```
1 with Ada.Command_Line; use Ada.Command_Line;
2 with Ada.Text_IO; use Ada.Text_IO;
3
4 procedure Sum_Angles is
5
6     DEGREES_MAX : constant := 360;
7     type Degrees is mod DEGREES_MAX;
8
9     type Degrees_List is array (Natural range <>) of Degrees;
10
11     function Add_Angles (List : Degrees_List) return Degrees
12     is
13         Sum : Degrees := 0;
14     begin
15         for I in List'Range loop
16             Sum := Sum + List (I);
17         end loop;
18
19         return Sum;
20     end Add_Angles;
21
22     List : Degrees_List (1 .. Argument_Count);
23 begin
24     for I in List'Range loop
25         List (I) := Degrees (Integer'Value (Argument (I)));
26     end loop;
27
28     Put_Line ("Sum:" & Add_Angles (List)'Img);
29 end Sum_Angles;
```

Build output

```
Compile
  [Ada]          sum_angles.adb
Bind
  [gprbind]      sum_angles.bexch
  [Ada]          sum_angles.ali
Link
  [link]         sum_angles.adb
```

Runtime output

```
Sum: 0
```

Here we have a piece of code in C and in Ada that takes some numbers from the command line and stores them in an array. We then sum all of the values in the array and print the result. The tricky part here is that we are working with values that model an angle in degrees. We know that angles are modular types, meaning that angles greater than 360° can also be represented as $\text{Angle} \bmod 360$. So if we have an angle of 400° , this is equivalent to 40° . In order to model this behavior in C we had to create the `MOD_DEGREES` macro, which performs the modulus operation. As we read values from the command line, we convert them to integers and perform the modulus before storing them into the array. We then call `add_angles` which returns the sum of the values in the array. Can you spot the problem with the C code?

Try running the Ada and C examples using the input sequence `340 2 50 70`. What does the C

program output? What does the Ada program output? Why are they different?

The problem with the C code is that we forgot to call `MOD_DEGREES` in the for loop of `add_angles`. This means that it is possible for `add_angles` to return values greater than `DEGREES_MAX`. Let's look at the equivalent Ada code now to see how Ada handles the situation. The first thing we do in the Ada code is to create the type `Degrees` which is a modular type. This means that the compiler is going to handle performing the modulus operation for us. If we use the same for loop in the `Add_Angles` function, we can see that we aren't doing anything special to make sure that our resulting value is within the 360° range we need it to be in.

The takeaway from this example is that Ada tries to abstract some concepts from the developer so that the developer can focus on solving the problem at hand using a data model that models the real world rather than using data types prescribed by the hardware. The main benefit of this is that the compiler takes some responsibility from the developer for generating correct code. In this example we forgot to put in a check in the C code. The compiler inserted the check for us in the Ada code because we told the compiler what we were trying to accomplish by defining strong types.

Ideally, we want all the power that the C programming language can give us to manipulate the hardware we are working on while also allowing us the ability to more accurately model data in a safe way. So, we have a dilemma; what can give us the power of operations like the C language, but also provide us with features that can minimize the potential for developer error? Since this course is about Ada, it's a good bet we're about to introduce the Ada language as the answer to this question...

Unlike C, the Ada language was designed as a higher level language from its conception; giving more responsibility to the compiler to generate correct code. As mentioned above, with C, developers are constantly shifting, masking, and accessing bits directly on memory pointers. In Ada, all of these operations are possible, but in most cases, there is a better way to perform these operations using higher level constructs that are less prone to mistakes, like off-by-one or unintentional buffer overflows. If we were to compare the same application written using C and with Ada using high level constructs, we would see similar performance in terms of speed and memory efficiency. If we compare the object code generated by both compilers, it's possible that they even look identical!

1.2 Ada — The Technical Details

Like C, Ada is a compiled language. This means that the compiler will parse the source code and emit machine code native to the target hardware. The Ada compiler we will be discussing in this course is the GNAT compiler. This compiler is based on the GCC technology like many C and C++ compilers available. When the GNAT compiler is invoked on Ada code, the GNAT front-end expands and translates the Ada code into an intermediate language which is passed to GCC where the code is optimized and translated to machine code. A C compiler based on GCC performs the same steps and uses the same intermediate GCC representation. This means that the optimizations we are used to seeing with a GCC based C compiler can also be applied to Ada code. The main difference between the two compilers is that the Ada compiler is expanding high level constructs into intermediate code. After expansion, the Ada code will be very similar to the equivalent C code.

It is possible to do a line-by-line translation of C code to Ada. This feels like a natural step for a developer used to C paradigms. However, there may be very little benefit to doing so. For the purpose of this course, we're going to assume that the choice of Ada over C is guided by considerations linked to reliability, safety or security. In order to improve upon the reliability, safety and security of our application, Ada paradigms should be applied in replacement of those usually applied in C. Constructs such as pointers, preprocessor macros, bitwise operations and defensive code typically get expressed in Ada in very different ways, improving the overall reliability and readability of the applications. Learning these new ways of coding, often, requires effort by the developer at first, but proves more efficient once the paradigms are understood.

In this course we will also introduce the SPARK subset of the Ada programming language. The SPARK subset removes a few features of the language, i.e., those that make proof difficult, such as

pointer aliasing. By removing these features we can write code that is fit for sound static analysis techniques. This means that we can run mathematical provers on the SPARK code to prove certain safety or security properties about the code.

THE C DEVELOPER'S PERSPECTIVE ON ADA

2.1 What we mean by Embedded Software

The Ada programming language is a general programming language, which means it can be used for many different types of applications. One type of application where it particularly shines is reliable and safety-critical embedded software; meaning, a platform with a microprocessor such as ARM, PowerPC, x86, or RISC-V. The application may be running on top of an embedded operating system, such as an embedded Linux, or directly on bare metal. And the application domain can range from small entities such as firmware or device controllers to flight management systems, communication based train control systems, or advanced driver assistance systems.

2.2 The GNAT Toolchain

The toolchain used throughout this course is called GNAT, which is a suite of tools with a compiler based on the GCC environment. It can be obtained from AdaCore, either as part of a commercial contract with [GNAT Pro](#)² or at no charge with the [GNAT Community edition](#)³. The information in this course will be relevant no matter which edition you're using. Most examples will be runnable on the native Linux or Windows version for convenience. Some will only be relevant in the context of a cross toolchain, in which case we'll be using the embedded ARM bare metal toolchain.

As for any Ada compiler, GNAT takes advantage of implementation permissions and offers a project management system. Because we're talking about embedded platforms, there are a lot of topics that we'll go over which will be specific to GNAT, and sometimes to specific platforms supported by GNAT. We'll try to make the distinction between what is GNAT-specific and Ada generic as much as possible throughout this course.

For an introduction to the GNAT Toolchain for the GNAT Community edition, you may refer to the [Introduction to GNAT Toolchain](#)⁴ course.

2.3 The GNAT Toolchain for Embedded Targets

When we're discussing embedded programming, our target device is often different from the host, which is the device we're using to actually write and build an application. In this case, we're talking about cross compilation platforms (concisely referred to as cross platforms).

The GNAT toolchain supports cross platform compilation for various target devices. This section provides a short introduction to the topic. For more details, please refer to the [GNAT User's Guide Supplement for Cross Platforms](#)⁵

² <https://www.adacore.com/gnatpro>

³ <https://www.adacore.com/community>

⁴ https://learn.adacore.com/courses/GNAT_Toolchain_Intro/index.html

⁵ https://docs.adacore.com/gnat_ugx-docs/html/gnat_ugx/gnat_ugx.html

GNAT supports two types of cross platforms:

- **cross targets**, where the target device has an embedded operating system.
 - ARM-Linux, which is commonly found in a Raspberry-Pi, is a prominent example.
- **bareboard targets**, where the run-times do not depend on an operating system.
 - In this case, the application has direct access to the system hardware.

For each platform, a set of run-time libraries is available. Run-time libraries implement a subset of the Ada language for different use cases, and they're different for each target platform. They may be selected via an attribute in the project's GPR project file or as a command-line switch to **GPRbuild**. Although the run-time libraries may vary from target to target, the user interface stays the same, providing portability for the application.

Run-time libraries consists of:

1. Files that are dependent on the target board.
 - These files are responsible for configuring and interacting with the hardware.
 - They are known as a Board Support Package — commonly referred to by their abbreviation *BSP*.
2. Code that is target-independent.
 - This code implements language-defined functionality.

The bareboard run-time libraries are provided as customized run-times that are configured to target a very specific micro-controller or processor. Therefore, for different micro-controllers and processors, the run-time libraries need to be ported to the specific target. These are some examples of what needs to be ported:

- startup code / scripts;
- clock frequency initializations;
- memory mapping / allocation;
- interrupts and interrupt priorities;
- register descriptions.

For more details on the topic, please refer to the following chapters of the [GNAT User's Guide Supplement for Cross Platforms](#)⁶:

- [Bareboard Topics](#)⁷
- [Customized Run-Time Libraries](#)⁸

2.4 Hello World in Ada

The first piece of code to translate from C to Ada is the usual Hello World program:

[C]

Listing 1: main.c

```
1 #include <stdio.h>
2
3 int main(int argc, const char * argv[])
4 {
```

(continues on next page)

⁶ https://docs.adacore.com/gnat_ugx-docs/html/gnat_ugx/gnat_ugx.html

⁷ http://docs.adacore.com/live/wave/gnat_ugx/html/gnat_ugx/gnat_ugx/bareboard_topics.html

⁸ http://docs.adacore.com/live/wave/gnat_ugx/html/gnat_ugx/gnat_ugx/customized_run-time_libraries.html

(continued from previous page)

```

5     printf("Hello World\n");
6     return 0;
7 }

```

Runtime output

Hello World

[Ada]

Listing 2: hello_world.adb

```

1  with Ada.Text_IO;
2
3  procedure Hello_World
4  is
5  begin
6      Ada.Text_IO.Put_Line ("Hello World");
7  end Hello_World;

```

Build output

Compile	
[Ada]	hello_world.adb
Bind	
[gprbind]	hello_world.bexch
[Ada]	hello_world.ali
Link	
[link]	hello_world.adb

Runtime output

Hello World

The resulting program will print Hello World on the screen. Let's now dissect the Ada version to describe what is going on:

The first line of the Ada code is giving us access to the `Ada.Text_IO` library which contains the `Put_Line` function we will use to print the text to the console. This is similar to C's `#include <stdio.h>`. We then create a procedure which executes `Put_Line` which prints to the console. This is similar to C's `printf` function. For now, we can assume these Ada and C features have similar functionality. In reality, they are very different. We will explore that more as we delve further into the Ada language.

You may have noticed that the Ada syntax is more verbose than C. Instead of using braces `{}` to declare scope, Ada uses keywords. `is` opens a declarative scope — which is empty here as there's no variable to declare. `begin` opens a sequence of statements. Within this sequence, we're calling the function `Put_Line`, prefixing explicitly with the name of the library unit where it's declared, `Ada.Text_IO`. The absence of the end of line `\n` can also be noted, as `Put_Line` always terminates by an end of line.

2.5 The Ada Syntax

Ada syntax might seem peculiar at first glance. Unlike many other languages, it's not derived from the popular C style of notation with its ample use of brackets; rather, it uses a more expository syntax coming from Pascal. In many ways, Ada is a more explicit language — its syntax was designed to increase readability and maintainability, rather than making it faster to write in a condensed manner. For example:

- full words like `begin` and `end` are used in place of curly braces.
- Conditions are written using `if`, `then`, `elsif`, `else`, and `end if`.
- Ada's assignment operator does not double as an expression, eliminating potential mistakes that could be caused by `=` being used where `==` should be.

All languages provide one or more ways to express comments. In Ada, two consecutive hyphens `--` mark the start of a comment that continues to the end of the line. This is exactly the same as using `//` for comments in C. Multi line comments like C's `/* */` do not exist in Ada.

Ada compilers are stricter with type and range checking than most C programmers are used to. Most beginning Ada programmers encounter a variety of warnings and error messages when coding, but this helps detect problems and vulnerabilities at compile time — early on in the development cycle. In addition, checks (such as array bounds checks) provide verification that could not be done at compile time but can be performed either at run-time, or through formal proof (with the SPARK tooling).

Ada identifiers and reserved words are case insensitive. The identifiers `VAR`, `var` and `VaR` are treated as the same identifier; likewise `begin`, `BEGIN`, `Begin`, etc. Identifiers may include letters, digits, and underscores, but must always start with a letter. There are 73 reserved keywords in Ada that may not be used as identifiers, and these are:

<code>abort</code>	<code>else</code>	<code>null</code>	<code>select</code>
<code>abs</code>	<code>elsif</code>	<code>of</code>	<code>separate</code>
<code>abstract</code>	<code>end</code>	<code>or</code>	<code>some</code>
<code>accept</code>	<code>entry</code>	<code>others</code>	<code>subtype</code>
<code>access</code>	<code>exception</code>	<code>out</code>	<code>synchronized</code>
<code>aliased</code>	<code>exit</code>	<code>overriding</code>	<code>tagged</code>
<code>all</code>	<code>for</code>	<code>package</code>	<code>task</code>
<code>and</code>	<code>function</code>	<code>pragma</code>	<code>terminate</code>
<code>array</code>	<code>generic</code>	<code>private</code>	<code>then</code>
<code>at</code>	<code>goto</code>	<code>procedure</code>	<code>type</code>
<code>begin</code>	<code>if</code>	<code>protected</code>	<code>until</code>
<code>body</code>	<code>in</code>	<code>raise</code>	<code>use</code>
<code>case</code>	<code>interface</code>	<code>range</code>	<code>when</code>
<code>constant</code>	<code>is</code>	<code>record</code>	<code>while</code>
<code>declare</code>	<code>limited</code>	<code>rem</code>	<code>with</code>
<code>delay</code>	<code>loop</code>	<code>renames</code>	<code>xor</code>
<code>delta</code>	<code>mod</code>	<code>requeue</code>	
<code>digits</code>	<code>new</code>	<code>return</code>	
<code>do</code>	<code>not</code>	<code>reverse</code>	

2.6 Compilation Unit Structure

Both C and Ada were designed with the idea that the code specification and code implementation could be separated into two files. In C, the specification typically lives in the .h, or header file, and the implementation lives in the .c file. Ada is superficially similar to C. With the GNAT toolchain, compilation units are stored in files with an .ads extension for specifications and with an .adb extension for implementations.

One main difference between the C and Ada compilation structure is that Ada compilation units are structured into something called packages.

2.7 Packages

The package is the basic modularization unit of the Ada language, as is the class for Java and the header and implementation pair for C. A specification defines a package and the implementation implements the package. We saw this in an earlier example when we included the Ada.Text_IO package into our application. The package specification has the structure:

[Ada]

```
-- my_package.ads
package My_Package is
    -- public declarations

private
    -- private declarations
end My_Package;
```

The package implementation, or body, has the structure:

```
-- my_package.adb
package body My_Package is
    -- implementation
end My_Package;
```

2.7.1 Declaration Protection

An Ada package contains three parts that, for GNAT, are separated into two files: .ads files contain public and private Ada specifications, and .adb files contain the implementation, or Ada bodies.

[Ada]

```
package Package_Name is
    -- public specifications
private
    -- private specifications
end Package_Name;

package body Package_Name is
    -- implementation
end Package_Name;
```

Private types are useful for preventing the users of a package's types from depending on the types' implementation details. Another use-case is the prevention of package users from accessing package state/data arbitrarily. The private reserved word splits the package spec into *public* and *private* parts. For example:

[Ada]

Listing 3: types.ads

```
1 package Types is
2     type Type_1 is private;
3     type Type_2 is private;
4     type Type_3 is private;
5     procedure P (X : Type_1);
6     -- ...
7 private
8     procedure Q (Y : Type_1);
9     type Type_1 is new Integer range 1 .. 1000;
10    type Type_2 is array (Integer range 1 .. 1000) of Integer;
11    type Type_3 is record
12        A, B : Integer;
13    end record;
14 end Types;
```

Subprograms declared above the `private` separator (such as `P`) will be visible to the package user, and the ones below (such as `Q`) will not. The body of the package, the implementation, has access to both parts. A package specification does not require a private section.

2.7.2 Hierarchical Packages

Ada packages can be organized into hierarchies. A child unit can be declared in the following way:

[Ada]

```
-- root-child.ads

package Root.Child is
    -- package spec goes here
end Root.Child;

-- root-child.adb

package body Root.Child is
    -- package body goes here
end Root.Child;
```

Here, `Root.Child` is a child package of `Root`. The public part of `Root.Child` has access to the public part of `Root`. The private part of `Child` has access to the private part of `Root`, which is one of the main advantages of child packages. However, there is no visibility relationship between the two bodies. One common way to use this capability is to define subsystems around a hierarchical naming scheme.

2.7.3 Using Entities from Packages

Entities declared in the visible part of a package specification can be made accessible using a `with` clause that references the package, which is similar to the C `#include` directive. After a `with` clause makes a package available, references to the package contents require the name of the package as a prefix, with a dot after the package name. This prefix can be omitted if a `use` clause is employed.

[Ada]

Listing 4: pck.ads

```
1  -- pck.ads
2
3  package Pck is
4      My_Glob : Integer;
5  end Pck;
```

Listing 5: main.adb

```
1  -- main.adb
2
3  with Pck;
4
5  procedure Main is
6  begin
7      Pck.My_Glob := 0;
8  end Main;
```

Build output

```
Compile
  [Ada]      main.adb
  [Ada]      pck.ads
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb
```

In contrast to C, the Ada `with` clause is a *semantic inclusion* mechanism rather than a *text inclusion* mechanism; for more information on this difference please refer to [Packages⁹](#).

2.8 Statements and Declarations

The following code samples are all equivalent, and illustrate the use of comments and working with integer variables:

[C]

Listing 6: main.c

```
1  #include <stdio.h>
2
3  int main(int argc, const char * argv[])
4  {
5      // variable declarations
```

(continues on next page)

⁹ https://learn.adacore.com/courses/intro-to-ada/chapters/modular_programming.html

(continued from previous page)

```
6  int a = 0, b = 0, c = 100, d;  
7  
8  // c shorthand for increment  
9  a++;  
10  
11 // regular addition  
12 d = a + b + c;  
13  
14 // printing the result  
15 printf("d = %d\n", d);  
16  
17 return 0;  
18 }
```

Runtime output

```
d = 101
```

[Ada]

Listing 7: main.adb

```
1  with Ada.Text_IO;  
2  
3  procedure Main  
4  is  
5      -- variable declaration  
6      A, B : Integer := 0;  
7      C    : Integer := 100;  
8      D    : Integer;  
9  begin  
10     -- Ada does not have a shortcut format for increment like in C  
11     A := A + 1;  
12  
13     -- regular addition  
14     D := A + B + C;  
15  
16     -- printing the result  
17     Ada.Text_IO.Put_Line ("D =" & D'Img);  
18 end Main;
```

Build output

```
Compile  
[Ada]          main.adb  
Bind  
[gprbind]      main.bexch  
[Ada]          main.ali  
Link  
[link]         main.adb
```

Runtime output

```
D = 101
```

You'll notice that, in both languages, statements are terminated with a semicolon. This means that you can have multi-line statements.

The shortcuts of incrementing and decrementing

You may have noticed that Ada does not have something similar to the `a++` or `a--` operators.

Instead you must use the full assignment `A := A + 1` or `A := A - 1`.

In the Ada example above, there are two distinct sections to the procedure `Main`. This first section is delimited by the `is` keyword and the `begin` keyword. This section is called the declarative block of the subprogram. The declarative block is where you will define all the local variables which will be used in the subprogram. C89 had something similar, where developers were required to declare their variables at the top of the scope block. Most C developers may have run into this before when trying to write a for loop:

[C]

Listing 8: main.c

```

1  /* The C89 version */
2
3  #include <stdio.h>
4
5  int average(int* list, int length)
6  {
7      int i;
8      int sum = 0;
9
10     for(i = 0; i < length; ++i) {
11         sum += list[i];
12     }
13     return (sum / length);
14 }
15
16 int main(int argc, const char * argv[])
17 {
18     int vals[] = { 2, 2, 4, 4 };
19
20     printf("Average: %d\n", average(vals, 4));
21
22     return 0;
23 }
```

Runtime output

```
Average: 3
```

[C]

Listing 9: main.c

```

1  // The modern C way
2
3  #include <stdio.h>
4
5  int average(int* list, int length)
6  {
7      int sum = 0;
8
9      for(int i = 0; i < length; ++i) {
10         sum += list[i];
11     }
12
13     return (sum / length);
14 }
15
16 int main(int argc, const char * argv[])
17 {
```

(continues on next page)

(continued from previous page)

```
18  int vals[] = { 2, 2, 4, 4 };
19
20  printf("Average: %d\n", average(vals, 4));
21
22  return 0;
23 }
```

Runtime output

```
Average: 3
```

For the fun of it, let's also see the Ada way to do this:

[Ada]

Listing 10: main.adb

```
1  with Ada.Text_IO;
2
3  procedure Main is
4      type Int_Array is array (Natural range <>) of Integer;
5
6      function Average (List : Int_Array) return Integer
7      is
8          Sum : Integer := 0;
9      begin
10         for I in List'Range loop
11             Sum := Sum + List (I);
12         end loop;
13
14         return (Sum / List'Length);
15     end Average;
16
17     Vals : constant Int_Array (1 .. 4) := (2, 2, 4, 4);
18 begin
19     Ada.Text_IO.Put_Line ("Average: " & Integer'Image (Average (Vals)));
20 end Main;
```

Build output

```
Compile
[Ada]      main.adb
Bind
[gprbind]  main.bexch
[Ada]      main.ali
Link
[link]     main.adb
```

Runtime output

```
Average:  3
```

We will explore more about the syntax of loops in Ada in a future section of this course; but for now, notice that the I variable used as the loop index is not declared in the declarative section!

Declaration Flippy Floppy

Something peculiar that you may have noticed about declarations in Ada is that they are backwards from the way C does declarations. The C language expects the type followed by the variable name. Ada expects the variable name followed by a semicolon and then the type.

The next block in the Ada example is between the `begin` and `end` keywords. This is where your statements will live. You can create new scopes by using the `declare` keyword:

[Ada]

Listing 11: main.adb

```

1  with Ada.Text_IO;
2
3  procedure Main
4  is
5      -- variable declaration
6      A, B : Integer := 0;
7      C   : Integer := 100;
8      D   : Integer;
9  begin
10     -- Ada does not have a shortcut format for increment like in C
11     A := A + 1;
12
13     -- regular addition
14     D := A + B + C;
15
16     -- printing the result
17     Ada.Text_IO.Put_Line ("D =" & D'Img);
18
19     declare
20         E : constant Integer := D * 100;
21     begin
22         -- printing the result
23         Ada.Text_IO.Put_Line ("E =" & E'Img);
24     end;
25
26 end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```

D = 101
E = 10100
```

Notice that we declared a new variable `E` whose scope only exists in our newly defined block. The equivalent C code is:

[C]

Listing 12: main.c

```

1  #include <stdio.h>
2
3  int main(int argc, const char * argv[])
4  {
5      // variable declarations
6      int a = 0, b = 0, c = 100, d;
7
```

(continues on next page)

(continued from previous page)

```

8      // c shorthand for increment
9      a++;
10
11     // regular addition
12     d = a + b + c;
13
14     // printing the result
15     printf("d = %d\n", d);
16
17     {
18         const int e = d * 100;
19         printf("e = %d\n", e);
20     }
21
22     return 0;
23 }

```

Runtime output

```

d = 101
e = 10100

```

Fun Fact about the C language assignment operator `=`: Did you know that an assignment in C can be used in an expression? Let's look at an example:

[C]

Listing 13: main.c

```

1  #include <stdio.h>
2
3  int main(int argc, const char * argv[])
4  {
5      int a = 0;
6
7      if (a = 10)
8          printf("True\n");
9      else
10         printf("False\n");
11
12     return 0;
13 }

```

Runtime output

```

True

```

Run the above code example. What does it output? Is that what you were expecting?

The author of the above code example probably meant to test if `a == 10` in the if statement but accidentally typed `=` instead of `==`. Because C treats assignment as an expression, it was able to evaluate `a = 10`.

Let's look at the equivalent Ada code:

[Ada]

Listing 14: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main

```

(continues on next page)

(continued from previous page)

```

4 is
5   A : Integer := 0;
6 begin
7
8   if A := 10 then
9     Put_Line ("True");
10  else
11    Put_Line ("False");
12  end if;
13 end Main;

```

The above code will not compile. This is because Ada does not allow assignment as an expression.

The "use" clause

You'll notice in the above code example, after `with Ada.Text_IO;` there is a new statement we haven't seen before — `use Ada.Text_IO;`. You may also notice that we are not using the `Ada.Text_IO` prefix before the `Put_Line` statements. When we add the `use` clause it tells the compiler that we won't be using the prefix in the call to subprograms of that package. The `use` clause is something to use with caution. For example: if we use the `Ada.Text_IO` package and we also have a `Put_Line` subprogram in our current compilation unit with the same signature, we have a (potential) collision!

2.9 Conditions

The syntax of an `if` statement:

[C]

Listing 15: main.c

```

1  #include <stdio.h>
2
3  int main(int argc, const char * argv[])
4  {
5      // try changing the initial value to change the
6      // output of the program
7      int v = 0;
8
9      if (v > 0) {
10         printf("Positive\n");
11     }
12     else if (v < 0) {
13         printf("Negative\n");
14     }
15     else {
16         printf("Zero\n");
17     }
18
19     return 0;
20 }

```

Runtime output

Zero

[Ada]

Listing 16: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main
4 is
5   -- try changing the initial value to change the
6   -- output of the program
7   V : constant Integer := 0;
8 begin
9   if V > 0 then
10    Put_Line ("Positive");
11   elsif V < 0 then
12    Put_Line ("Negative");
13   else
14    Put_Line ("Zero");
15   end if;
16 end Main;
```

Build output

```
Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```
Zero
```

In Ada, everything that appears between the `if` and `then` keywords is the conditional expression, no parentheses are required. Comparison operators are the same except for:

Operator	C	Ada
Equality	<code>==</code>	<code>=</code>
Inequality	<code>!=</code>	<code>/=</code>
Not	<code>!</code>	<code>not</code>
And	<code>&&</code>	<code>and</code>
Or	<code> </code>	<code>or</code>

The syntax of a switch/case statement:

[C]

Listing 17: main.c

```
1 #include <stdio.h>
2
3 int main(int argc, const char * argv[])
4 {
5   // try changing the initial value to change the
6   // output of the program
7   int v = 0;
8
9   switch(v) {
10     case 0:
11       printf("Zero\n");
```

(continues on next page)

(continued from previous page)

```

12         break;
13     case 1: case 2: case 3: case 4: case 5:
14     case 6: case 7: case 8: case 9:
15         printf("Positive\n");
16         break;
17     case 10: case 12: case 14: case 16: case 18:
18         printf("Even number between 10 and 18\n");
19         break;
20     default:
21         printf("Something else\n");
22         break;
23 }
24
25 return 0;
26 }
```

Runtime output

Zero

[Ada]

Listing 18: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main
4  is
5      -- try changing the initial value to change the
6      -- output of the program
7      V : constant Integer := 0;
8  begin
9      case V is
10         when 0 =>
11             Put_Line ("Zero");
12         when 1 .. 9 =>
13             Put_Line ("Positive");
14         when 10 | 12 | 14 | 16 | 18 =>
15             Put_Line ("Even number between 10 and 18");
16         when others =>
17             Put_Line ("Something else");
18     end case;
19 end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

Zero

Switch or Case?

A switch statement in C is the same as a case statement in Ada. This may be a little strange because

C uses both keywords in the statement syntax. Let's make an analogy between C and Ada: C's `switch` is to Ada's `case` as C's `case` is to Ada's `when`.

Notice that in Ada, the case statement does not use the `break` keyword. In C, we use `break` to stop the execution of a case branch from falling through to the next branch. Here is an example:

[C]

Listing 19: main.c

```
1 #include <stdio.h>
2
3 int main(int argc, const char * argv[])
4 {
5     int v = 0;
6
7     switch(v) {
8         case 0:
9             printf("Zero\n");
10        case 1:
11            printf("One\n");
12        default:
13            printf("Other\n");
14    }
15
16    return 0;
17 }
```

Runtime output

```
Zero
One
Other
```

Run the above code with `v = 0`. What prints? What prints when we change the assignment to `v = 1`?

When `v = 0` the program outputs the strings `Zero` then `One` then `Other`. This is called fall through. If you add the `break` statements back into the `switch` you can stop this fall through behavior from happening. The reason why fall through is allowed in C is to allow the behavior from the previous example where we want a specific branch to execute for multiple inputs. Ada solves this a different way because it is possible, or even probable, that the developer might forget a `break` statement accidentally. So Ada does not allow fall through. Instead, you can use Ada's syntax to identify when a specific branch can be executed by more than one input. If you want a range of values for a specific branch you can use the `First .. Last` notation. If you want a few non-consecutive values you can use the `Value1 | Value2 | Value3` notation.

Instead of using the word `default` to denote the catch-all case, Ada uses the `others` keyword.

2.10 Loops

Let's start with some syntax:

[C]

Listing 20: main.c

```
1 #include <stdio.h>
2
```

(continues on next page)

(continued from previous page)

```

3  int main(int argc, const char * argv[])
4  {
5      int v;
6
7      // this is a while loop
8      v = 1;
9      while(v < 100) {
10         v *= 2;
11     }
12     printf("v = %d\n", v);
13
14     // this is a do while loop
15     v = 1;
16     do {
17         v *= 2;
18     } while(v < 200);
19     printf("v = %d\n", v);
20
21     // this is a for loop
22     v = 0;
23     for(int i = 0; i < 5; ++i) {
24         v += (i * i);
25     }
26     printf("v = %d\n", v);
27
28     // this is a forever loop with a conditional exit
29     v = 0;
30     while(1) {
31         // do stuff here
32         v += 1;
33         if(v == 10)
34             break;
35     }
36     printf("v = %d\n", v);
37
38     // this is a loop over an array
39     {
40         #define ARR_SIZE (10)
41         const int arr[ARR_SIZE] = { 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 };
42         int sum = 0;
43
44         for(int i = 0; i < ARR_SIZE; ++i) {
45             sum += arr[i];
46         }
47         printf("sum = %d\n", sum);
48     }
49
50     return 0;
51 }

```

Runtime output

```

v = 128
v = 256
v = 30
v = 10
sum = 55

```

[Ada]

Listing 21: main.adb

```

1  with Ada.Text_IO;
2
3  procedure Main is
4      V : Integer;
5  begin
6      -- this is a while loop
7      V := 1;
8      while V < 100 loop
9          V := V * 2;
10     end loop;
11     Ada.Text_IO.Put_Line ("V = " & Integer'Image (V));
12
13     -- Ada doesn't have an explicit do while loop
14     -- instead you can use the loop and exit keywords
15     V := 1;
16     loop
17         V := V * 2;
18         exit when V >= 200;
19     end loop;
20     Ada.Text_IO.Put_Line ("V = " & Integer'Image (V));
21
22     -- this is a for loop
23     V := 0;
24     for I in 0 .. 4 loop
25         V := V + (I * I);
26     end loop;
27     Ada.Text_IO.Put_Line ("V = " & Integer'Image (V));
28
29     -- this is a forever loop with a conditional exit
30     V := 0;
31     loop
32         -- do stuff here
33         V := V + 1;
34         exit when V = 10;
35     end loop;
36     Ada.Text_IO.Put_Line ("V = " & Integer'Image (V));
37
38     -- this is a loop over an array
39     declare
40         type Int_Array is array (Natural range 1 .. 10) of Integer;
41
42         Arr : constant Int_Array := (1, 2, 3, 4, 5, 6, 7, 8, 9, 10);
43         Sum : Integer := 0;
44     begin
45         for I in Arr'Range loop
46             Sum := Sum + Arr (I);
47         end loop;
48         Ada.Text_IO.Put_Line ("Sum = " & Integer'Image (Sum));
49     end;
50 end Main;

```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
V = 128
V = 256
V = 30
V = 10
Sum = 55
```

The loop syntax in Ada is pretty straightforward. The `loop` and `end loop` keywords are used to open and close the loop scope. Instead of using the `break` keyword to exit the loop, Ada has the `exit` statement. The `exit` statement can be combined with a logic expression using the `exit when` syntax.

The major deviation in loop syntax is regarding for loops. You'll notice, in C, that you sometimes declare, and at least initialize a loop counter variable, specify a loop predicate, or an expression that indicates when the loop should continue executing or complete, and last you specify an expression to update the loop counter.

[C]

```
for (initialization expression; loop predicate; update expression) {
    // some statements
}
```

In Ada, you don't declare or initialize a loop counter or specify an update expression. You only name the loop counter and give it a range to loop over. The loop counter is **read-only**! You cannot modify the loop counter inside the loop like you can in C. And the loop counter will increment consecutively along the specified range. But what if you want to loop over the range in reverse order?

[C]

Listing 22: main.c

```
1  #include <stdio.h>
2
3  #define MY_RANGE (10)
4
5  int main(int argc, const char * argv[])
6  {
7
8      for (int i = MY_RANGE; i >= 0; --i) {
9          printf("%d\n", i);
10     }
11
12     return 0;
13 }
```

Runtime output

```
10
9
8
7
6
5
4
3
2
1
0
```

[Ada]

Listing 23: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main
4 is
5   My_Range : constant := 10;
6 begin
7   for I in reverse 0 .. My_Range loop
8     Put_Line (I'Img);
9   end loop;
10 end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
10
9
8
7
6
5
4
3
2
1
0
```

Tick Image

Strangely enough, Ada people call the single apostrophe symbol, `'`, "tick". This "tick" says the we are accessing an attribute of the variable. When we do `'Img` on a variable of a numerical type, we are going to return the string version of that numerical type. So in the for loop above, `I'Img`, or "I tick image" will return the string representation of the numerical value stored in `I`. We have to do this because `Put_Line` is expecting a string as an input parameter.

We'll discuss attributes in more details [later in this chapter](#) (page 41).

In the above example, we are traversing over the range in reverse order. In Ada, we use the `reverse` keyword to accomplish this.

In many cases, when we are writing a for loop, it has something to do with traversing an array. In C, this is a classic location for off-by-one errors. Let's see an example in action:

[C]

Listing 24: main.c

```
1 #include <stdio.h>
2
3 #define LIST_LENGTH (100)
4
```

(continues on next page)

(continued from previous page)

```

5  int main(int argc, const char * argv[])
6  {
7      int list[LIST_LENGTH];
8
9      for(int i = LIST_LENGTH; i > 0; --i) {
10         list[i] = LIST_LENGTH - i;
11     }
12
13     for (int i = 0; i < LIST_LENGTH; ++i)
14     {
15         printf("%d ", list[i]);
16     }
17     printf("\n");
18
19     return 0;
20 }

```

Runtime output

```

2 99 98 97 96 95 94 93 92 91 90 89 88 87 86 85 84 83 82 81 80 79 78 77 76 75 74 73
↪ 72 71 70 69 68 67 66 65 64 63 62 61 60 59 58 57 56 55 54 53 52 51 50 49 48 47 46
↪ 45 44 43 42 41 40 39 38 37 36 35 34 33 32 31 30 29 28 27 26 25 24 23 22 21 20 19
↪ 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

```

[Ada]

Listing 25: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main
4  is
5      type Int_Array is array (Natural range 1 .. 100) of Integer;
6
7      List : Int_Array;
8  begin
9
10     for I in reverse List'Range loop
11         List (I) := List'Last - I;
12     end loop;
13
14     for I in List'Range loop
15         Put (List (I)'Img & " ");
16     end loop;
17
18     New_Line;
19 end Main;

```

Build output

```

Compile
[Ada]          main.adb
Bind
[gnatbind]     main.bexch
[Ada]          main.ali
Link
[link]         main.adb

```

Runtime output

```

99 98 97 96 95 94 93 92 91 90 89 88 87 86 85 84 83 82 81 80
↪ 79 78 77 76 75 74 73 72 71 70 69 68 67 66 65 64 63 62 61 60
↪ 59 58 57 56 55 54 53 52 51 50 49 48 47 46 45 44 43 42 41 40
↪ 39 38 37 36 35 34 33 32 31 30 29 28 27 26 25 24 23 22 21 20
↪ 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0

```

(continued from previous page)

The above Ada and C code should initialize an array using a for loop. The initial values in the array should be contiguously decreasing from 99 to 0 as we index from the first index to the last index. In other words, the first index has a value of 99, the next has 98, the next 97 ... the last has a value of 0.

If you run both the C and Ada code above you'll notice that the outputs of the two programs are different. Can you spot why?

In the C code there are two problems:

1. There's a buffer overflow in the first iteration of the loop. We would need to modify the loop initialization to `int i = LIST_LENGTH - 1;`. The loop predicate should be modified to `i >= 0;`
2. The C code also has another off-by-one problem in the math to compute the value stored in `list[i]`. The expression should be changed to be `list[i] = LIST_LENGTH - i - 1;`

These are typical off-by-one problems that plagues C programs. You'll notice that we didn't have this problem with the Ada code because we aren't defining the loop with arbitrary numeric literals. Instead we are accessing attributes of the array we want to manipulate and are using a keyword to determine the indexing direction.

We can actually simplify the Ada for loop a little further using iterators:

[Ada]

Listing 26: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main
4 is
5     type Int_Array is array (Natural range 1 .. 100) of Integer;
6
7     List : Int_Array;
8 begin
9
10    for I in reverse List'Range loop
11        List (I) := List'Last - I;
12    end loop;
13
14    for I of List loop
15        Put (I'Img & " ");
16    end loop;
17
18    New_Line;
19 end Main;
```

Build output

```

Compile
[Ada]      main.adb
Bind
[gnatbind] main.bexch
[Ada]      main.ali
Link
[link]     main.adb
```

Runtime output

99	98	97	96	95	94	93	92	91	90	89	88	87	86	85	84	83	82	81	80	79	78	77	76	75	74	73	72	71	70	69	68	67	66	65	64	63	62	61	60	59	58	57	56	55	54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	---	---	---	---	---	---	---	---	---	---

Chapter 2. The C Developer's Perspective on Ada

(continued from previous page)

In the second for loop, we changed the syntax to `for I of List`. Instead of `I` being the index counter, it is now an iterator that references the underlying element. This example of Ada code is identical to the last bit of Ada code. We just used a different method to index over the second for loop. There is no C equivalent to this Ada feature, but it is similar to C++'s range based for loop.

2.11 Type System

2.11.1 Strong Typing

Ada is considered a "strongly typed" language. This means that the language does not define any implicit type conversions. C does define implicit type conversions, sometimes referred to as *integer promotion*. The rules for promotion are fairly straightforward in simple expressions but can get confusing very quickly. Let's look at a typical place of confusion with implicit type conversion:

[C]

Listing 27: main.c

```

1  #include <stdio.h>
2
3  int main(int argc, const char * argv[])
4  {
5      unsigned char a = 0xFF;
6      char b = 0xFF;
7
8      printf("Does a == b?\n");
9      if(a == b)
10         printf("Yes.\n");
11     else
12         printf("No.\n");
13
14     printf("a: 0x%08X, b: 0x%08X\n", a, b);
15
16     return 0;
17 }
```

Runtime output

```

Does a == b?
No.
a: 0x000000FF, b: 0xFFFFFFFF
```

Run the above code. You will notice that `a != b`! If we look at the output of the last `printf` statement we will see the problem. `a` is an unsigned number where `b` is a signed number. We stored a value of `0xFF` in both variables, but `a` treated this as the decimal number 255 while `b` treated this as the decimal number -1. When we compare the two variables, of course they aren't equal; but that's not very intuitive. Let's look at the equivalent Ada example:

[Ada]

Listing 28: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main
4  is
```

(continues on next page)

(continued from previous page)

```

5  type Char is range 0 .. 255;
6  type Unsigned_Char is mod 256;
7
8  A : Char := 16#FF#;
9  B : Unsigned_Char := 16#FF#;
10 begin
11
12   Put_Line ("Does A = B?");
13
14   if A = B then
15     Put_Line ("Yes");
16   else
17     Put_Line ("No");
18   end if;
19
20 end Main;

```

Compilation output

```

main.adb:14:09: error: invalid operand types for operator "="
main.adb:14:09: error: left operand has type "Char" defined at line 5
main.adb:14:09: error: right operand has type "Unsigned_Char" defined at line 6

```

If you try to run this Ada example you will get a compilation error. This is because the compiler is telling you that you cannot compare variables of two different types. We would need to explicitly cast one side to make the comparison against two variables of the same type. By enforcing the explicit cast we can't accidentally end up in a situation where we assume something will happen implicitly when, in fact, our assumption is incorrect.

Another example: you can't divide an integer by a float. You need to perform the division operation using values of the same type, so one value must be explicitly converted to match the type of the other (in this case the more likely conversion is from integer to float). Ada is designed to guarantee that what's done by the program is what's meant by the programmer, leaving as little room for compiler interpretation as possible. Let's have a look at the following example:

[Ada]

Listing 29: strong_typing.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Strong_Typing is
4     Alpha : constant Integer := 1;
5     Beta  : constant Integer := 10;
6     Result : Float;
7  begin
8     Result := Float (Alpha) / Float (Beta);
9
10    Put_Line (Float'Image (Result));
11 end Strong_Typing;

```

Build output

Compile	
[Ada]	strong_typing.adb
Bind	
[gprbind]	strong_typing.bexch
[Ada]	strong_typing.ali
Link	
[link]	strong_typing.adb

Runtime output

```
1.00000E-01
```

[C]

Listing 30: main.c

```

1  #include <stdio.h>
2
3  void weakTyping (void) {
4      const int    alpha = 1;
5      const int    beta  = 10;
6      float result;
7
8      result = alpha / beta;
9
10     printf("%f\n", result);
11 }
12
13 int main(int argc, const char * argv[])
14 {
15     weakTyping();
16
17     return 0;
18 }

```

Runtime output

```
0.000000
```

Are the three programs above equivalent? It may seem like Ada is just adding extra complexity by forcing you to make the conversion from Integer to Float explicit. In fact, it significantly changes the behavior of the computation. While the Ada code performs a floating point operation $1.0/10.0$ and stores 0.1 in Result, the C version instead store 0.0 in result. This is because the C version perform an integer operation between two integer variables: $1/10$ is 0. The result of the integer division is then converted to a float and stored. Errors of this sort can be very hard to locate in complex pieces of code, and systematic specification of how the operation should be interpreted helps to avoid this class of errors. If an integer division was actually intended in the Ada case, it is still necessary to explicitly convert the final result to Float:

[Ada]

```

-- Perform an Integer division then convert to Float
Result := Float (Alpha / Beta);

```

The complete example would then be:

[Ada]

Listing 31: strong_typing.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Strong_Typing is
4      Alpha : constant Integer := 1;
5      Beta  : constant Integer := 10;
6      Result : Float;
7  begin
8      Result := Float (Alpha / Beta);
9
10     Put_Line (Float'Image (Result));
11 end Strong_Typing;

```

Build output

```
Compile
  [Ada]          strong_typing.adb
Bind
  [gprbind]      strong_typing.bexch
  [Ada]          strong_typing.ali
Link
  [link]         strong_typing.adb
```

Runtime output

```
0.000000E+00
```

Floating Point Literals

In Ada, a floating point literal must be written with both an integral and decimal part. `10` is not a valid literal for a floating point value, while `10.0` is.

2.11.2 Language-Defined Types

The principal scalar types predefined by Ada are Integer, Float, Boolean, and Character. These correspond to `int`, `float`, `int` (when used for Booleans), and `char`, respectively. The names for these types are not reserved words; they are regular identifiers. There are other language-defined integer and floating-point types as well. All have implementation-defined ranges and precision.

2.11.3 Application-Defined Types

Ada's type system encourages programmers to think about data at a high level of abstraction. The compiler will at times output a simple efficient machine instruction for a full line of source code (and some instructions can be eliminated entirely). The careful programmer's concern that the operation really makes sense in the real world would be satisfied, and so would the programmer's concern about performance.

The next example below defines two different metrics: area and distance. Mixing these two metrics must be done with great care, as certain operations do not make sense, like adding an area to a distance. Others require knowledge of the expected semantics; for example, multiplying two distances. To help avoid errors, Ada requires that each of the binary operators `+`, `-`, `*`, and `/` for integer and floating-point types take operands of the same type and return a value of that type.

[Ada]

Listing 32: main.adb

```
1  procedure Main is
2    type Distance is new Float;
3    type Area is new Float;
4
5    D1 : Distance := 2.0;
6    D2 : Distance := 3.0;
7    A  : Area;
8  begin
9    D1 := D1 + D2; -- OK
10   D1 := D1 + A;  -- NOT OK: incompatible types for "+"
11   A  := D1 * D2; -- NOT OK: incompatible types for ":="
12   A  := Area (D1 * D2); -- OK
13 end Main;
```

Compilation output

```
main.adb:10:13: error: invalid operand types for operator "+"
main.adb:10:13: error: left operand has type "Distance" defined at line 2
main.adb:10:13: error: right operand has type "Area" defined at line 3
main.adb:11:13: error: expected type "Area" defined at line 3
main.adb:11:13: error: found type "Distance" defined at line 2
```

Even though the `Distance` and `Area` types above are just `Float`, the compiler does not allow arbitrary mixing of values of these different types. An explicit conversion (which does not necessarily mean any additional object code) is necessary.

The predefined Ada rules are not perfect; they admit some problematic cases (for example multiplying two `Distance` yields a `Distance`) and prohibit some useful cases (for example multiplying two `Distances` should deliver an `Area`). These situations can be handled through other mechanisms. A predefined operation can be identified as abstract to make it unavailable; overloading can be used to give new interpretations to existing operator symbols, for example allowing an operator to return a value from a type different from its operands; and more generally, GNAT has introduced a facility that helps perform dimensionality checking.

Ada enumerations work similarly to C `enum`:

[Ada]

Listing 33: main.adb

```
1 procedure Main is
2   type Day is
3     (Monday,
4      Tuesday,
5      Wednesday,
6      Thursday,
7      Friday,
8      Saturday,
9      Sunday);
10
11   D : Day := Monday;
12 begin
13   null;
14 end Main;
```

Build output

```
Compile
[Ada]      main.adb
Bind
[gnatbind] main.bexch
[Ada]      main.ali
Link
[link]     main.adb
```

[C]

Listing 34: main.c

```
1 enum Day {
2   Monday,
3   Tuesday,
4   Wednesday,
5   Thursday,
6   Friday,
7   Saturday,
8   Sunday
```

(continues on next page)

(continued from previous page)

```
9 };
10
11 int main(int argc, const char * argv[])
12 {
13     enum Day d = Monday;
14
15     return 0;
16 }
```

But even though such enumerations may be implemented by the compiler as numeric values, at the language level Ada will not confuse the fact that `Monday` is a `Day` and is not an `Integer`. You can compare a `Day` with another `Day`, though. To specify implementation details like the numeric values that correspond with enumeration values in C you include them in the original enum declaration:

[C]

Listing 35: main.c

```
1 #include <stdio.h>
2
3 enum Day {
4     Monday    = 10,
5     Tuesday   = 11,
6     Wednesday = 12,
7     Thursday  = 13,
8     Friday    = 14,
9     Saturday  = 15,
10    Sunday     = 16
11 };
12
13 int main(int argc, const char * argv[])
14 {
15     enum Day d = Monday;
16
17     printf("d = %d\n", d);
18
19     return 0;
20 }
```

Runtime output

```
d = 10
```

But in Ada you must use both a type definition for `Day` as well as a separate representation clause for it like:

[Ada]

Listing 36: main.adb

```
1 with Ada.Text_IO;
2
3 procedure Main is
4     type Day is
5         (Monday,
6          Tuesday,
7          Wednesday,
8          Thursday,
9          Friday,
10         Saturday,
```

(continues on next page)

(continued from previous page)

```

11     Sunday);
12
13     -- Representation clause for Day type:
14     for Day use
15         (Monday    => 10,
16          Tuesday   => 11,
17          Wednesday => 12,
18          Thursday  => 13,
19          Friday    => 14,
20          Saturday  => 15,
21          Sunday    => 16);
22
23     D : Day := Monday;
24     V : Integer;
25 begin
26     V := Day'Enum_Rep (D);
27     Ada.Text_IO.Put_Line (Integer'Image (V));
28 end Main;

```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
10
```

Note that however, unlike C, values for enumerations in Ada have to be unique.

2.11.4 Type Ranges

Contracts can be associated with types and variables, to refine values and define what are considered valid values. The most common kind of contract is a *range constraint* introduced with the range reserved word, for example:

[Ada]

Listing 37: main.adb

```

1  procedure Main is
2      type Grade is range 0 .. 100;
3
4      G1, G2 : Grade;
5      N      : Integer;
6  begin
7      -- ...           -- Initialization of N
8      G1 := 80;         -- OK
9      G1 := N;          -- Illegal (type mismatch)
10     G1 := Grade (N);   -- Legal, run-time range check
11     G2 := G1 + 10;     -- Legal, run-time range check
12     G1 := (G1 + G2) / 2; -- Legal, run-time range check
13 end Main;

```

Compilation output

```
main.adb:9:10: error: expected type "Grade" defined at line 2
main.adb:9:10: error: found type "Standard.Integer"
```

In the above example, `Grade` is a new integer type associated with a range check. Range checks are dynamic and are meant to enforce the property that no object of the given type can have a value outside the specified range. In this example, the first assignment to `G1` is correct and will not raise a run-time exception. Assigning `N` to `G1` is illegal since `Grade` is a different type than `Integer`. Converting `N` to `Grade` makes the assignment legal, and a range check on the conversion confirms that the value is within `0 .. 100`. Assigning `G1 + 10` to `G2` is legal since `+` for `Grade` returns a `Grade` (note that the literal `10` is interpreted as a `Grade` value in this context), and again there is a range check.

The final assignment illustrates an interesting but subtle point. The subexpression `G1 + G2` may be outside the range of `Grade`, but the final result will be in range. Nevertheless, depending on the representation chosen for `Grade`, the addition may overflow. If the compiler represents `Grade` values as signed 8-bit integers (i.e., machine numbers in the range `-128 .. 127`) then the sum `G1 + G2` may exceed `127`, resulting in an integer overflow. To prevent this, you can use explicit conversions and perform the computation in a sufficiently large integer type, for example:

[Ada]

Listing 38: main.adb

```
1 with Ada.Text_IO;
2
3 procedure Main is
4   type Grade is range 0 .. 100;
5
6   G1, G2 : Grade := 99;
7 begin
8   G1 := Grade ((Integer (G1) + Integer (G2)) / 2);
9   Ada.Text_IO.Put_Line (Grade'Image (G1));
10 end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

99

Range checks are useful for detecting errors as early as possible. However, there may be some impact on performance. Modern compilers do know how to remove redundant checks, and you can deactivate these checks altogether if you have sufficient confidence that your code will function correctly.

Types can be derived from the representation of any other type. The new derived type can be associated with new constraints and operations. Going back to the `Day` example, one can write:

[Ada]

Listing 39: main.adb

```
1 procedure Main is
2   type Day is
```

(continues on next page)

(continued from previous page)

```

3      (Monday,
4        Tuesday,
5        Wednesday,
6        Thursday,
7        Friday,
8        Saturday,
9        Sunday);
10
11     type Business_Day is new Day range Monday .. Friday;
12     type Weekend_Day is new Day range Saturday .. Sunday;
13 begin
14     null;
15 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Since these are new types, implicit conversions are not allowed. In this case, it's more natural to create a new set of constraints for the same type, instead of making completely new ones. This is the idea behind *subtypes* in Ada. A subtype is a type with optional additional constraints. For example:

[Ada]

Listing 40: main.adb

```

1  procedure Main is
2      type Day is
3          (Monday,
4            Tuesday,
5            Wednesday,
6            Thursday,
7            Friday,
8            Saturday,
9            Sunday);
10
11     subtype Business_Day is Day range Monday .. Friday;
12     subtype Weekend_Day is Day range Saturday .. Sunday;
13     subtype Dice_Throw is Integer range 1 .. 6;
14 begin
15     null;
16 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

These declarations don't create new types, just new names for constrained ranges of their base types.

The purpose of numeric ranges is to express some application-specific constraint that we want the compiler to help us enforce. More importantly, we want the compiler to tell us when that constraint cannot be met — when the underlying hardware cannot support the range given. There are two things to consider:

- just a range constraint, such as `A : Integer range 0 .. 10;`, or
- a type declaration, such as `type Result is range 0 .. 1_000_000_000;`.

Both represent some sort of application-specific constraint, but in addition, the type declaration promotes portability because it won't compile on targets that do not have a sufficiently large hardware numeric type. That's a definition of portability that is preferable to having something compile anywhere but not run correctly, as in C.

2.11.5 Unsigned And Modular Types

Unsigned integer numbers are quite common in embedded applications. In C, you can use them by declaring unsigned `int` variables. In Ada, you have two options:

- declare custom *unsigned* range types;
 - In addition, you can declare custom range *subtypes* or use existing subtypes such as `Natural`.
- declare custom modular types.

The following table presents the main features of each type. We discuss these types right after.

Feature	[C] unsigned int	[Ada] Unsigned range	[Ada] Modular
Excludes negative value	✓	✓	✓
Wraparound	✓		✓

When declaring custom range types in Ada, you may use the full range in the same way as in C. For example, this is the declaration of a 32-bit unsigned integer type and the X variable in Ada:

[Ada]

Listing 41: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   type Unsigned_Int_32 is range 0 .. 2 ** 32 - 1;
5
6   X : Unsigned_Int_32 := 42;
7 begin
8   Put_Line ("X = " & Unsigned_Int_32'Image (X));
9 end Main;
```

Build output

```

Compile
  [Ada]      main.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb
```

Runtime output

```
X = 42
```

In C, when `unsigned int` has a size of 32 bits, this corresponds to the following declaration:
[C]

Listing 42: main.c

```

1  #include <stdio.h>
2  #include <limits.h>
3
4  int main(int argc, const char * argv[])
5  {
6      unsigned int x = 42;
7      printf("x = %u\n", x);
8
9      return 0;
10 }
```

Runtime output

```
x = 42
```

Another strategy is to declare subtypes for existing signed types and specify just the range that excludes negative numbers. For example, let's declare a custom 32-bit signed type and its unsigned subtype:

[Ada]

Listing 43: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4      type Signed_Int_32 is range -2 ** 31 .. 2 ** 31 - 1;
5
6      subtype Unsigned_Int_31 is Signed_Int_32 range 0 .. Signed_Int_32'Last;
7      -- Equivalent to:
8      -- subtype Unsigned_Int_31 is Signed_Int_32 range 0 .. 2 ** 31 - 1;
9
10     X : Unsigned_Int_31 := 42;
11 begin
12     Put_Line ("X = " & Unsigned_Int_31'Image (X));
13 end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
X = 42
```

In this case, we're just skipping the sign bit of the `Signed_Int_32` type. In other words, while `Signed_Int_32` has a size of 32 bits, `Unsigned_Int_31` has a range of 31 bits, even if the base type has 32 bits.

Note that the declaration above is actually similar to the existing `Natural` subtype. Ada provides the following standard subtypes:

```
subtype Natural is Integer range 0..Integer'Last;
subtype Positive is Integer range 1..Integer'Last;
```

Since they're standard subtypes, you can declare variables of those subtypes directly in your implementation, in the same way as you can declare `Integer` variables.

As indicated in the table above, however, there is a difference in behavior for the variables we just declared, which occurs in case of overflow. Let's consider this C example:

[C]

Listing 44: main.c

```
1 #include <stdio.h>
2 #include <limits.h>
3
4 int main(int argc, const char * argv[])
5 {
6     unsigned int x = UINT_MAX + 1;
7     /* Now: x == 0 */
8
9     printf("x = %u\n", x);
10
11     return 0;
12 }
```

Runtime output

```
x = 0
```

The corresponding code in Ada raises an exception:

[Ada]

Listing 45: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4     type Unsigned_Int_32 is range 0 .. 2 ** 32 - 1;
5
6     X : Unsigned_Int_32 := Unsigned_Int_32'Last + 1;
7     -- Overflow: exception is raised!
8 begin
9     Put_Line ("X = " & Unsigned_Int_32'Image (X));
10 end Main;
```

Build output

```
Compile
  [Ada]          main.adb
main.adb:6:48: warning: value not in range of type "Unsigned_Int_32" defined at
↳ line 4 [enabled by default]
main.adb:6:48: warning: "Constraint_Error" will be raised at run time [enabled by
↳ default]
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```
raised CONSTRAINT_ERROR : main.adb:6 range check failed
```

While the C uses modulo arithmetic for unsigned integer, Ada doesn't use it for the `Unsigned_Int_32` type. Ada does, however, support modular types via type definitions using the `mod` keyword. In this example, we declare a 32-bit modular type:

[Ada]

Listing 46: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   type Unsigned_32 is mod 2**32;
5
6   X : Unsigned_32 := Unsigned_32'Last + 1;
7   -- Now: X = 0
8 begin
9   Put_Line ("X = " & Unsigned_32'Image (X));
10 end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
X = 0
```

In this case, the behavior is the same as in the C declaration above.

Modular types, unlike Ada's signed integers, also provide bit-wise operations, a typical application for unsigned integers in C. In Ada, you can use operators such as `and`, `or`, `xor` and `not`. You can also use typical bit-shifting operations, such as `Shift_Left`, `Shift_Right`, `Shift_Right_Arithmetic`, `Rotate_Left` and `Rotate_Right`.

2.11.6 Attributes

Attributes start with a single apostrophe ("tick"), and they allow you to query properties of, and perform certain actions on, declared entities such as types, objects, and subprograms. For example, you can determine the first and last bounds of scalar types, get the sizes of objects and types, and convert values to and from strings. This section provides an overview of how attributes work. For more information on the many attributes defined by the language, you can refer directly to the Ada Language Reference Manual.

The `'Image` and `'Value` attributes allow you to transform a scalar value into a `String` and vice-versa. For example:

[Ada]

Listing 47: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   A : Integer := 10;
5 begin
6   Put_Line (Integer'Image (A));
7   A := Integer'Value ("99");
8   Put_Line (Integer'Image (A));
9 end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
10
99
```

Important

Semantically, attributes are equivalent to subprograms. For example, `Integer'Image` is defined as follows:

```
function Integer'Image(Arg : Integer'Base) return String;
```

Certain attributes are provided only for certain kinds of types. For example, the `'Val` and `'Pos` attributes for an enumeration type associates a discrete value with its position among its peers. One circuitous way of moving to the next character of the ASCII table is:

[Ada]

Listing 48: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   C : Character := 'a';
5 begin
6   Put (C);
7   C := Character'Val (Character'Pos (C) + 1);
8   Put (C);
9 end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali

(continues on next page)

(continued from previous page)

```
Link
  [link]      main.adb
```

Runtime output

```
ab
```

A more concise way to get the next value in Ada is to use the 'Succ attribute:

[Ada]

Listing 49: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   C : Character := 'a';
5 begin
6   Put (C);
7   C := Character'Succ (C);
8   Put (C);
9 end Main;
```

Build output

```
Compile
  [Ada]      main.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]      main.adb
```

Runtime output

```
ab
```

You can get the previous value using the 'Pred attribute. Here is the equivalent in C:

[C]

Listing 50: main.c

```
1 #include <stdio.h>
2
3 int main(int argc, const char * argv[])
4 {
5     char c = 'a';
6     printf("%c", c);
7     c++;
8     printf("%c", c);
9
10    return 0;
11 }
```

Runtime output

```
ab
```

Other interesting examples are the 'First and 'Last attributes which, respectively, return the first and last values of a scalar type. Using 32-bit integers, for instance, Integer'First returns -2^{31} and Integer'Last returns $2^{31} - 1$.

2.11.7 Arrays and Strings

C arrays are pointers with offsets, but the same is not the case for Ada. Arrays in Ada are not interchangeable with operations on pointers, and array types are considered first-class citizens. They have dedicated semantics such as the availability of the array's boundaries at run-time. Therefore, unhandled array overflows are impossible unless checks are suppressed. Any discrete type can serve as an array index, and you can specify both the starting and ending bounds — the lower bound doesn't necessarily have to be 0. Most of the time, array types need to be explicitly declared prior to the declaration of an object of that array type.

Here's an example of declaring an array of 26 characters, initializing the values from 'a' to 'z':

[Ada]

Listing 51: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   type Arr_Type is array (Integer range <>) of Character;
5   Arr : Arr_Type (1 .. 26);
6   C : Character := 'a';
7 begin
8   for I in Arr'Range loop
9     Arr (I) := C;
10    C := Character'Succ (C);
11
12    Put (Arr (I) & " ");
13  end loop;
14 end Main;
```

Build output

```
Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
```

[C]

Listing 52: main.c

```
1 #include <stdio.h>
2
3 int main(int argc, const char * argv[])
4 {
5     char Arr [26];
6     char C = 'a';
7
8     for (int I = 0; I < 26; ++I) {
9         Arr [I] = C++;
10        printf ("%c ", Arr [I]);
11    }
12
13    return 0;
14 }
```

Runtime output

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
```

In C, only the size of the array is given during declaration. In Ada, array index ranges are specified using two values of a discrete type. In this example, the array type declaration specifies the use of Integer as the index type, but does not provide any constraints (use <>, pronounced *box*, to specify "no constraints"). The constraints are defined in the object declaration to be 1 to 26, inclusive. Arrays have an attribute called 'Range'. In our example, Arr'Range can also be expressed as Arr'First .. Arr'Last; both expressions will resolve to 1 .. 26. So the 'Range attribute supplies the bounds for our for loop. There is no risk of stating either of the bounds incorrectly, as one might do in C where `I <= 26` may be specified as the end-of-loop condition.

As in C, Ada String is an array of Character. Ada strings, importantly, are not delimited with the special character '\0' like they are in C. It is not necessary because Ada uses the array's bounds to determine where the string starts and stops.

Ada's predefined String type is very straightforward to use:

[Ada]

Listing 53: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   My_String : String (1 .. 19) := "This is an example!";
5 begin
6   Put_Line (My_String);
7 end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
This is an example!
```

Unlike C, Ada does not offer escape sequences such as '\n'. Instead, explicit values from the ASCII package must be concatenated (via the concatenation operator, &). Here for example, is how to initialize a line of text ending with a new line:

[Ada]

Listing 54: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   My_String : String := "This is a line" & ASCII.LF;
5 begin
6   Put (My_String);
7 end Main;
```

Build output

```
Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```
This is a line
```

You see here that no constraints are necessary for this variable definition. The initial value given allows the automatic determination of `My_String`'s bounds.

Ada offers high-level operations for copying, slicing, and assigning values to arrays. We'll start with assignment. In C, the assignment operator doesn't make a copy of the value of an array, but only copies the address or reference to the target variable. In Ada, the actual array contents are duplicated. To get the above behavior, actual pointer types would have to be defined and used.

[Ada]

Listing 55: main.adb

```
1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4      type Arr_Type is array (Integer range <>) of Integer;
5      A1 : Arr_Type (1 .. 2);
6      A2 : Arr_Type (1 .. 2);
7  begin
8      A1 (1) := 0;
9      A1 (2) := 1;
10
11     A2 := A1;
12
13     for I in A2'Range loop
14         Put_Line (Integer'Image (A2 (I)));
15     end loop;
16 end Main;
```

Build output

```
Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```
0
1
```

[C]

Listing 56: main.c

```

1  #include <stdio.h>
2  #include <string.h>
3
4  int main(int argc, const char * argv[])
5  {
6      int A1 [2];
7      int A2 [2];
8
9      A1 [0] = 0;
10     A1 [1] = 1;
11
12     memcpy (A2, A1, sizeof (int) * 2);
13
14     for (int i = 0; i < 2; i++) {
15         printf("%d\n", A2[i]);
16     }
17
18     return 0;
19 }

```

Runtime output

```

0
1

```

In all of the examples above, the source and destination arrays must have precisely the same number of elements. Ada allows you to easily specify a portion, or slice, of an array. So you can write the following:

[Ada]

Listing 57: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4      type Arr_Type is array (Integer range <>) of Integer;
5      A1 : Arr_Type (1 .. 10) := (1, 2, 3, 4, 5, 6, 7, 8, 9, 10);
6      A2 : Arr_Type (1 .. 5)  := (1, 2, 3, 4, 5);
7  begin
8      A2 (1 .. 3) := A1 (4 .. 6);
9
10     for I in A2'Range loop
11         Put_Line (Integer'Image (A2 (I)));
12     end loop;
13 end Main;

```

Build output

Compile	
[Ada]	main.adb
Bind	
[gnatbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
4
5
6
4
5
```

This assigns the 4th, 5th, and 6th elements of A1 into the 1st, 2nd, and 3rd elements of A2. Note that only the length matters here: the values of the indexes don't have to be equal; they slide automatically.

Ada also offers high level comparison operations which compare the contents of arrays as opposed to their addresses:

[Ada]

Listing 58: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   type Arr_Type is array (Integer range <>) of Integer;
5   A1 : Arr_Type (1 .. 2) := (10, 20);
6   A2 : Arr_Type (1 .. 2) := (10, 20);
7 begin
8   if A1 = A2 then
9     Put_Line ("A1 = A2");
10  else
11    Put_Line ("A1 /= A2");
12  end if;
13 end Main;
```

Build output

```
Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```
A1 = A2
```

[C]

Listing 59: main.c

```
1 #include <stdio.h>
2
3 int main(int argc, const char * argv[])
4 {
5     int A1 [2] = { 10, 20 };
6     int A2 [2] = { 10, 20 };
7
8     int eq = 1;
9
10    for (int i = 0; i < 2; ++i) {
11        if (A1 [i] != A2 [i]) {
12            eq = 0;
13            break;
14        }
15    }
```

(continues on next page)

(continued from previous page)

```

14     }
15 }
16
17 if (eq) {
18     printf("A1 == A2\n");
19 }
20 else {
21     printf("A1 != A2\n");
22 }
23
24 return 0;
25 }

```

Runtime output

```
A1 == A2
```

You can assign to all the elements of an array in each language in different ways. In Ada, the number of elements to assign can be determined by looking at the right-hand side, the left-hand side, or both sides of the assignment. When bounds are known on the left-hand side, it's possible to use the others expression to define a default value for all the unspecified array elements. Therefore, you can write:

[Ada]

Listing 60: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4      type Arr_Type is array (Integer range <>) of Integer;
5      A1 : Arr_Type (-2 .. 42) := (others => 0);
6  begin
7      -- use a slice to assign A1 elements 11 .. 19 to 1
8      A1 (11 .. 19) := (others => 1);
9
10     Put_Line ("---- A1 ----");
11     for I in A1'Range loop
12         Put_Line (Integer'Image (I) & " => " &
13                 Integer'Image (A1 (I)));
14     end loop;
15 end Main;

```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```

---- A1 ----
-2 => 0
-1 => 0
0 => 0
1 => 0
2 => 0

```

(continues on next page)

(continued from previous page)

```

3 => 0
4 => 0
5 => 0
6 => 0
7 => 0
8 => 0
9 => 0
10 => 0
11 => 1
12 => 1
13 => 1
14 => 1
15 => 1
16 => 1
17 => 1
18 => 1
19 => 1
20 => 0
21 => 0
22 => 0
23 => 0
24 => 0
25 => 0
26 => 0
27 => 0
28 => 0
29 => 0
30 => 0
31 => 0
32 => 0
33 => 0
34 => 0
35 => 0
36 => 0
37 => 0
38 => 0
39 => 0
40 => 0
41 => 0
42 => 0

```

In this example, we're specifying that A1 has a range between -2 and 42. We use (others => 0) to initialize all array elements with zero. In the next example, the number of elements is determined by looking at the right-hand side:

[Ada]

Listing 61: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   type Arr_Type is array (Integer range <>) of Integer;
5   A1 : Arr_Type := (1, 2, 3, 4, 5, 6, 7, 8, 9);
6 begin
7   A1 := (1, 2, 3, others => 10);
8
9   Put_Line ("---- A1 ----");
10  for I in A1'Range loop
11    Put_Line (Integer'Image (I) & " => " &
12              Integer'Image (A1 (I)));

```

(continues on next page)

(continued from previous page)

```

13   end loop;
14 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```

---- A1 ----
-2147483648 => 1
-2147483647 => 2
-2147483646 => 3
-2147483645 => 10
-2147483644 => 10
-2147483643 => 10
-2147483642 => 10
-2147483641 => 10
-2147483640 => 10

```

Since A1 is initialized with an aggregate of 9 elements, A1 automatically has 9 elements. Also, we're not specifying any range in the declaration of A1. Therefore, the compiler uses the default range of the underlying array type `Arr_Type`, which has an unconstrained range based on the `Integer` type. The compiler selects the first element of that type (`Integer'First`) as the start index of A1. If you replaced `Integer range <>` in the declaration of the `Arr_Type` by `Positive range <>`, then A1's start index would be `Positive'First` — which corresponds to one.

2.11.8 Heterogeneous Data Structures

The structure corresponding to a C `struct` is an Ada record. Here are some simple records:

[Ada]

Listing 62: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   type R is record
5     A, B : Integer;
6     C   : Float;
7   end record;
8
9   V : R;
10 begin
11   V.A := 0;
12   Put_Line ("V.A = " & Integer'Image (V.A));
13 end Main;

```

Build output

```

Compile
  [Ada]          main.adb

```

(continues on next page)

(continued from previous page)

```

Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```
V.A = 0
```

```
[C]
```

Listing 63: main.c

```

1  #include <stdio.h>
2
3  struct R {
4      int A, B;
5      float C;
6  };
7
8  int main(int argc, const char * argv[])
9  {
10     struct R V;
11     V.A = 0;
12     printf("V.A = %d\n", V.A);
13
14     return 0;
15 }

```

Runtime output

```
V.A = 0
```

Ada allows specification of default values for fields just like C. The values specified can take the form of an ordered list of values, a named list of values, or an incomplete list followed by others => <> to specify that fields not listed will take their default values. For example:

```
[Ada]
```

Listing 64: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4
5      type R is record
6          A, B : Integer := 0;
7          C    : Float   := 0.0;
8      end record;
9
10     procedure Put_R (V : R; Name : String) is
11     begin
12         Put_Line (Name & " = ("
13                 & Integer'Image (V.A) & ", "
14                 & Integer'Image (V.B) & ", "
15                 & Float'Image (V.C) & ")");
16     end Put_R;
17
18     V1 : constant R := (1, 2, 1.0);
19     V2 : constant R := (A => 1, B => 2, C => 1.0);

```

(continues on next page)

(continued from previous page)

```

20  V3 : constant R := (C => 1.0, A => 1, B => 2);
21  V4 : constant R := (C => 1.0, others => <>);
22
23  begin
24    Put_R (V1, "V1");
25    Put_R (V2, "V2");
26    Put_R (V3, "V3");
27    Put_R (V4, "V4");
28  end Main;

```

Build output

```

Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```

V1 = ( 1,  2,  1.00000E+00)
V2 = ( 1,  2,  1.00000E+00)
V3 = ( 1,  2,  1.00000E+00)
V4 = ( 0,  0,  1.00000E+00)

```

2.11.9 Pointers

As a foreword to the topic of pointers, it's important to keep in mind the fact that most situations that would require a pointer in C do not in Ada. In the vast majority of cases, indirect memory management can be hidden from the developer and thus saves from many potential errors. However, there are situations that do require the use of pointers, or said differently that require to make memory indirection explicit. This section will present Ada access types, the equivalent of C pointers. A further section will provide more details as to how situations that require pointers in C can be done without access types in Ada.

We'll continue this section by explaining the difference between objects allocated on the stack and objects allocated on the heap using the following example:

[Ada]

Listing 65: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4    type R is record
5      A, B : Integer;
6    end record;
7
8    procedure Put_R (V : R; Name : String) is
9    begin
10      Put_Line (Name & " = ("
11                & Integer'Image (V.A) & ", "
12                & Integer'Image (V.B) & ")");
13    end Put_R;
14
15    V1, V2 : R;

```

(continues on next page)

(continued from previous page)

```

16
17 begin
18   V1.A := 0;
19   V2 := V1;
20   V2.A := 1;
21
22   Put_R (V1, "V1");
23   Put_R (V2, "V2");
24 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```

V1 = ( 0,  0)
V2 = ( 1,  0)

```

[C]

Listing 66: main.c

```

1  #include <stdio.h>
2
3  struct R {
4      int A, B;
5  };
6
7  void print_r(const struct R *v,
8              const char *name)
9  {
10     printf("%s = (%d, %d)\n", name, v->A, v->B);
11 }
12
13 int main(int argc, const char * argv[])
14 {
15     struct R V1, V2;
16     V1.A = 0;
17     V2 = V1;
18     V2.A = 1;
19
20     print_r(&V1, "V1");
21     print_r(&V2, "V2");
22
23     return 0;
24 }

```

Runtime output

```

V1 = (0, 0)
V2 = (1, 0)

```

There are many commonalities between the Ada and C semantics above. In Ada and C, objects are allocated on the stack and are directly accessed. V1 and V2 are two different objects and the assignment statement copies the value of V1 into V2. V1 and V2 are two distinct objects.

Here's now a similar example, but using heap allocation instead:

[Ada]

Listing 67: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4      type R is record
5          A, B : Integer;
6      end record;
7
8      type R_Access is access R;
9
10     procedure Put_R (V : R; Name : String) is
11     begin
12         Put_Line (Name & " = ("
13                     & Integer'Image (V.A) & ", "
14                     & Integer'Image (V.B) & ")");
15     end Put_R;
16
17     V1 : R_Access;
18     V2 : R_Access;
19 begin
20     V1 := new R;
21     V1.A := 0;
22     V2 := V1;
23     V2.A := 1;
24
25     Put_R (V1.all, "V1");
26     Put_R (V2.all, "V2");
27 end Main;

```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```

V1 = ( 1,  0)
V2 = ( 1,  0)

```

[C]

Listing 68: main.c

```

1  #include <stdio.h>
2  #include <stdlib.h>
3
4  struct R {
5      int A, B;
6  };
7
8  void print_r(const struct R *v,
9              const char *name)
10 {

```

(continues on next page)

(continued from previous page)

```

11     printf("%s = (%d, %d)\n", name, v->A, v->B);
12 }
13
14 int main(int argc, const char * argv[])
15 {
16     struct R * V1, * V2;
17     V1 = malloc(sizeof(struct R));
18     V1->A = 0;
19     V2 = V1;
20     V2->A = 1;
21
22     print_r(V1, "V1");
23     print_r(V2, "V2");
24
25     return 0;
26 }

```

Runtime output

```

V1 = (1, 0)
V2 = (1, 0)

```

In this example, an object of type R is allocated on the heap. The same object is then referred to through V1 and V2. As in C, there's no garbage collector in Ada, so objects allocated by the new operator need to be expressly freed (which is not the case here).

Dereferencing is performed automatically in certain situations, for instance when it is clear that the type required is the dereferenced object rather than the pointer itself, or when accessing record members via a pointer. To explicitly dereference an access variable, append `.all`. The equivalent of `V1->A` in C can be written either as `V1.A` or `V1.all.A`.

Pointers to scalar objects in Ada and C look like:

[Ada]

Listing 69: main.adb

```

1 procedure Main is
2     type A_Int is access Integer;
3     Var : A_Int := new Integer;
4 begin
5     Var.all := 0;
6 end Main;

```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

[C]

Listing 70: main.c

```

1 #include <stdlib.h>
2
3 int main(int argc, const char * argv[])
4 {

```

(continues on next page)

(continued from previous page)

```

5  int * Var = malloc (sizeof(int));
6  *Var = 0;
7  return 0;
8  }

```

In Ada, an initializer can be specified with the allocation by appending ' (value):

[Ada]

Listing 71: main.adb

```

1  procedure Main is
2      type A_Int is access Integer;
3
4      Var : A_Int := new Integer'(0);
5  begin
6      null;
7  end Main;

```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

When using Ada pointers to reference objects on the stack, the referenced objects must be declared as being aliased. This directs the compiler to implement the object using a memory region, rather than using registers or eliminating it entirely via optimization. The access type needs to be declared as either access all (if the referenced object needs to be assigned to) or access constant (if the referenced object is a constant). The 'Access attribute works like the C & operator to get a pointer to the object, but with a *scope accessibility* check to prevent references to objects that have gone out of scope. For example:

[Ada]

Listing 72: main.adb

```

1  procedure Main is
2      type A_Int is access all Integer;
3      Var : aliased Integer;
4      Ptr : A_Int := Var'Access;
5  begin
6      null;
7  end Main;

```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

[C]

Listing 73: main.c

```
1 int main(int argc, const char * argv[])
2 {
3     int Var;
4     int * Ptr = &Var;
5
6     return 0;
7 }
```

To deallocate objects from the heap in Ada, it is necessary to use a deallocation subprogram that accepts a specific access type. A generic procedure is provided that can be customized to fit your needs, it's called `Ada.Unchecked_Deallocation`. To create your customized deallocator (that is, to instantiate this generic), you must provide the object type as well as the access type as follows:

[Ada]

Listing 74: main.adb

```
1 with Ada.Unchecked_Deallocation;
2
3 procedure Main is
4     type Integer_Access is access all Integer;
5     procedure Free is new Ada.Unchecked_Deallocation (Integer, Integer_Access);
6     My_Pointer : Integer_Access := new Integer;
7 begin
8     Free (My_Pointer);
9 end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

[C]

Listing 75: main.c

```

1  #include <stdlib.h>
2
3  int main(int argc, const char * argv[])
4  {
5      int * my_pointer = malloc (sizeof(int));
6      free (my_pointer);
7
8      return 0;
9  }

```

We'll discuss generics later [in this section](#) (page 155).

2.12 Functions and Procedures

2.12.1 General Form

Subroutines in C are always expressed as functions which may or may not return a value. Ada explicitly differentiates between functions and procedures. Functions must return a value and procedures must not. Ada uses the more general term *subprogram* to refer to both functions and procedures.

Parameters can be passed in three distinct modes:

- **in**, which is the default, is for input parameters, whose value is provided by the caller and cannot be changed by the subprogram.
- **out** is for output parameters, with no initial value, to be assigned by the subprogram and returned to the caller.
- **in out** is a parameter with an initial value provided by the caller, which can be modified by the subprogram and returned to the caller (more or less the equivalent of a non-constant pointer in C).

Ada also provides access and aliased parameters, which are in effect explicit pass-by-reference indicators.

In Ada, the programmer specifies how the parameter will be used and in general the compiler decides how it will be passed (i.e., by copy or by reference). C has the programmer specify how to pass the parameter.

Important

There are some exceptions to the "general" rule in Ada. For example, parameters of scalar types are always passed by copy, for all three modes.

Here's a first example:

[Ada]

Listing 76: proc.ads

```

1  procedure Proc
2  (Var1 : Integer;
3   Var2 : out Integer;
4   Var3 : in out Integer);

```

Listing 77: func.ads

```
1 function Func (Var : Integer) return Integer;
```

Listing 78: proc.adb

```
1 with Func;  
2  
3 procedure Proc  
4   (Var1 : Integer;  
5    Var2 : out Integer;  
6    Var3 : in out Integer)  
7 is  
8 begin  
9   Var2 := Func (Var1);  
10  Var3 := Var3 + 1;  
11 end Proc;
```

Listing 79: func.adb

```
1 function Func (Var : Integer) return Integer  
2 is  
3 begin  
4   return Var + 1;  
5 end Func;
```

Listing 80: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;  
2 with Proc;  
3  
4 procedure Main is  
5   V1, V2 : Integer;  
6 begin  
7   V2 := 2;  
8   Proc (5, V1, V2);  
9  
10  Put_Line ("V1: " & Integer'Image (V1));  
11  Put_Line ("V2: " & Integer'Image (V2));  
12 end Main;
```

Build output

```
Compile  
  [Ada]      main.adb  
  [Ada]      proc.adb  
  [Ada]      func.adb  
Bind  
  [gprbind]  main.bexch  
  [Ada]      main.ali  
Link  
  [link]     main.adb
```

Runtime output

```
V1:  6  
V2:  3
```

[C]

Listing 81: proc.h

```

1 void Proc
2   (int  Var1,
3    int * Var2,
4    int * Var3);

```

Listing 82: func.h

```

1 int Func (int Var);

```

Listing 83: proc.c

```

1 #include "func.h"
2
3 void Proc
4   (int  Var1,
5    int * Var2,
6    int * Var3)
7 {
8     *Var2 = Func (Var1);
9     *Var3 += 1;
10 }

```

Listing 84: func.c

```

1 int Func (int Var)
2 {
3     return Var + 1;
4 }

```

Listing 85: main.c

```

1 #include <stdio.h>
2 #include "proc.h"
3
4 int main(int argc, const char * argv[])
5 {
6     int v1, v2;
7
8     v2 = 2;
9     Proc (5, &v1, &v2);
10
11     printf("v1: %d\n", v1);
12     printf("v2: %d\n", v2);
13
14     return 0;
15 }

```

Runtime output

```

v1: 6
v2: 3

```

The first two declarations for Proc and Func are specifications of the subprograms which are being provided later. Although optional here, it's still considered good practice to separately define specifications and implementations in order to make it easier to read the program. In Ada and C, a function that has not yet been seen cannot be used. Here, Proc can call Func because its specification has been declared.

Parameters in Ada subprogram declarations are separated with semicolons, because commas are

reserved for listing multiple parameters of the same type. Parameter declaration syntax is the same as variable declaration syntax (except for the modes), including default values for parameters. If there are no parameters, the parentheses must be omitted entirely from both the declaration and invocation of the subprogram.

In Ada 202X

Ada 202X allows for using static expression functions, which are evaluated at compile time. To achieve this, we can use an aspect — we'll discuss aspects *later in this chapter* (page 65).

An expression function is static when the `Static` aspect is specified. For example:

```
procedure Main is
    X1 : constant := (if True then 37 else 42);

    function If_Then_Else (Flag : Boolean; X, Y : Integer)
        return Integer is
        (if Flag then X else Y) with Static;

    X2 : constant := If_Then_Else (True, 37, 42);

begin
    null;
end Main;
```

In this example, we declare `X1` using an expression. In the declaration of `X2`, we call the static expression function `If_Then_Else`. Both `X1` and `X2` have the same constant value.

2.12.2 Overloading

In C, function names must be unique. Ada allows overloading, in which multiple subprograms can share the same name as long as the subprogram signatures (the parameter types, and function return types) are different. The compiler will be able to resolve the calls to the proper routines or it will reject the calls. For example:

[Ada]

Listing 86: machine.ads

```
1 package Machine is
2     type Status is (Off, On);
3     type Code is new Integer range 0 .. 3;
4     type Threshold is new Float range 0.0 .. 10.0;
5
6     function Get (S : Status) return Code;
7     function Get (S : Status) return Threshold;
8
9 end Machine;
```

Listing 87: machine.adb

```
1 package body Machine is
2
3     function Get (S : Status) return Code is
4     begin
5         case S is
6             when Off => return 1;
7             when On  => return 3;
```

(continues on next page)

(continued from previous page)

```

8     end case;
9 end Get;

10
11 function Get (S : Status) return Threshold is
12 begin
13     case S is
14         when Off => return 2.0;
15         when On  => return 10.0;
16     end case;
17 end Get;
18
19 end Machine;

```

Listing 88: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Machine;     use Machine;
3
4 procedure Main is
5     S : Status;
6     C : Code;
7     T : Threshold;
8 begin
9     S := On;
10    C := Get (S);
11    T := Get (S);
12
13    Put_Line ("S: " & Status'Image (S));
14    Put_Line ("C: " & Code'Image (C));
15    Put_Line ("T: " & Threshold'Image (T));
16 end Main;

```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      machine.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb

```

Runtime output

```

S: ON
C: 3
T: 1.00000E+01

```

The Ada compiler knows that an assignment to C requires a Code value. So, it chooses the Get function that returns a Code to satisfy this requirement.

Operators in Ada are functions too. This allows you to define local operators that override operators defined at an outer scope, and provide overloaded operators that operate on and compare different types. To declare an operator as a function, enclose its "name" in quotes:

```
[Ada]
```

Listing 89: machine_2.ads

```
1 package Machine_2 is
2   type Status is (Off, Waiting, On);
3   type Input is new Float range 0.0 .. 10.0;
4
5   function Get (I : Input) return Status;
6
7   function "=" (Left : Input; Right : Status) return Boolean;
8
9 end Machine_2;
```

Listing 90: machine_2.adb

```
1 package body Machine_2 is
2
3   function Get (I : Input) return Status is
4   begin
5     if I >= 0.0 and I < 3.0 then
6       return Off;
7     elsif I >= 3.0 and I < 6.5 then
8       return Waiting;
9     else
10      return On;
11    end if;
12  end Get;
13
14  function "=" (Left : Input; Right : Status) return Boolean is
15  begin
16    return Get (Left) = Right;
17  end "=";
18
19 end Machine_2;
```

Listing 91: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Machine_2;   use Machine_2;
3
4 procedure Main is
5   I : Input;
6 begin
7   I := 3.0;
8   if I = Off then
9     Put_Line ("Machine is off.");
10  else
11    Put_Line ("Machine is not off.");
12  end if;
13 end Main;
```

Build output

```
Compile
  [Ada]      main.adb
  [Ada]      machine_2.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb
```

Runtime output

```
Machine is not off.
```

2.12.3 Aspects

Aspect specifications allow you to define certain characteristics of a declaration using the `with` keyword after the declaration:

```
procedure Some_Procedure is <procedure_definition>
  with Some_Aspect => <aspect_specification>;

function Some_Function is <function_definition>
  with Some_Aspect => <aspect_specification>;

type Some_Type is <type_definition>
  with Some_Aspect => <aspect_specification>;

Obj : Some_Type with Some_Aspect => <aspect_specification>;
```

For example, you can inline a subprogram by specifying the `Inline` aspect:

[Ada]

Listing 92: float_arrays.ads

```
1 package Float_Arrays is
2
3   type Float_Array is array (Positive range <>) of Float;
4
5   function Average (Data : Float_Array) return Float
6     with Inline;
7
8 end Float_Arrays;
```

We'll discuss inlining *later in this course* (page 204).

Aspect specifications were introduced in Ada 2012. In previous versions of Ada, you had to use a pragma instead. The previous example would be written as follows:

[Ada]

Listing 93: float_arrays.ads

```
1 package Float_Arrays is
2
3   type Float_Array is array (Positive range <>) of Float;
4
5   function Average (Data : Float_Array) return Float;
6
7   pragma Inline (Average);
8
9 end Float_Arrays;
```

Aspects and attributes might refer to the same kind of information. For example, we can use the `Size` aspect to define the expected minimum size of objects of a certain type:

[Ada]

Listing 94: my_device_types.ads

```
1 package My_Device_Types is
2
```

(continues on next page)

(continued from previous page)

```
3   type UInt10 is mod 2 ** 10
4     with Size => 10;
5
6 end My_Device_Types;
```

In the same way, we can use the size attribute to retrieve the size of a type or of an object:

[Ada]

Listing 95: show_device_types.adb

```
1 with Ada.Text_IO;      use Ada.Text_IO;
2
3 with My_Device_Types; use My_Device_Types;
4
5 procedure Show_Device_Types is
6   UInt10_Obj : constant UInt10 := 0;
7 begin
8   Put_Line ("Size of UInt10 type:  " & Positive'Image (UInt10'Size));
9   Put_Line ("Size of UInt10 object: " & Positive'Image (UInt10_Obj'Size));
10 end Show_Device_Types;
```

We'll explain both Size aspect and Size attribute *later in this course* (page 95).

CONCURRENCY AND REAL-TIME

3.1 Understanding the various options

Concurrent and real-time programming are standard parts of the Ada language. As such, they have the same semantics, whether executing on a native target with an OS such as Linux, on a real-time operating system (RTOS) such as VxWorks, or on a bare metal target with no OS or RTOS at all.

For resource-constrained systems, two subsets of the Ada concurrency facilities are defined, known as the Ravenscar and Jorvik profiles. Though restricted, these subsets have highly desirable properties, including: efficiency, predictability, analyzability, absence of deadlock, bounded blocking, absence of priority inversion, a real-time scheduler, and a small memory footprint. On bare metal systems, this means in effect that Ada comes with its own real-time kernel.

For further information

We'll discuss the Ravenscar profile *later in this chapter* (page 78). Details about the Jorvik profile can be found elsewhere [[Jorvik](#)].

Enhanced portability and expressive power are the primary advantages of using the standard concurrency facilities, potentially resulting in considerable cost savings. For example, with little effort, it is possible to migrate from Windows to Linux to a bare machine without requiring any changes to the code. Thread management and synchronization is all done by the implementation, transparently. However, in some situations, it's critical to be able to access directly the services provided by the platform. In this case, it's always possible to make direct system calls from Ada code. Several targets of the GNAT compiler provide this sort of API by default, for example win32ada for Windows and Florist for POSIX systems.

On native and RTOS-based platforms GNAT typically provides the full concurrency facilities. In contrast, on bare metal platforms GNAT typically provides the two standard subsets: Ravenscar and Jorvik.

3.2 Tasks

Ada offers a high level construct called a *task* which is an independent thread of execution. In GNAT, tasks are either mapped to the underlying OS threads, or use a dedicated kernel when not available.

The following example will display the 26 letters of the alphabet twice, using two concurrent tasks. Since there is no synchronization between the two threads of control in any of the examples, the output may be interspersed.

[Ada]

Listing 1: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is -- implicitly called by the environment task
4     subtype A_To_Z is Character range 'A' .. 'Z';
5
6     task My_Task;
7
8     task body My_Task is
9     begin
10         for I in A_To_Z'Range loop
11             Put (I);
12         end loop;
13         New_Line;
14     end My_Task;
15 begin
16     for I in A_To_Z'Range loop
17         Put (I);
18     end loop;
19     New_Line;
20 end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
ABCDEFGHIJKLMNOPQRSTUVWXYZ
ABCDEFGHIJKLMNOPQRSTUVWXYZ
```

Any number of Ada tasks may be declared in any declarative region. A task declaration is very similar to a procedure or package declaration. They all start automatically when control reaches the begin. A block will not exit until all sequences of statements defined within that scope, including those in tasks, have been completed.

A task type is a generalization of a task object; each object of a task type has the same behavior. A declared object of a task type is started within the scope where it is declared, and control does not leave that scope until the task has terminated.

Task types can be parameterized; the parameter serves the same purpose as an argument to a constructor in Java. The following example creates 10 tasks, each of which displays a subset of the alphabet contained between the parameter and the 'Z' Character. As with the earlier example, since there is no synchronization among the tasks, the output may be interspersed depending on the underlying implementation of the task scheduling algorithm.

[Ada]

Listing 2: my_tasks.ads

```
1 package My_Tasks is
2
3     task type My_Task (First : Character);
4
5 end My_Tasks;
```

Listing 3: my_tasks.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body My_Tasks is
4
5     task body My_Task is
6     begin
7         for I in First .. 'Z' loop
8             Put (I);
9         end loop;
10        New_Line;
11    end My_Task;
12
13 end My_Tasks;

```

Listing 4: main.adb

```

1 with My_Tasks; use My_Tasks;
2
3 procedure Main is
4     Dummy_Tab : array (0 .. 3) of My_Task ('W');
5 begin
6     null;
7 end Main;

```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      my_tasks.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb

```

Runtime output

```

WXYZ
WXYZ
WXYZ
WXYZ

```

In Ada, a task may be dynamically allocated rather than declared statically. The task will then start as soon as it has been allocated, and terminates when its work is completed.

[Ada]

Listing 5: main.adb

```

1 with My_Tasks; use My_Tasks;
2
3 procedure Main is
4     type Ptr_Task is access My_Task;
5
6     T : Ptr_Task;
7 begin
8     T := new My_Task ('W');
9 end Main;

```

Build output

```
Compile
  [Ada]          main.adb
  [Ada]          my_tasks.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```
WXYZ
```

3.3 Rendezvous

A rendezvous is a synchronization between two tasks, allowing them to exchange data and coordinate execution. Let's consider the following example:

[Ada]

Listing 6: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4
5     task After is
6         entry Go;
7     end After;
8
9     task body After is
10    begin
11        accept Go;
12        Put_Line ("After");
13    end After;
14
15 begin
16     Put_Line ("Before");
17     After.Go;
18 end Main;
```

Build output

```
Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```
Before
After
```

The Go entry declared in After is the client interface to the task. In the task body, the accept statement causes the task to wait for a call on the entry. This particular entry and accept pair simply causes the task to wait until Main calls After.Go. So, even though the two tasks start

simultaneously and execute independently, they can coordinate via Go. Then, they both continue execution independently after the rendezvous.

The entry/accept pair can take/pass parameters, and the accept statement can contain a sequence of statements; while these statements are executed, the caller is blocked.

Let's look at a more ambitious example. The rendezvous below accepts parameters and executes some code:

[Ada]

Listing 7: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4
5     task After is
6         entry Go (Text : String);
7     end After;
8
9     task body After is
10    begin
11        accept Go (Text : String) do
12            Put_Line ("After: " & Text);
13        end Go;
14    end After;
15
16 begin
17     Put_Line ("Before");
18     After.Go ("Main");
19 end Main;
```

Build output

```

Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```

Before
After: Main
```

In the above example, the Put_Line is placed in the accept statement. Here's a possible execution trace, assuming a uniprocessor:

1. At the begin of Main, task After is started and the main procedure is suspended.
2. After reaches the accept statement and is suspended, since there is no pending call on the Go entry.
3. The main procedure is awakened and executes the Put_Line invocation, displaying the string "Before".
4. The main procedure calls the Go entry. Since After is suspended on its accept statement for this entry, the call succeeds.
5. The main procedure is suspended, and the task After is awakened to execute the body of the accept statement. The actual parameter "Main" is passed to the accept statement, and the Put_Line invocation is executed. As a result, the string "After: Main" is displayed.

6. When the accept statement is completed, both the After task and the main procedure are ready to run. Suppose that the Main procedure is given the processor. It reaches its end, but the local task After has not yet terminated. The main procedure is suspended.
7. The After task continues, and terminates since it is at its end. The main procedure is resumed, and it too can terminate since its dependent task has terminated.

The above description is a conceptual model; in practice the implementation can perform various optimizations to avoid unnecessary context switches.

3.4 Selective Rendezvous

The accept statement by itself can only wait for a single event (call) at a time. The select statement allows a task to listen for multiple events simultaneously, and then to deal with the first event to occur. This feature is illustrated by the task below, which maintains an integer value that is modified by other tasks that call Increment, Decrement, and Get:

[Ada]

Listing 8: counters.ads

```
1 package Counters is
2
3     task Counter is
4         entry Get (Result : out Integer);
5         entry Increment;
6         entry Decrement;
7     end Counter;
8
9 end Counters;
```

Listing 9: counters.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body Counters is
4
5     task body Counter is
6         Value : Integer := 0;
7     begin
8         loop
9             select
10                 accept Increment do
11                     Value := Value + 1;
12                 end Increment;
13             or
14                 accept Decrement do
15                     Value := Value - 1;
16                 end Decrement;
17             or
18                 accept Get (Result : out Integer) do
19                     Result := Value;
20                 end Get;
21             or
22                 delay 5.0;
23                 Put_Line ("Exiting Counter task...");
24                 exit;
25             end select;
26         end loop;
27     end Counter;
```

(continues on next page)

(continued from previous page)

```

28
29 end Counters;

```

Listing 10: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2  with Counters;    use Counters;
3
4  procedure Main is
5      V : Integer;
6  begin
7      Put_Line ("Main started.");
8
9      Counter.Get (V);
10     Put_Line ("Got value. Value = " & Integer'Image (V));
11
12     Counter.Increment;
13     Put_Line ("Incremented value.");
14
15     Counter.Increment;
16     Put_Line ("Incremented value.");
17
18     Counter.Get (V);
19     Put_Line ("Got value. Value = " & Integer'Image (V));
20
21     Counter.Decrement;
22     Put_Line ("Decrement value.");
23
24     Counter.Get (V);
25     Put_Line ("Got value. Value = " & Integer'Image (V));
26
27     Put_Line ("Main finished.");
28 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
  [Ada]          counters.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```

Main started.
Got value. Value = 0
Incremented value.
Incremented value.
Got value. Value = 2
Decrement value.
Got value. Value = 1
Main finished.
Exiting Counter task...

```

When the task's statement flow reaches the select, it will wait for all four events — three entries and a delay — in parallel. If the delay of five seconds is exceeded, the task will execute the statements following the delay statement (and in this case will exit the loop, in effect terminating the task). The accept bodies for the Increment, Decrement, or Get entries will be otherwise executed as

they're called. These four sections of the select statement are mutually exclusive: at each iteration of the loop, only one will be invoked. This is a critical point; if the task had been written as a package, with procedures for the various operations, then a *race condition* could occur where multiple tasks simultaneously calling, say, `Increment`, cause the value to only get incremented once. In the tasking version, if multiple tasks simultaneously call `Increment` then only one at a time will be accepted, and the value will be incremented by each of the tasks when it is accepted.

More specifically, each entry has an associated queue of pending callers. If a task calls one of the entries and `Counter` is not ready to accept the call (i.e., if `Counter` is not suspended at the select statement) then the calling task is suspended, and placed in the queue of the entry that it is calling. From the perspective of the `Counter` task, at any iteration of the loop there are several possibilities:

- There is no call pending on any of the entries. In this case `Counter` is suspended. It will be awakened by the first of two events: a call on one of its entries (which will then be immediately accepted), or the expiration of the five second delay (whose effect was noted above).
- There is a call pending on exactly one of the entries. In this case control passes to the select branch with an `accept` statement for that entry.
- There are calls pending on more than one entry. In this case one of the entries with pending callers is chosen, and then one of the callers is chosen to be de-queued. The choice of which caller to accept depends on the queuing policy, which can be specified via a `pragma` defined in the Real-Time Systems Annex of the Ada standard; the default is *First-In First-Out*.

3.5 Protected Objects

Although the rendezvous may be used to implement mutually exclusive access to a shared data object, an alternative (and generally preferable) style is through a protected object, an efficiently implementable mechanism that makes the effect more explicit. A protected object has a public interface (its protected operations) for accessing and manipulating the object's components (its private part). Mutual exclusion is enforced through a conceptual lock on the object, and encapsulation ensures that the only external access to the components are through the protected operations.

Two kinds of operations can be performed on such objects: read-write operations by procedures or entries, and read-only operations by functions. The lock mechanism is implemented so that it's possible to perform concurrent read operations but not concurrent write or read/write operations.

Let's reimplement our earlier tasking example with a protected object called `Counter`:

[Ada]

Listing 11: counters.ads

```
1 package Counters is
2
3     protected Counter is
4         function Get return Integer;
5         procedure Increment;
6         procedure Decrement;
7     private
8         Value : Integer := 0;
9     end Counter;
10
11 end Counters;
```

Listing 12: counters.adb

```
1 package body Counters is
2
3     protected body Counter is
```

(continues on next page)

(continued from previous page)

```

4      function Get return Integer is
5      begin
6          return Value;
7      end Get;
8
9      procedure Increment is
10     begin
11         Value := Value + 1;
12     end Increment;
13
14     procedure Decrement is
15     begin
16         Value := Value - 1;
17     end Decrement;
18 end Counter;
19
20 end Counters;

```

Having two completely different ways to implement the same paradigm might seem complicated. However, in practice the actual problem to solve usually drives the choice between an active structure (a task) or a passive structure (a protected object).

A protected object can be accessed through prefix notation:

[Ada]

Listing 13: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2  with Counters;    use Counters;
3
4  procedure Main is
5  begin
6      Counter.Increment;
7      Counter.Decrement;
8      Put_Line (Integer'Image (Counter.Get));
9  end Main;

```

Build output

Compile	
[Ada]	main.adb
[Ada]	counters.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

0

A protected object may look like a package syntactically, since it contains declarations that can be accessed externally using prefix notation. However, the declaration of a protected object is extremely restricted; for example, no public data is allowed, no types can be declared inside, etc. And besides the syntactic differences, there is a critical semantic distinction: a protected object has a conceptual lock that guarantees mutual exclusion; there is no such lock for a package.

Like tasks, it's possible to declare protected types that can be instantiated several times:

```
declare
  protected type Counter is
    -- as above
  end Counter;

  protected body Counter is
    -- as above
  end Counter;

  C1 : Counter;
  C2 : Counter;
begin
  C1.Increment;
  C2.Decrement;
  . . .
end;
```

Protected objects and types can declare a procedure-like operation known as an *entry*. An entry is somewhat similar to a procedure but includes a so-called barrier condition that must be true in order for the entry invocation to succeed. Calling a protected entry is thus a two step process: first, acquire the lock on the object, and then evaluate the barrier condition. If the condition is true then the caller will execute the entry body. If the condition is false, then the caller is placed in the queue for the entry, and relinquishes the lock. Barrier conditions (for entries with non-empty queues) are reevaluated upon completion of protected procedures and protected entries.

Here's an example illustrating protected entries: a protected type that models a binary semaphore / persistent signal.

[Ada]

Listing 14: binary_semaphores.ads

```
1 package Binary_Semaphores is
2
3   protected type Binary_Semaphore is
4     entry Wait;
5     procedure Signal;
6   private
7     Signaled : Boolean := False;
8   end Binary_Semaphore;
9
10 end Binary_Semaphores;
```

Listing 15: binary_semaphores.adb

```
1 package body Binary_Semaphores is
2
3   protected body Binary_Semaphore is
4     entry Wait when Signaled is
5     begin
6       Signaled := False;
7     end Wait;
8
9     procedure Signal is
10    begin
11      Signaled := True;
12    end Signal;
13  end Binary_Semaphore;
14
15 end Binary_Semaphores;
```

Listing 16: main.adb

```

1  with Ada.Text_IO;          use Ada.Text_IO;
2  with Binary_Semaphores; use Binary_Semaphores;
3
4  procedure Main is
5      B : Binary_Semaphore;
6
7      task T1;
8      task T2;
9
10     task body T1 is
11     begin
12         Put_Line ("Task T1 waiting...");
13         B.Wait;
14
15         Put_Line ("Task T1.");
16         delay 1.0;
17
18         Put_Line ("Task T1 will signal...");
19         B.Signal;
20
21         Put_Line ("Task T1 finished.");
22     end T1;
23
24     task body T2 is
25     begin
26         Put_Line ("Task T2 waiting...");
27         B.Wait;
28
29         Put_Line ("Task T2");
30         delay 1.0;
31
32         Put_Line ("Task T2 will signal...");
33         B.Signal;
34
35         Put_Line ("Task T2 finished.");
36     end T2;
37
38     begin
39         Put_Line ("Main started.");
40         B.Signal;
41         Put_Line ("Main finished.");
42     end Main;

```

Build output

```

Compile
  [Ada]          main.adb
  [Ada]          binary_semaphores.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```

Task T1 waiting...
Main started.
Main finished.

```

(continues on next page)

(continued from previous page)

```
Task T2 waiting...
Task T1.
Task T1 will signal...
Task T1 finished.
Task T2
Task T2 will signal...
Task T2 finished.
```

Ada concurrency features provide much further generality than what's been presented here. For additional information please consult one of the works cited in the *References* section.

3.6 Ravenscar

The Ravenscar profile is a subset of the Ada concurrency facilities that supports determinism, schedulability analysis, constrained memory utilization, and certification to the highest integrity levels. Four distinct application domains are intended:

- hard real-time applications requiring predictability,
- safety-critical systems requiring formal, stringent certification,
- high-integrity applications requiring formal static analysis and verification,
- embedded applications requiring both a small memory footprint and low execution overhead.

Tasking constructs that preclude analysis, either technically or economically, are disallowed. You can use the pragma `Profile` (Ravenscar) to indicate that the Ravenscar restrictions must be observed in your program.

Some of the examples we've seen above will be rejected by the compiler when using the Ravenscar profile. For example:

[Ada]

Listing 17: my_tasks.ads

```
1 package My_Tasks is
2
3     task type My_Task (First : Character);
4
5 end My_Tasks;
```

Listing 18: my_tasks.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body My_Tasks is
4
5     task body My_Task is
6     begin
7         for C in First .. 'Z' loop
8             Put (C);
9         end loop;
10        New_Line;
11    end My_Task;
12
13 end My_Tasks;
```

Listing 19: main.adb

```

1  pragma Profile (Ravenscar);
2
3  with My_Tasks; use My_Tasks;
4
5  procedure Main is
6      Tab : array (0 .. 3) of My_Task ('W');
7  begin
8      null;
9  end Main;

```

Compilation output

```

main.adb:6:04: error: violation of restriction "No_Task_Hierarchy"
main.adb:6:04: error: from profile "Ravenscar" at line 1

```

This code violates the *No_Task_Hierarchy* restriction of the Ravenscar profile. This is due to the declaration of `Tab` in the `Main` procedure. Ravenscar requires task declarations to be done at the library level. Therefore, a simple solution is to create a separate package and reference it in the main application:

[Ada]

Listing 20: my_task_inst.ads

```

1  with My_Tasks; use My_Tasks;
2
3  package My_Task_Inst is
4
5      Tab : array (0 .. 3) of My_Task ('W');
6
7  end My_Task_Inst;

```

Listing 21: main.adb

```

1  pragma Profile (Ravenscar);
2
3  with My_Task_Inst;
4
5  procedure Main is
6  begin
7      null;
8  end Main;

```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      my_task_inst.ads
  [Ada]      my_tasks.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb

```

Runtime output

```

WXYZ
WXYZ

```

(continues on next page)

(continued from previous page)

```
WXYZ  
WXYZ
```

Also, Ravenscar prohibits entries for tasks. For example, we're not allowed to write this declaration:

```
task type My_Task (First : Character) is  
  entry Start;  
end My_Task;
```

You can use, however, one entry per protected object. As an example, the declaration of the `Binary_Semaphore` type that we've discussed before compiles fine with Ravenscar:

```
protected type Binary_Semaphore is  
  entry Wait;  
  procedure Signal;  
private  
  Signaled : Boolean := False;  
end Binary_Semaphore;
```

We could add more procedures and functions to the declaration of `Binary_Semaphore`, but we wouldn't be able to add another entry when using Ravenscar.

Similar to the previous example with the task array declaration, objects of `Binary_Semaphore` cannot be declared in the main application:

```
procedure Main is  
  B : Binary_Semaphore;  
begin  
  null;  
end Main;
```

This violates the *No_Local_Protected_Objects* restriction. Again, Ravenscar expects this declaration to be done on a library level, so a solution to make this code compile is to have this declaration in a separate package and reference it in the `Main` procedure.

Ravenscar offers many additional restrictions. Covering those would exceed the scope of this chapter. You can find more examples using the Ravenscar profile on [this blog post](https://blog.adacore.com/theres-a-mini-rtos-in-my-language)¹⁰.

¹⁰ <https://blog.adacore.com/theres-a-mini-rtos-in-my-language>

WRITING ADA ON EMBEDDED SYSTEMS

4.1 Understanding the Ada Run-Time

Ada supports a high level of abstractness and expressiveness. In some cases, the compiler translates those constructs directly into machine code. However, there are many high-level constructs for which a direct compilation would be difficult. In those cases, the compiler links to a library containing an implementation of those high-level constructs: this is the so-called run-time library.

One typical example of high-level constructs that can be cumbersome for direct machine code generation is Ada source-code using tasking. In this case, linking to a low-level implementation of multithreading support — for example, an implementation using POSIX threads — is more straightforward than trying to make the compiler generate all the machine code.

In the case of GNAT, the run-time library is implemented using both C and Ada source-code. Also, depending on the operating system, the library will interface with low-level functionality from the target operating system.

There are basically two types of run-time libraries:

- the **standard** run-time library: in many cases, this is the run-time library available on desktop operating systems or on some embedded platforms (such as ARM-Linux on a Raspberry-Pi).
- the **configurable** run-time library: this is a capability that is used to create custom run-time libraries for specific target devices.

Configurable run-time libraries are usually used for constrained target devices where support for the full library would be difficult or even impossible. In this case, configurable run-time libraries may support just a subset of the full Ada language. There are many reasons that speak for this approach:

- Some aspects of the Ada language may not translate well to limited operating systems.
- Memory constraints may require reducing the size of the run-time library, so that developers may need to replace or even remove parts of the library.
- When certification is required, those parts of the library that would require too much certification effort can be removed.

When using a configurable run-time library, the compiler checks whether the library supports certain features of the language. If a feature isn't supported, the compiler will give an error message.

You can find further information about the run-time library on [this chapter of the GNAT User's Guide Supplement for Cross Platforms](#)¹¹

¹¹ https://docs.adacore.com/gnat_ugx-docs/html/gnat_ugx/gnat_ugx/the_gnat_configurable_run_time_facility.html

4.2 Low Level Programming

4.2.1 Representation Clauses

We've seen in the previous chapters how Ada can be used to describe high level semantics and architecture. The beauty of the language, however, is that it can be used all the way down to the lowest levels of the development, including embedded assembly code or bit-level data management.

One very interesting feature of the language is that, unlike C, for example, there are no data representation constraints unless specified by the developer. This means that the compiler is free to choose the best trade-off in terms of representation vs. performance. Let's start with the following example:

[Ada]

```
type R is record
  V : Integer range 0 .. 255;
  B1 : Boolean;
  B2 : Boolean;
end record
with Pack;
```

[C]

```
struct R {
    unsigned int v:8;
    bool b1;
    bool b2;
};
```

The Ada and the C code above both represent efforts to create an object that's as small as possible. Controlling data size is not possible in Java, but the language does specify the size of values for the primitive types.

Although the C and Ada code are equivalent in this particular example, there's an interesting semantic difference. In C, the number of bits required by each field needs to be specified. Here, we're stating that `v` is only 8 bits, effectively representing values from 0 to 255. In Ada, it's the other way around: the developer specifies the range of values required and the compiler decides how to represent things, optimizing for speed or size. The `Pack` aspect declared at the end of the record specifies that the compiler should optimize for size even at the expense of decreased speed in accessing record components. We'll see more details about the `Pack` aspect in the sections about *bitwise operations* (page 135) and *mapping structures to bit-fields* (page 137) in chapter 6.

Other representation clauses can be specified as well, along with compile-time consistency checks between requirements in terms of available values and specified sizes. This is particularly useful when a specific layout is necessary; for example when interfacing with hardware, a driver, or a communication protocol. Here's how to specify a specific data layout based on the previous example:

[Ada]

```
type R is record
  V : Integer range 0 .. 255;
  B1 : Boolean;
  B2 : Boolean;
end record;

for R use record
  -- Occupy the first bit of the first byte.
  B1 at 0 range 0 .. 0;
```

(continues on next page)

(continued from previous page)

```

-- Occupy the last 7 bits of the first byte,
-- as well as the first bit of the second byte.
V at 0 range 1 .. 8;

-- Occupy the second bit of the second byte.
B2 at 1 range 1 .. 1;
end record;
```

We omit the `with Pack` directive and instead use a record representation clause following the record declaration. The compiler is directed to spread objects of type `R` across two bytes. The layout we're specifying here is fairly inefficient to work with on any machine, but you can have the compiler construct the most efficient methods for access, rather than coding your own machine-dependent bit-level methods manually.

4.2.2 Embedded Assembly Code

When performing low-level development, such as at the kernel or hardware driver level, there can be times when it is necessary to implement functionality with assembly code.

Every Ada compiler has its own conventions for embedding assembly code, based on the hardware platform and the supported assembler(s). Our examples here will work with GNAT and GCC on the x86 architecture.

All x86 processors since the Intel Pentium offer the `rdtsc` instruction, which tells us the number of cycles since the last processor reset. It takes no inputs and places an unsigned 64-bit value split between the `edx` and `eax` registers.

GNAT provides a subprogram called `System.Machine_Code.Asm` that can be used for assembly code insertion. You can specify a string to pass to the assembler as well as source-level variables to be used for input and output:

[Ada]

Listing 1: `get_processor_cycles.adb`

```

1 with System.Machine_Code; use System.Machine_Code;
2 with Interfaces;          use Interfaces;
3
4 function Get_Processor_Cycles return Unsigned_64 is
5   Low, High : Unsigned_32;
6   Counter   : Unsigned_64;
7 begin
8   Asm ("rdtsc",
9       Outputs =>
10         (Unsigned_32'Asm_Output ("=a", High),
11          Unsigned_32'Asm_Output ("=d", Low)),
12       Volatile => True);
13
14   Counter :=
15     Unsigned_64 (High) * 2 ** 32 +
16     Unsigned_64 (Low);
17
18   return Counter;
19 end Get_Processor_Cycles;
```

The `Unsigned_32'Asm_Output` clauses above provide associations between machine registers and source-level variables to be updated. `=a` and `=d` refer to the `eax` and `edx` machine registers, respectively. The use of the `Unsigned_32` and `Unsigned_64` types from package `Interfaces` ensures correct representation of the data. We assemble the two 32-bit values to form a single 64-bit value.

We set the `Volatile` parameter to `True` to tell the compiler that invoking this instruction multiple times with the same inputs can result in different outputs. This eliminates the possibility that the compiler will optimize multiple invocations into a single call.

With optimization turned on, the GNAT compiler is smart enough to use the `eax` and `edx` registers to implement the `High` and `Low` variables, resulting in zero overhead for the assembly interface.

The machine code insertion interface provides many features beyond what was shown here. More information can be found in the GNAT User's Guide, and the GNAT Reference manual.

4.3 Interrupt Handling

Handling interrupts is an important aspect when programming embedded devices. Interrupts are used, for example, to indicate that a hardware or software event has happened. Therefore, by handling interrupts, an application can react to external events.

Ada provides built-in support for handling interrupts. We can process interrupts by attaching a handler — which must be a protected procedure — to it. In the declaration of the protected procedure, we use the `Attach_Handler` aspect and indicate which interrupt we want to handle.

Let's look into a code example that *traps* the quit interrupt (`SIGQUIT`) on Linux:

[Ada]

Listing 2: `signal_handlers.ads`

```
1 with System.OS_Interface;
2
3 package Signal_Handlers is
4
5     protected type Quit_Handler is
6         function Requested return Boolean;
7     private
8         Quit_Request : Boolean := False;
9
10        --
11        -- Declaration of an interrupt handler for the "quit" interrupt:
12        --
13        procedure Handle_Quit_Signal
14            with Attach_Handler => System.OS_Interface.SIGQUIT;
15    end Quit_Handler;
16
17 end Signal_Handlers;
```

Listing 3: `signal_handlers.adb`

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body Signal_Handlers is
4
5     protected body Quit_Handler is
6
7         function Requested return Boolean is
8             (Quit_Request);
9
10        procedure Handle_Quit_Signal is
11            begin
12                Put_Line ("Quit request detected!");
13                Quit_Request := True;
14            end Handle_Quit_Signal;
```

(continues on next page)

(continued from previous page)

```

15
16     end Quit_Handler;
17
18 end Signal_Handlers;

```

Listing 4: test_quit_handler.adb

```

1  with Ada.Text_IO;      use Ada.Text_IO;
2  with Signal_Handlers;
3
4  procedure Test_Quit_Handler is
5      Quit : Signal_Handlers.Quit_Handler;
6
7  begin
8      while True loop
9          delay 1.0;
10         exit when Quit.Requested;
11     end loop;
12
13     Put_Line ("Exiting application...");
14 end Test_Quit_Handler;

```

The specification of the `Signal_Handlers` package from this example contains the declaration of `Quit_Handler`, which is a protected type. In the private part of this protected type, we declare the `Handle_Quit_Signal` procedure. By using the `Attach_Handler` aspect in the declaration of `Handle_Quit_Signal` and indicating the quit interrupt (`System.OS_Interface.SIGQUIT`), we're instructing the operating system to call this procedure for any quit request. So when the user presses CTRL+\ on their keyboard, for example, the application will behave as follows:

- the operating system calls the `Handle_Quit_Signal` procedure, which displays a message to the user ("Quit request detected!") and sets a Boolean variable — `Quit_Request`, which is declared in the `Quit_Handler` type;
- the main application checks the status of the quit handler by calling the `Requested` function as part of the `while True` loop;
 - This call is in the `exit when Quit.Requested` line.
 - The `Requested` function returns `True` in this case because the `Quit_Request` flag was set by the `Handle_Quit_Signal` procedure.
- the main application exits the loop, displays a message and finishes.

Note that the code example above isn't portable because it makes use of interrupts from the Linux operating system. When programming embedded devices, we would use instead the interrupts available on those specific devices.

Also note that, in the example above, we're declaring a static handler at compilation time. If you need to make use of dynamic handlers, which can be configured at runtime, you can use the subprograms from the `Ada.Interrupts` package. This package includes not only a version of `Attach_Handler` as a procedure, but also other procedures such as:

- `Exchange_Handler`, which lets us exchange, at runtime, the current handler associated with a specific interrupt by a different handler;
- `Detach_Handler`, which we can use to remove the handler currently associated with a given interrupt.

Details about the `Ada.Interrupts` package are out of scope for this course. We'll discuss them in a separate, more advanced course in the future. You can find some information about it in the [Interrupts appendix of the Ada Reference Manual](https://www.adaic.org/resources/add_content/standards/12aarm/html/AA-C-3-2.html)¹².

¹² https://www.adaic.org/resources/add_content/standards/12aarm/html/AA-C-3-2.html

4.4 Dealing with Absence of FPU with Fixed Point

Many numerical applications typically use floating-point types to compute values. However, in some platforms, a floating-point unit may not be available. Other platforms may have a floating-point unit, but using it in certain numerical algorithms can be prohibitive in terms of performance. For those cases, fixed-point arithmetic can be a good alternative.

The difference between fixed-point and floating-point types might not be so obvious when looking at this code snippet:

[Ada]

Listing 5: fixed_definitions.ads

```

1 package Fixed_Definitions is
2
3   D : constant := 2.0 ** (-31);
4
5   type Fixed is delta D range -1.0 .. 1.0 - D;
6
7 end Fixed_Definitions;
```

Listing 6: show_float_and_fixed_point.adb

```

1 with Ada.Text_IO;      use Ada.Text_IO;
2
3 with Fixed_Definitions; use Fixed_Definitions;
4
5 procedure Show_Float_And_Fixed_Point is
6   Float_Value : Float := 0.25;
7   Fixed_Value : Fixed := 0.25;
8 begin
9
10  Float_Value := Float_Value + 0.25;
11  Fixed_Value := Fixed_Value + 0.25;
12
13  Put_Line ("Float_Value = " & Float'Image (Float_Value));
14  Put_Line ("Fixed_Value = " & Fixed'Image (Fixed_Value));
15 end Show_Float_And_Fixed_Point;
```

Build output

```

Compile
[Ada]      show_float_and_fixed_point.adb
[Ada]      fixed_definitions.ads
Bind
[gprbind]  show_float_and_fixed_point.bexch
[Ada]      show_float_and_fixed_point.ali
Link
[link]     show_float_and_fixed_point.adb
```

Runtime output

```

Float_Value = 5.00000E-01
Fixed_Value = 0.5000000000
```

In this example, the application will show the value 0.5 for both `Float_Value` and `Fixed_Value`.

The major difference between floating-point and fixed-point types is in the way the values are stored. Values of ordinary fixed-point types are, in effect, scaled integers. The scaling used for ordinary fixed-point types is defined by the type's `small`, which is derived from the specified `delta` and, by default, is a power of two. Therefore, ordinary fixed-point types are sometimes called

binary fixed-point types. In that sense, ordinary fixed-point types can be thought of being close to the actual representation on the machine. In fact, ordinary fixed-point types make use of the available integer shift instructions, for example.

Another difference between floating-point and fixed-point types is that Ada doesn't provide standard fixed-point types — except for the `Duration` type, which is used to represent an interval of time in seconds. While the Ada standard specifies floating-point types such as `Float` and `Long_Float`, we have to declare our own fixed-point types. Note that, in the previous example, we have used a fixed-point type named `Fixed`: this type isn't part of the standard, but must be declared somewhere in the source-code of our application.

The syntax for an ordinary fixed-point type is

```
type <type_name> is delta <delta_value> range <lower_bound> .. <upper_bound>;
```

By default, the compiler will choose a scale factor, or `small`, that is a power of 2 no greater than `<delta_value>`.

For example, we may define a normalized range between -1.0 and 1.0 as following:

[Ada]

Listing 7: `normalized_fixed_point_type.adb`

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Normalized_Fixed_Point_Type is
4   D : constant := 2.0 ** (-31);
5   type TQ31 is delta D range -1.0 .. 1.0 - D;
6 begin
7   Put_Line ("TQ31 requires " & Integer'Image (TQ31'Size) & " bits");
8   Put_Line ("The delta value of TQ31 is " & TQ31'Image (TQ31'Delta));
9   Put_Line ("The minimum value of TQ31 is " & TQ31'Image (TQ31'First));
10  Put_Line ("The maximum value of TQ31 is " & TQ31'Image (TQ31'Last));
11 end Normalized_Fixed_Point_Type;
```

Build output

Compile	
[Ada]	<code>normalized_fixed_point_type.adb</code>
Bind	
[gprbind]	<code>normalized_fixed_point_type.bexch</code>
[Ada]	<code>normalized_fixed_point_type.ali</code>
Link	
[link]	<code>normalized_fixed_point_type.adb</code>

Runtime output

```
TQ31 requires 32 bits
The delta value of TQ31 is 0.0000000005
The minimum value of TQ31 is -1.0000000000
The maximum value of TQ31 is 0.9999999995
```

In this example, we are defining a 32-bit fixed-point data type for our normalized range. When running the application, we notice that the upper bound is close to one, but not exactly one. This is a typical effect of fixed-point data types — you can find more details in this discussion about the [Q format](https://en.wikipedia.org/wiki/Q_(number_format))¹³. We may also rewrite this code with an exact type definition:

[Ada]

¹³ [https://en.wikipedia.org/wiki/Q_\(number_format\)](https://en.wikipedia.org/wiki/Q_(number_format))

Listing 8: normalized_adapted_fixed_point_type.ads

```
1 package Normalized_Adapted_Fixed_Point_Type is
2
3     type TQ31 is delta 2.0 ** (-31) range -1.0 .. 1.0 - 2.0 ** (-31);
4
5 end Normalized_Adapted_Fixed_Point_Type;
```

We may also use any other range. For example:

[Ada]

Listing 9: custom_fixed_point_range.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Numerics; use Ada.Numerics;
3
4 procedure Custom_Fixed_Point_Range is
5     type Inv_Trig is delta 2.0 ** (-15) * Pi range -Pi / 2.0 .. Pi / 2.0;
6 begin
7     Put_Line ("Inv_Trig requires " & Integer'Image (Inv_Trig'Size)
8               & " bits");
9     Put_Line ("The delta value of Inv_Trig is "
10              & Inv_Trig'Image (Inv_Trig'Delta));
11     Put_Line ("The minimum value of Inv_Trig is "
12              & Inv_Trig'Image (Inv_Trig'First));
13     Put_Line ("The maximum value of Inv_Trig is "
14              & Inv_Trig'Image (Inv_Trig'Last));
15 end Custom_Fixed_Point_Range;
```

Build output

Compile	
[Ada]	custom_fixed_point_range.adb
Bind	
[gprbind]	custom_fixed_point_range.bexch
[Ada]	custom_fixed_point_range.ali
Link	
[link]	custom_fixed_point_range.adb

Runtime output

```
Inv_Trig requires 16 bits
The delta value of Inv_Trig is 0.00006
The minimum value of Inv_Trig is -1.57080
The maximum value of Inv_Trig is 1.57080
```

In this example, we are defining a 16-bit type called `Inv_Trig`, which has a range from $-\pi/2$ to $\pi/2$.

All standard operations are available for fixed-point types. For example:

[Ada]

Listing 10: fixed_point_op.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Fixed_Point_Op is
4     type TQ31 is delta 2.0 ** (-31) range -1.0 .. 1.0 - 2.0 ** (-31);
5
6     A, B, R : TQ31;
```

(continues on next page)

(continued from previous page)

```

7 begin
8   A := 0.25;
9   B := 0.50;
10  R := A + B;
11  Put_Line ("R is " & TQ31'Image (R));
12 end Fixed_Point_Op;

```

Build output

```

Compile
  [Ada]          fixed_point_op.adb
Bind
  [gprbind]      fixed_point_op.bexch
  [Ada]          fixed_point_op.ali
Link
  [link]         fixed_point_op.adb

```

Runtime output

```
R is 0.7500000000
```

As expected, R contains 0.75 after the addition of A and B.

In the case of C, since the language doesn't support fixed-point arithmetic, we need to emulate it using integer types and custom operations via functions. Let's look at this very rudimentary example:

[C]

Listing 11: main.c

```

1  #include <stdio.h>
2  #include <math.h>
3
4  #define SHIFT_FACTOR 32
5
6  #define TO_FIXED(x)    ((int) ((x) * pow (2.0, SHIFT_FACTOR - 1)))
7  #define TO_FLOAT(x)    ((float) ((double)(x) * (double)pow (2.0, -(SHIFT_FACTOR -
8  ↪1))))
9
10 typedef int fixed;
11
12 fixed add (fixed a, fixed b)
13 {
14     return a + b;
15 }
16
17 fixed mult (fixed a, fixed b)
18 {
19     return (fixed)(((long)a * (long)b) >> (SHIFT_FACTOR - 1));
20 }
21
22 void display_fixed (fixed x)
23 {
24     printf("value (integer) = %d\n", x);
25     printf("value (float)   = %3.5f\n\n", TO_FLOAT (x));
26 }
27
28 int main(int argc, const char * argv[])
29 {
30     int fixed_value = TO_FIXED(0.25);

```

(continues on next page)

(continued from previous page)

```
30
31     printf("Original value\n");
32     display_fixed(fixed_value);
33
34     printf("... + 0.25\n");
35     fixed_value = add(fixed_value, TO_FIXED(0.25));
36     display_fixed(fixed_value);
37
38     printf("... * 0.5\n");
39     fixed_value = mult(fixed_value, TO_FIXED(0.5));
40     display_fixed(fixed_value);
41
42     return 0;
43 }
```

Runtime output

```
Original value
value (integer) = 536870912
value (float)   = 0.25000

... + 0.25
value (integer) = 1073741824
value (float)   = 0.50000

... * 0.5
value (integer) = 536870912
value (float)   = 0.25000
```

Here, we declare the fixed-point type `fixed` based on `int` and two operations for it: addition (via the `add` function) and multiplication (via the `mult` function). Note that, while fixed-point addition is quite straightforward, multiplication requires right-shifting to match the correct internal representation. In Ada, since fixed-point operations are part of the language specification, they don't need to be emulated. Therefore, no extra effort is required from the programmer.

Also note that the example above is very rudimentary, so it doesn't take some of the side-effects of fixed-point arithmetic into account. In C, you have to manually take all side-effects deriving from fixed-point arithmetic into account, while in Ada, the compiler takes care of selecting the right operations for you.

4.5 Volatile and Atomic data

Ada has built-in support for handling both volatile and atomic data. Let's start by discussing volatile objects.

4.5.1 Volatile

A [volatile](#)¹⁴ object can be described as an object in memory whose value may change between two consecutive memory accesses of a process A — even if process A itself hasn't changed the value. This situation may arise when an object in memory is being shared by multiple threads. For example, a thread B may modify the value of that object between two read accesses of a thread A. Another typical example is the one of [memory-mapped I/O](#)¹⁵, where the hardware might be constantly changing the value of an object in memory.

Because the value of a volatile object may be constantly changing, a compiler cannot generate code that stores the value of that object into a register and use the value from the register in subsequent operations. Storing into a register is avoided because, if the value is stored there, it would be outdated if another process had changed the volatile object in the meantime. Instead, the compiler generates code in such a way that the process must read the value of the volatile object from memory for each access.

Let's look at a simple example of a volatile variable in C:

[C]

Listing 12: main.c

```

1  #include <stdio.h>
2
3  int main(int argc, const char * argv[])
4  {
5      volatile double val = 0.0;
6      int i;
7
8      for (i = 0; i < 1000; i++)
9      {
10         val += i * 2.0;
11     }
12     printf ("val: %5.3f\n", val);
13
14     return 0;
15 }
```

Runtime output

```
val: 999000.000
```

In this example, `val` has the modifier `volatile`, which indicates that the compiler must handle `val` as a volatile object. Therefore, each read and write access in the loop is performed by accessing the value of `val` in then memory.

This is the corresponding implementation in Ada:

[Ada]

Listing 13: show_volatile_object.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Show_Volatile_Object is
4      Val : Long_Float with Volatile;
5  begin
6      Val := 0.0;
7      for I in 0 .. 999 loop
8          Val := Val + 2.0 * Long_Float (I);
```

(continues on next page)

¹⁴ [https://en.wikipedia.org/wiki/Volatile_\(computer_programming\)](https://en.wikipedia.org/wiki/Volatile_(computer_programming))

¹⁵ https://en.wikipedia.org/wiki/Memory-mapped_I/O

(continued from previous page)

```

9   end loop;
10
11   Put_Line ("Val: " & Long_Float'Image (Val));
12 end Show_Volatile_Object;

```

Build output

```

Compile
  [Ada]          show_volatile_object.adb
Bind
  [gprbind]      show_volatile_object.bexch
  [Ada]          show_volatile_object.ali
Link
  [link]         show_volatile_object.adb

```

Runtime output

```
Val:  9.990000000000000E+05
```

In this example, Val has the `Volatile` aspect, which makes the object volatile. We can also use the `Volatile` aspect in type declarations. For example:

```
[Ada]
```

Listing 14: show_volatile_type.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Show_Volatile_Type is
4    type Volatile_Long_Float is new Long_Float with Volatile;
5
6    Val : Volatile_Long_Float;
7  begin
8    Val := 0.0;
9    for I in 0 .. 999 loop
10     Val := Val + 2.0 * Volatile_Long_Float (I);
11   end loop;
12
13   Put_Line ("Val: " & Volatile_Long_Float'Image (Val));
14 end Show_Volatile_Type;

```

Build output

```

Compile
  [Ada]          show_volatile_type.adb
Bind
  [gprbind]      show_volatile_type.bexch
  [Ada]          show_volatile_type.ali
Link
  [link]         show_volatile_type.adb

```

Runtime output

```
Val:  9.990000000000000E+05
```

Here, we're declaring a new type `Volatile_Long_Float` based on the `Long_Float` type and using the `Volatile` aspect. Any object of this type is automatically volatile.

In addition to that, we can declare components of an array to be volatile. In this case, we can use the `Volatile_Components` aspect in the array declaration. For example:

```
[Ada]
```

Listing 15: show_volatile_array_components.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Volatile_Array_Components is
4   Arr : array (1 .. 2) of Long_Float with Volatile_Components;
5 begin
6   Arr := (others => 0.0);
7
8   for I in 0 .. 999 loop
9     Arr (1) := Arr (1) + 2.0 * Long_Float (I);
10    Arr (2) := Arr (2) + 10.0 * Long_Float (I);
11  end loop;
12
13  Put_Line ("Arr (1): " & Long_Float'Image (Arr (1)));
14  Put_Line ("Arr (2): " & Long_Float'Image (Arr (2)));
15 end Show_Volatile_Array_Components;

```

Build output

```

Compile
  [Ada]          show_volatile_array_components.adb
Bind
  [gprbind]      show_volatile_array_components.bexch
  [Ada]          show_volatile_array_components.ali
Link
  [link]         show_volatile_array_components.adb

```

Runtime output

```

Arr (1):  9.990000000000000E+05
Arr (2):  4.995000000000000E+06

```

Note that it's possible to use the `Volatile` aspect for the array declaration as well:

```
[Ada]
```

```
Arr : array (1 .. 2) of Long_Float with Volatile;
```

4.5.2 Atomic

An atomic object is an object that only accepts atomic reads and updates. The Ada standard specifies that "for an atomic object (including an atomic component), all reads and updates of the object as a whole are indivisible." In this case, the compiler must generate Assembly code in such a way that reads and updates of an atomic object must be done in a single instruction, so that no other instruction could execute on that same object before the read or update completes.

In other contexts

Generally, we can say that operations are said to be atomic when they can be completed without interruptions. This is an important requirement when we're performing operations on objects in memory that are shared between multiple processes.

This definition of atomicity above is used, for example, when implementing databases. However, for this section, we're using the term "atomic" differently. Here, it really means that reads and updates must be performed with a single Assembly instruction.

For example, if we have a 32-bit object composed of four 8-bit bytes, the compiler cannot generate code to read or update the object using four 8-bit store / load instructions, or even two 16-bit store

/ load instructions. In this case, in order to maintain atomicity, the compiler must generate code using one 32-bit store / load instruction.

Because of this strict definition, we might have objects for which the `Atomic` aspect cannot be specified. Lots of machines support integer types that are larger than the native word-sized integer. For example, a 16-bit machine probably supports both 16-bit and 32-bit integers, but only 16-bit integer objects can be marked as atomic — or, more generally, only objects that fit into at most 16 bits.

Atomicity may be important, for example, when dealing with shared hardware registers. In fact, for certain architectures, the hardware may require that memory-mapped registers are handled atomically. In Ada, we can use the `Atomic` aspect to indicate that an object is atomic. This is how we can use the aspect to declare a shared hardware register:

[Ada]

Listing 16: `show_shared_hw_register.adb`

```
1 with System;
2
3 procedure Show_Shared_HW_Register is
4     R    : Integer
5         with Atomic, Address => System'To_Address (16#FFF00A0#);
6 begin
7     null;
8 end Show_Shared_HW_Register;
```

Note that the `Address` aspect allows for assigning a variable to a specific location in the memory. In this example, we're using this aspect to specify the address of the memory-mapped register. We'll discuss more about the `Address` aspect later in the section about *mapping structures to bit-fields* (page 137) (in chapter 6).

In addition to atomic objects, we can declare atomic types and atomic array components — similarly to what we've seen before for volatile objects. For example:

[Ada]

Listing 17: show_shared_hw_register.adb

```

1 with System;
2
3 procedure Show_Shared_HW_Register is
4   type Atomic_Integer is new Integer with Atomic;
5
6   R : Atomic_Integer with Address => System'To_Address (16#FFFF00A0#);
7
8   Arr : array (1 .. 2) of Integer with Atomic_Components;
9 begin
10  null;
11 end Show_Shared_HW_Register;

```

In this example, we're declaring the `Atomic_Integer` type, which is an atomic type. Objects of this type — such as `R` in this example — are automatically atomic. This example also includes the declaration of the `Arr` array, which has atomic components.

4.6 Interfacing with Devices

Previously, we've seen that we can use *representation clauses* (page 82) to specify a particular layout for a record type. As mentioned before, this is useful when interfacing with hardware, drivers, or communication protocols. In this section, we'll extend this concept for two specific use-cases: register overlays and data streams. Before we discuss those use-cases, though, we'll first explain the `Size` aspect and the `Size` attribute.

4.6.1 Size aspect and attribute

The `Size` aspect indicates the minimum number of bits required to represent an object. When applied to a type, the `Size` aspect is telling the compiler to not make record or array components of a type `T` any smaller than `X` bits. Therefore, a common usage for this aspect is to just confirm expectations: developers specify `'Size` to tell the compiler that `T` should fit `X` bits, and the compiler will tell them if they are right (or wrong).

When the specified size value is larger than necessary, it can cause objects to be bigger in memory than they would be otherwise. For example, for some enumeration types, we could say `for type Enum'Size use 32;` when the number of literals would otherwise have required only a byte. That's useful for unchecked conversions because the sizes of the two types need to be the same. Likewise, it's useful for interfacing with C, where enum types are just mapped to the `int` type, and thus larger than Ada might otherwise require. We'll discuss unchecked conversions *later in the course* (page 150).

Let's look at an example from an earlier chapter:

[Ada]

Listing 18: my_device_types.ads

```

1 package My_Device_Types is
2
3   type UInt10 is mod 2 ** 10
4     with Size => 10;
5
6 end My_Device_Types;

```

Here, we're saying that objects of type `UInt10` must have at least 10 bits. In this case, if the code compiles, it is a confirmation that such values can be represented in 10 bits when packed into an

enclosing record or array type.

If the size specified was larger than what the compiler would use by default, then it could affect the size of objects. For example, for `UInt10`, anything up to and including 16 would make no difference on a typical machine. However, anything over 16 would then push the compiler to use a larger object representation. That would be important for unchecked conversions, for example.

The `Size` attribute indicates the number of bits required to represent a type or an object. We can use the size attribute to retrieve the size of a type or of an object:

[Ada]

Listing 19: `show_device_types.adb`

```

1  with Ada.Text_IO;      use Ada.Text_IO;
2
3  with My_Device_Types; use My_Device_Types;
4
5  procedure Show_Device_Types is
6      UInt10_Obj : constant UInt10 := 0;
7  begin
8      Put_Line ("Size of UInt10 type:  " & Positive'Image (UInt10'Size));
9      Put_Line ("Size of UInt10 object: " & Positive'Image (UInt10_Obj'Size));
10 end Show_Device_Types;
```

Build output

```

Compile
  [Ada]      show_device_types.adb
  [Ada]      my_device_types.ads
Bind
  [gprbind]  show_device_types.bexch
  [Ada]      show_device_types.ali
Link
  [link]     show_device_types.adb
```

Runtime output

```

Size of UInt10 type:    10
Size of UInt10 object:  16
```

Here, we're retrieving the actual sizes of the `UInt10` type and an object of that type. Note that the sizes don't necessarily need to match. For example, although the size of `UInt10` type is expected to be 10 bits, the size of `UInt10_Obj` may be 16 bits, depending on the platform. Also, components of this type within composite types (arrays, records) will probably be 16 bits as well unless they are packed.

4.6.2 Register overlays

Register overlays make use of representation clauses to create a structure that facilitates manipulating bits from registers. Let's look at a simplified example of a power management controller containing registers such as a system clock enable register. Note that this example is based on an actual architecture:

[Ada]

Listing 20: `registers.ads`

```

1  with System;
2
3  package Registers is
```

(continues on next page)

(continued from previous page)

```

4  type Bit      is mod 2 ** 1
5      with Size => 1;
6  type UInt5    is mod 2 ** 5
7      with Size => 5;
8  type UInt10   is mod 2 ** 10
9      with Size => 10;
10
11
12  subtype USB_Clock_Enable is Bit;
13
14  -- System Clock Enable Register
15  type PMC_SCER_Register is record
16      -- Reserved bits
17      Reserved_0_4      : UInt5          := 16#0#;
18      -- Write-only. Enable USB FS Clock
19      USBCLK            : USB_Clock_Enable := 16#0#;
20      -- Reserved bits
21      Reserved_6_15    : UInt10         := 16#0#;
22  end record
23  with
24      Volatile,
25      Size      => 16,
26      Bit_Order => System.Low_Order_First;
27
28  for PMC_SCER_Register use record
29      Reserved_0_4      at 0 range 0 .. 4;
30      USBCLK            at 0 range 5 .. 5;
31      Reserved_6_15    at 0 range 6 .. 15;
32  end record;
33
34  -- Power Management Controller
35  type PMC_Peripheral is record
36      -- System Clock Enable Register
37      PMC_SCER          : aliased PMC_SCER_Register;
38      -- System Clock Disable Register
39      PMC_SCDR          : aliased PMC_SCER_Register;
40  end record
41  with Volatile;
42
43  for PMC_Peripheral use record
44      -- 16-bit register at byte 0
45      PMC_SCER          at 16#0# range 0 .. 15;
46      -- 16-bit register at byte 2
47      PMC_SCDR          at 16#2# range 0 .. 15;
48  end record;
49
50  -- Power Management Controller
51  PMC_Periph : aliased PMC_Peripheral
52      with Import, Address => System'To_Address (16#400E0600#);
53
54  end Registers;

```

First, we declare the system clock enable register — this is `PMC_SCER_Register` type in the code example. Most of the bits in that register are reserved. However, we're interested in bit #5, which is used to activate or deactivate the system clock. To achieve a correct representation of this bit, we do the following:

- We declare the `USBCLK` component of this record using the `USB_Clock_Enable` type, which has a size of one bit; and
- we use a representation clause to indicate that the `USBCLK` component is specifically at bit #5 of byte #0.

After declaring the system clock enable register and specifying its individual bits as components of a record type, we declare the power management controller type — `PMC_Peripheral` record type in the code example. Here, we declare two 16-bit registers as record components of `PMC_Peripheral`. These registers are used to enable or disable the system clock. The strategy we use in the declaration is similar to the one we've just seen above:

- We declare these registers as components of the `PMC_Peripheral` record type;
- we use a representation clause to specify that the `PMC_SCER` register is at byte #0 and the `PMC_SCDR` register is at byte #2.
 - Since these registers have 16 bits, we use a range of bits from 0 to 15.

The actual power management controller becomes accessible by the declaration of the `PMC_Periph` object of `PMC_Peripheral` type. Here, we specify the actual address of the memory-mapped registers (`400E0600` in hexadecimal) using the `Address` aspect in the declaration. When we use the `Address` aspect in an object declaration, we're indicating the address in memory of that object.

Because we specify the address of the memory-mapped registers in the declaration of `PMC_Periph`, this object is now an overlay for those registers. This also means that any operation on this object corresponds to an actual operation on the registers of the power management controller. We'll discuss more details about overlays in the section about *mapping structures to bit-fields* (page 137) (in chapter 6).

Finally, in a test application, we can access any bit of any register of the power management controller with simple record component selection. For example, we can set the `USBCLK` bit of the `PMC_SCER` register by using `PMC_Periph.PMC_SCER.USBCLK`:

[Ada]

Listing 21: `enable_usb_clock.adb`

```
1 with Registers;  
2  
3 procedure Enable_USB_Clock is  
4 begin  
5   Registers.PMC_Periph.PMC_SCER.USBCLK := 1;  
6 end Enable_USB_Clock;
```

This code example makes use of many aspects and keywords of the Ada language. One of them is the `Volatile` aspect, which we've discussed in the section about *volatile and atomic objects* (page 90). Using the `Volatile` aspect for the `PMC_SCER_Register` type ensures that objects of this type won't be stored in a register.

In the declaration of the `PMC_SCER_Register` record type of the example, we use the `Bit_Order` aspect to specify the bit ordering of the record type. Here, we can select one of these options:

- `High_Order_First`: first bit of the record is the most significant bit;
- `Low_Order_First`: first bit of the record is the least significant bit.

The declarations from the `Registers` package also makes use of the `Import`, which is sometimes necessary when creating overlays. When used in the context of object declarations, it avoids default initialization (for data types that have it.). Aspect `Import` will be discussed in the section that explains how to *map structures to bit-fields* (page 137) in chapter 6. Please refer to that chapter for more details.

Details about 'Size

In the example above, we're using the `Size` aspect in the declaration of the `PMC_SCER_Register` type. In this case, the effect is that it has the compiler confirm that the record type will fit into the expected 16 bits.

That's what the aspect does for type `PMC_SCER_Register` in the example above, as well as for the types `Bit`, `UInt5` and `UInt10`. For example, we may declare a stand-alone object of type `Bit`:

Listing 22: `show_bit_declaration.adb`

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Bit_Declaration is
4
5     type Bit is mod 2 ** 1
6         with Size => 1;
7
8     B : constant Bit := 0;
9     -- ^ Although Bit'Size is 1, B'Size is almost certainly 8
10 begin
11     Put_Line ("Bit'Size = " & Positive'Image (Bit'Size));
12     Put_Line ("B'Size = " & Positive'Image (B'Size));
13 end Show_Bit_Declaration;

```

Build output

Compile	
[Ada]	show_bit_declaration.adb
Bind	
[gprbind]	show_bit_declaration.bexch
[Ada]	show_bit_declaration.ali
Link	
[link]	show_bit_declaration.adb

Runtime output

```

Bit'Size = 1
B'Size   = 8

```

In this case, `B` is almost certainly going to be 8-bits wide on a typical machine, even though the language requires that `Bit'Size` is 1 by default.

In the declaration of the components of the `PMC_Peripheral` record type, we use the `aliased` keyword to specify that those record components are accessible via other paths besides the component name. Therefore, the compiler won't store them in registers. This makes sense because we want to ensure that we're accessing specific memory-mapped registers, and not registers assigned by the compiler. Note that, for the same reason, we also use the `aliased` keyword in the declaration of the `PMC_Periph` object.

4.6.3 Data streams

Creating data streams — in the context of interfacing with devices — means the serialization of arbitrary information and its transmission over a communication channel. For example, we might want to transmit the content of memory-mapped registers as byte streams using a serial port. To do this, we first need to get a serialized representation of those registers as an array of bytes, which we can then transmit over the serial port.

Serialization of arbitrary record types — including register overlays — can be achieved by declaring an array of bytes as an overlay. By doing this, we're basically interpreting the information from those record types as bytes while ignoring their actual structure — i.e. their components and representation clause. We'll discuss details about overlays in the section about *mapping structures to bit-fields* (page 137) (in chapter 6).

Let's look at a simple example of serialization of an arbitrary record type:

[Ada]

Listing 23: arbitrary_types.ads

```
1 package Arbitrary_Types is
2
3     type Arbitrary_Record is record
4         A : Integer;
5         B : Integer;
6         C : Integer;
7     end record;
8
9 end Arbitrary_Types;
```

Listing 24: serialize_data.ads

```
1 with Arbitrary_Types;
2
3 procedure Serialize_Data (Some_Object : Arbitrary_Types.Arbitrary_Record);
```

Listing 25: serialize_data.adb

```
1 with Arbitrary_Types;
2
3 procedure Serialize_Data (Some_Object : Arbitrary_Types.Arbitrary_Record) is
4     type UByte is new Natural range 0 .. 255
5         with Size => 8;
6
7     type UByte_Array is array (Positive range <>) of UByte;
8
9     --
10    -- We can access the serialized data in Raw_TX, which is our overlay
11    --
12    Raw_TX : UByte_Array (1 .. Some_Object'Size / 8)
13        with Address => Some_Object'Address;
14 begin
15     null;
16     --
17     -- Now, we could stream the data from Some_Object.
18     --
19     -- For example, we could send the bytes (from Raw_TX) via the
20     -- serial port.
21     --
22 end Serialize_Data;
```

Listing 26: data_stream_declaration.adb

```
1 with Arbitrary_Types;
2 with Serialize_Data;
3
4 procedure Data_Stream_Declaration is
5     Dummy_Object : Arbitrary_Types.Arbitrary_Record;
6
7 begin
8     Serialize_Data (Dummy_Object);
9 end Data_Stream_Declaration;
```

Build output

```
Compile
  [Ada]      data_stream_declaration.adb
  [Ada]      arbitrary_types.ads
```

(continues on next page)

(continued from previous page)

```

[Ada]          serialize_data.adb
Bind
  [gprbind]     data_stream_declaration.bexch
  [Ada]         data_stream_declaration.ali
Link
  [link]        data_stream_declaration.adb

```

The most important part of this example is the implementation of the `Serialize_Data` procedure, where we declare `Raw_TX` as an overlay for our arbitrary object (`Some_Object` of `Arbitrary_Record` type). In simple terms, by writing with `Address => Some_Object'Address;` in the declaration of `Raw_TX`, we're specifying that `Raw_TX` and `Some_Object` have the same address in memory. Here, we are:

- taking the address of `Some_Object` — using the `Address` attribute —, and then
- using it as the address of `Raw_TX` — which is specified with the `Address` aspect.

By doing this, we're essentially saying that both `Raw_TX` and `Some_Object` are different representations of the same object in memory.

Because the `Raw_TX` overlay is completely agnostic about the actual structure of the record type, the `Arbitrary_Record` type could really be anything. By declaring `Raw_TX`, we create an array of bytes that we can use to stream the information from `Some_Object`.

We can use this approach and create a data stream for the register overlay example that we've seen before. This is the corresponding implementation:

[Ada]

Listing 27: registers.ads

```

1  with System;
2
3  package Registers is
4
5      type Bit is mod 2 ** 1
6          with Size => 1;
7      type UInt5 is mod 2 ** 5
8          with Size => 5;
9      type UInt10 is mod 2 ** 10
10         with Size => 10;
11
12     subtype USB_Clock_Enable is Bit;
13
14     -- System Clock Register
15     type PMC_SCER_Register is record
16         -- Reserved bits
17         Reserved_0_4 : UInt5 := 16#0#;
18         -- Write-only. Enable USB FS Clock
19         USBCLK       : USB_Clock_Enable := 16#0#;
20         -- Reserved bits
21         Reserved_6_15 : UInt10 := 16#0#;
22     end record
23     with
24         Volatile,
25         Size      => 16,
26         Bit_Order => System.Low_Order_First;
27
28     for PMC_SCER_Register use record
29         Reserved_0_4 at 0 range 0 .. 4;
30         USBCLK       at 0 range 5 .. 5;
31         Reserved_6_15 at 0 range 6 .. 15;

```

(continues on next page)

(continued from previous page)

```

32  end record;
33
34  -- Power Management Controller
35  type PMC_Peripheral is record
36    -- System Clock Enable Register
37    PMC_SCER      : aliased PMC_SCER_Register;
38    -- System Clock Disable Register
39    PMC_SCDR      : aliased PMC_SCER_Register;
40  end record
41    with Volatile;
42
43  for PMC_Peripheral use record
44    -- 16-bit register at byte 0
45    PMC_SCER      at 16#0# range 0 .. 15;
46    -- 16-bit register at byte 2
47    PMC_SCDR      at 16#2# range 0 .. 15;
48  end record;
49
50  -- Power Management Controller
51  PMC_Periph : aliased PMC_Peripheral;
52  -- with Import, Address => System'To_Address (16#400E0600#);
53
54  end Registers;

```

Listing 28: serial_ports.ads

```

1  package Serial_Ports is
2
3    type UByte is new Natural range 0 .. 255
4      with Size => 8;
5
6    type UByte_Array is array (Positive range <>) of UByte;
7
8    type Serial_Port is null record;
9
10   procedure Read (Port : in out Serial_Port;
11                  Data : out UByte_Array);
12
13   procedure Write (Port : in out Serial_Port;
14                   Data : UByte_Array);
15
16 end Serial_Ports;

```

Listing 29: serial_ports.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  package body Serial_Ports is
4
5    procedure Display (Data : UByte_Array) is
6    begin
7      Put_Line ("---- Data ----");
8      for E of Data loop
9        Put_Line (UByte'Image (E));
10     end loop;
11     Put_Line ("-----");
12 end Display;
13
14 procedure Read (Port : in out Serial_Port;
15                Data : out UByte_Array) is

```

(continues on next page)

(continued from previous page)

```

16     pragma Unreferenced (Port);
17 begin
18     Put_Line ("Reading data...");
19     Data := (0, 0, 32, 0);
20 end Read;
21
22 procedure Write (Port : in out Serial_Port;
23                Data : UByte_Array) is
24     pragma Unreferenced (Port);
25 begin
26     Put_Line ("Writing data...");
27     Display (Data);
28 end Write;
29
30 end Serial_Ports;

```

Listing 30: data_stream.ads

```

1 with Serial_Ports; use Serial_Ports;
2 with Registers;    use Registers;
3
4 package Data_Stream is
5
6     procedure Send (Port : in out Serial_Port;
7                   PMC : PMC_Peripheral);
8
9     procedure Receive (Port : in out Serial_Port;
10                     PMC : out PMC_Peripheral);
11
12 end Data_Stream;

```

Listing 31: data_stream.adb

```

1 package body Data_Stream is
2
3     procedure Send (Port : in out Serial_Port;
4                   PMC : PMC_Peripheral)
5     is
6         Raw_TX : UByte_Array (1 .. PMC'Size / 8)
7         with Address => PMC'Address;
8     begin
9         Write (Port => Port,
10              Data => Raw_TX);
11     end Send;
12
13     procedure Receive (Port : in out Serial_Port;
14                     PMC : out PMC_Peripheral)
15     is
16         Raw_TX : UByte_Array (1 .. PMC'Size / 8)
17         with Address => PMC'Address;
18     begin
19         Read (Port => Port,
20              Data => Raw_TX);
21     end Receive;
22
23 end Data_Stream;

```

Listing 32: test_data_stream.adb

```

1  with Ada.Text_IO;
2
3  with Registers;
4  with Data_Stream;
5  with Serial_Ports;
6
7  procedure Test_Data_Stream is
8
9      procedure Display_Registers is
10         use Ada.Text_IO;
11         begin
12             Put_Line ("---- Registers ----");
13             Put_Line ("PMC_SCER.USBCLK: "
14                 & Registers.PMC_Periph.PMC_SCER.USBCLK'Image);
15             Put_Line ("PMC_SCDR.USBCLK: "
16                 & Registers.PMC_Periph.PMC_SCDR.USBCLK'Image);
17             Put_Line ("-----");
18         end Display_Registers;
19
20         Port : Serial_Ports.Serial_Port;
21     begin
22         Registers.PMC_Periph.PMC_SCER.USBCLK := 1;
23         Registers.PMC_Periph.PMC_SCDR.USBCLK := 1;
24
25         Display_Registers;
26
27         Data_Stream.Send (Port => Port,
28                         PMC => Registers.PMC_Periph);
29
30         Data_Stream.Receive (Port => Port,
31                             PMC => Registers.PMC_Periph);
32
33         Display_Registers;
34     end Test_Data_Stream;

```

Build output

```

Compile
  [Ada]          test_data_stream.adb
  [Ada]          data_stream.adb
  [Ada]          registers.ads
  [Ada]          serial_ports.adb
Bind
  [gprbind]      test_data_stream.bexch
  [Ada]          test_data_stream.ali
Link
  [link]         test_data_stream.adb

```

Runtime output

```

---- Registers ----
PMC_SCER.USBCLK: 1
PMC_SCDR.USBCLK: 1
-----
Writing data...
---- Data ----
32
0
32

```

(continues on next page)

(continued from previous page)

```

0
-----
Reading data...
---- Registers ----
PMC_SCER.USBCLK:  0
PMC_SCDR.USBCLK:  1
-----

```

In this example, we can find the overlay in the implementation of the Send and Receive procedures from the `Data_Stream` package. Because the overlay doesn't need to know the internals of the `PMC_Peripheral` type, we're declaring it in the same way as in the previous example (where we created an overlay for `Some_Object`). In this case, we're creating an overlay for the `PMC` parameter.

Note that, for this section, we're not really interested in the details about the serial port. Thus, package `Serial_Ports` in this example is just a stub. However, because the `Serial_Port` type in that package only *sees* arrays of bytes, after implementing an actual serial port interface for a specific device, we could create data streams for any type.

4.7 ARM and svd2ada

As we've seen in the previous section about *interfacing with devices* (page 95), Ada offers powerful features to describe low-level details about the hardware architecture without giving up its strong typing capabilities. However, it can be cumbersome to create a specification for all those low-level details when you have a complex architecture. Fortunately, for ARM Cortex-M devices, the GNAT toolchain offers an Ada binding generator called **svd2ada**, which takes CMSIS-SVD descriptions for those devices and creates Ada specifications that match the architecture. CMSIS-SVD description files are based on the Cortex Microcontroller Software Interface Standard (CMSIS), which is a hardware abstraction layer for ARM Cortex microcontrollers.

Please refer to the [svd2ada project page](https://github.com/AdaCore/svd2ada)¹⁶ for details about this tool.

¹⁶ <https://github.com/AdaCore/svd2ada>

ENHANCING VERIFICATION WITH SPARK AND ADA

5.1 Understanding Exceptions and Dynamic Checks

In Ada, several common programming errors that are not already detected at compile-time are detected instead at run-time, triggering "exceptions" that interrupt the normal flow of execution. For example, an exception is raised by an attempt to access an array component via an index that is out of bounds. This simple check precludes exploits based on buffer overflow. Several other cases also raise language-defined exceptions, such as scalar range constraint violations and null pointer dereferences. Developers may declare and raise their own application-specific exceptions too. (Exceptions are software artifacts, although an implementation may map hardware events to exceptions.)

Exceptions are raised during execution of what we will loosely define as a "frame." A frame is a language construct that has a call stack entry when called, for example a procedure or function body. There are a few other constructs that are also pertinent but this definition will suffice for now.

Frames have a sequence of statements implementing their functionality. They can also have optional "exception handlers" that specify the response when exceptions are "raised" by those statements. These exceptions could be raised directly within the statements, or indirectly via calls to other procedures and functions.

For example, the frame below is a procedure including three exceptions handlers:

Listing 1: p.adb

```
1 procedure P is
2 begin
3   Statements_That_Might_Raise_Exceptions;
4 exception
5   when A =>
6     Handle_A;
7   when B =>
8     Handle_B;
9   when C =>
10    Handle_C;
11 end P;
```

The three exception handlers each start with the word `when` (lines 5, 7, and 9). Next comes one or more exception identifiers, followed by the so-called "arrow." In Ada, the arrow always associates something on the left side with something on the right side. In this case, the left side is the exception name and the right side is the handler's code for that exception.

Each handler's code consists of an arbitrary sequence of statements, in this case specific procedures called in response to those specific exceptions. If exception A is raised we call procedure `Handle_A` (line 6), dedicated to doing the actual work of handling that exception. The other two exceptions are dealt with similarly, on lines 8 and 10.

Structurally, the exception handlers are grouped together and textually separated from the rest of the code in a frame. As a result, the sequence of statements representing the normal flow of execution is distinct from the section representing the error handling. The reserved word `exception` separates these two sections (line 4 above). This separation helps simplify the overall flow, increasing understandability. In particular, status result codes are not required so there is no mixture of error checking and normal processing. If no exception is raised the exception handler section is automatically skipped when the frame exits.

Note how the syntactic structure of the exception handling section resembles that of an Ada case statement. The resemblance is intentional, to suggest similar behavior. When something in the statements of the normal execution raises an exception, the corresponding exception handler for that specific exception is executed. After that, the routine completes. The handlers do not "fall through" to the handlers below. For example, if exception B is raised, procedure `Handle_B` is called but `Handle_C` is not called. There's no need for a `break` statement, just as there is no need for it in a case statement. (There's no `break` statement in Ada anyway.)

So far, we've seen a frame with three specific exceptions handled. What happens if a frame has no handler for the actual exception raised? In that case the run-time library code goes "looking" for one.

Specifically, the active exception is propagated up the dynamic call chain. At each point in the chain, normal execution in that caller is abandoned and the handlers are examined. If that caller has a handler for the exception, the handler is executed. That caller then returns normally to its caller and execution continues from there. Otherwise, propagation goes up one level in the call chain and the process repeats. The search continues until a matching handler is found or no callers remain. If a handler is never found the application terminates abnormally. If the search reaches the main procedure and it has a matching handler it will execute the handler, but, as always, the routine completes so once again the application terminates.

For a concrete example, consider the following:

Listing 2: arrays.ads

```
1 package Arrays is
2
3     type List is array (Natural range <>) of Integer;
4
5     function Value (A : List; X, Y : Integer) return Integer;
6
7 end Arrays;
```

Listing 3: arrays.adb

```
1 package body Arrays is
2
3     function Value (A : List; X, Y : Integer) return Integer is
4     begin
5         return A (X + Y * 10);
6     end Value;
7
8 end Arrays;
```

Listing 4: some_process.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Arrays;      use Arrays;
3
4 procedure Some_Process is
5     L : constant List (1 .. 100) := (others => 42);
6 begin
7     Put_Line (Integer'Image (Value (L, 1, 10)));
```

(continues on next page)

(continued from previous page)

```

8 exception
9   when Constraint_Error =>
10     Put_Line ("Constraint_Error caught in Some_Process");
11     Put_Line ("Some_Process completes normally");
12 end Some_Process;

```

Listing 5: main.adb

```

1 with Some_Process;
2 with Ada.Text_IO; use Ada.Text_IO;
3
4 procedure Main is
5 begin
6   Some_Process;
7   Put_Line ("Main completes normally");
8 end Main;

```

Procedure Main calls `Some_Process`, which in turn calls function `Value` (line 7). `Some_Process` declares the array object `L` of type `List` on line 5, with bounds 1 through 100. The call to `Value` has arguments, including variable `L`, leading to an attempt to access an array component via an out-of-bounds index ($1 + 10 * 10 = 101$, beyond the last index of `L`). This attempt will trigger an exception in `Value` prior to actually accessing the array object's memory. Function `Value` doesn't have any exception handlers so the exception is propagated up to the caller `Some_Process`. Procedure `Some_Process` has an exception handler for `Constraint_Error` and it so happens that `Constraint_Error` is the exception raised in this case. As a result, the code for that handler will be executed, printing some messages on the screen. Then procedure `Some_Process` will return to `Main` normally. `Main` then continues to execute normally after the call to `Some_Process` and prints its completion message.

If procedure `Some_Process` had also not had a handler for `Constraint_Error`, that procedure call would also have returned abnormally and the exception would have been propagated further up the call chain to procedure `Main`. Normal execution in `Main` would likewise be abandoned in search of a handler. But `Main` does not have any handlers so `Main` would have completed abnormally, immediately, without printing its closing message.

This semantic model is the same as with many other programming languages, in which the execution of a frame's sequence of statements is unavoidably abandoned when an exception becomes active. The model is a direct reaction to the use of status codes returned from functions as in C, where it is all too easy to forget (intentionally or otherwise) to check the status values returned. With the exception model errors cannot be ignored.

However, full exception propagation as described above is not the norm for embedded applications when the highest levels of integrity are required. The run-time library code implementing exception propagation can be rather complex and expensive to certify. Those problems apply to the application code too, because exception propagation is a form of control flow without any explicit construct in the source. Instead of the full exception model, designers of high-integrity applications often take alternative approaches.

One alternative consists of deactivating exceptions altogether, or more precisely, deactivating language-defined checks, which means that the compiler will not generate code checking for conditions giving rise to exceptions. Of course, this makes the code vulnerable to attacks, such as buffer overflow, unless otherwise verified (e.g. through static analysis). Deactivation can be applied at the unit level, through the `-gnatp` compiler switch, or locally within a unit via the pragma `Suppress`. (Refer to the [GNAT User's Guide for Native Platforms](https://docs.adacore.com/gnat_ugn-docs/html/gnat_ugn/gnat_ugn/building_executable_programs_with_gnat.html)¹⁷ for more details about the switch.)

For example, we can write the following. Note the pragma on line 4 of `arrays.adb` within function `Value`:

¹⁷ https://docs.adacore.com/gnat_ugn-docs/html/gnat_ugn/gnat_ugn/building_executable_programs_with_gnat.html

Listing 6: arrays.ads

```
1 package Arrays is
2
3     type List is array (Natural range <>) of Integer;
4
5     function Value (A : List; X, Y : Integer) return Integer;
6
7 end Arrays;
```

Listing 7: arrays.adb

```
1 package body Arrays is
2
3     function Value (A : List; X, Y : Integer) return Integer is
4         pragma Suppress (All_Checks);
5     begin
6         return A (X + Y * 10);
7     end Value;
8
9 end Arrays;
```

Listing 8: some_process.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Arrays;      use Arrays;
3
4 procedure Some_Process is
5     L : constant List (1 .. 100) := (others => 42);
6 begin
7     Put_Line (Integer'Image (Value (L, 1, 10)));
8 exception
9     when Constraint_Error =>
10         Put_Line ("FAILURE");
11 end Some_Process;
```

This placement of the pragma will only suppress checks in the function body. However, that is where the exception would otherwise have been raised, leading to incorrect and unpredictable execution. (Run the program more than once. If it prints the right answer (42), or even the same value each time, it's just a coincidence.) As you can see, suppressing checks negates the guarantee of errors being detected and addressed at run-time.

Another alternative is to leave checks enabled but not retain the dynamic call-chain propagation. There are a couple of approaches available in this alternative.

The first approach is for the run-time library to invoke a global "last chance handler" (LCH) when any exception is raised. Instead of the sequence of statements of an ordinary exception handler, the LCH is actually a procedure intended to perform "last-wishes" before the program terminates. No exception handlers are allowed. In this scheme "propagation" is simply a direct call to the LCH procedure. The default LCH implementation provided by GNAT does nothing other than loop infinitely. Users may define their own replacement implementation.

The availability of this approach depends on the run-time library. Typically, *Zero Footprint* and *Ravenscar SFP* run-times will provide this mechanism because they are intended for certification.

A user-defined LCH handler can be provided either in C or in Ada, with the following profiles:

[Ada]

```
procedure Last_Chance_Handler (Source_Location : System.Address; Line : Integer);
pragma Export (C,
```

(continues on next page)

(continued from previous page)

```
Last_Chance_Handler,
  "__gnat_last_chance_handler");
```

[C]

```
void __gnat_last_chance_handler (char *source_location,
                                int line);
```

We'll go into the details of the pragma Export in a further section on language interfacing. For now, just know that the symbol `__gnat_last_chance_handler` is what the run-time uses to branch immediately to the last-chance handler. Pragma Export associates that symbol with this replacement procedure so it will be invoked instead of the default routine. As a consequence, the actual procedure name in Ada is immaterial.

Here is an example implementation that simply blinks an LED forever on the target:

```
procedure Last_Chance_Handler (Msg : System.Address; Line : Integer) is
  pragma Unreferenced (Msg, Line);

  Next_Release : Time := Clock;
  Period       : constant Time_Span := Milliseconds (500);
begin
  Initialize_LEDs;
  All_LEDs_Off;

  loop
    Toggle (LCH_LED);
    Next_Release := Next_Release + Period;
    delay until Next_Release;
  end loop;
end Last_Chance_Handler;
```

The `LCH_LED` is a constant referencing the LED used by the last-chance handler, declared elsewhere. The infinite loop is necessary because a last-chance handler must never return to the caller (hence the term "last-chance"). The LED changes state every half-second.

Unlike the approach in which there is only the last-chance handler routine, the other approach allows exception handlers, but in a specific, restricted manner. Whenever an exception is raised, the only handler that can apply is a matching handler located in the same frame in which the exception is raised. Propagation in this context is simply an immediate branch instruction issued by the compiler, going directly to the matching handler's sequence of statements. If there is no matching local handler the last chance handler is invoked. For example consider the body of function `Value` in the body of package `Arrays`:

Listing 9: arrays.ads

```
1 package Arrays is
2
3   type List is array (Natural range <>) of Integer;
4
5   function Value (A : List; X, Y : Integer) return Integer;
6
7 end Arrays;
```

Listing 10: arrays.adb

```
1 package body Arrays is
2
3   function Value (A : List; X, Y : Integer) return Integer is
4   begin
```

(continues on next page)

(continued from previous page)

```

5     return A (X + Y * 10);
6 exception
7     when Constraint_Error =>
8         return 0;
9 end Value;
10
11 end Arrays;

```

Listing 11: some_process.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Arrays;      use Arrays;
3
4 procedure Some_Process is
5     L : constant List (1 .. 100) := (others => 42);
6 begin
7     Put_Line (Integer'Image (Value (L, 1, 10)));
8 exception
9     when Constraint_Error =>
10         Put_Line ("FAILURE");
11 end Some_Process;

```

In both procedure `Some_Process` and function `Value` we have an exception handler for `Constraint_Error`. In this example the exception is raised in `Value` because the index check fails there. A local handler for that exception is present so the handler applies and the function returns zero, normally. Because the call to the function returns normally, the execution of `Some_Process` prints zero and then completes normally.

Let's imagine, however, that function `Value` did *not* have a handler for `Constraint_Error`. In the context of full exception propagation, the function call would return to the caller, i.e., `Some_Process`, and would be handled in that procedure's handler. But only local handlers are allowed under the second alternative so the lack of a local handler in `Value` would result in the last-chance handler being invoked. The handler for `Constraint_Error` in `Some_Process` under this alternative approach.

So far we've only illustrated handling the `Constraint_Error` exception. It's possible to handle other language-defined and user-defined exceptions as well, of course. It is even possible to define a single handler for all other exceptions that might be encountered in the handled sequence of statements, beyond those explicitly named. The "name" for this otherwise anonymous exception is the Ada reserved word `others`. As in case statements, it covers all other choices not explicitly mentioned, and so must come last. For example:

Listing 12: arrays.ads

```

1 package Arrays is
2
3     type List is array (Natural range <>) of Integer;
4
5     function Value (A : List; X, Y : Integer) return Integer;
6
7 end Arrays;

```

Listing 13: arrays.adb

```

1 package body Arrays is
2
3     function Value (A : List; X, Y : Integer) return Integer is
4     begin
5         return A (X + Y * 10);
6     exception

```

(continues on next page)

(continued from previous page)

```

7      when Constraint_Error =>
8          return 0;
9      when others =>
10         return -1;
11     end Value;
12
13 end Arrays;

```

Listing 14: some_process.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2  with Arrays;      use Arrays;
3
4  procedure Some_Process is
5      L : constant List (1 .. 100) := (others => 42);
6  begin
7      Put_Line (Integer'Image (Value (L, 1, 10)));
8  exception
9      when Constraint_Error =>
10         Put_Line ("FAILURE");
11 end Some_Process;

```

In the code above, the `Value` function has a handler specifically for `Constraint_Error` as before, but also now has a handler for all other exceptions. For any exception other than `Constraint_Error`, function `Value` returns -1. If you remove the function's handler for `Constraint_Error` (lines 7 and 8) then the other "anonymous" handler will catch the exception and -1 will be returned instead of zero.

There are additional capabilities for exceptions, but for now you have a good basic understanding of how exceptions work, especially their dynamic nature at run-time.

5.2 Understanding Dynamic Checks versus Formal Proof

So far, we have discussed language-defined checks inserted by the compiler for verification at run-time, leading to exceptions being raised. We saw that these dynamic checks verified semantic conditions ensuring proper execution, such as preventing writing past the end of a buffer, or exceeding an application-specific integer range constraint, and so on. These checks are defined by the language because they apply generally and can be expressed in language-defined terms.

Developers can also define dynamic checks. These checks specify component-specific or application-specific conditions, expressed in terms defined by the component or application. We will refer to these checks as "user-defined" for convenience. (Be sure you understand that we are not talking about user-defined *exceptions* here.)

Like the language-defined checks, user-defined checks must be true at run-time. All checks consist of Boolean conditions, which is why we can refer to them as assertions: their conditions are asserted to be true by the compiler or developer.

Assertions come in several forms, some relatively low-level, such as a simple pragma `Assert`, and some high-level, such as type invariants and contracts. These forms will be presented in detail in a later section, but we will illustrate some of them here.

User-defined checks can be enabled at run-time in GNAT with the `-gnata` switch, as well as with pragma `Assertion_Policy`. The switch enables all forms of these assertions, whereas the pragma can be used to control specific forms. The switch is typically used but there are reasonable use-cases in which some user-defined checks are enabled, and others, although defined, are disabled.

By default in GNAT, language-defined checks are enabled but user-defined checks are disabled. Here's an example of a simple program employing a low-level assertion. We can use it to show the effects of the switches, including the defaults:

Listing 15: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   X : Positive := 10;
5 begin
6   X := X * 5;
7   pragma Assert (X > 99);
8   X := X - 99;
9   Put_Line (Integer'Image (X));
10 end Main;

```

If we compiled this code we would get a warning about the assignment on line 8 after the pragma Assert, but not one about the Assert itself on line 7.

```

gprbuild -q -P main.gpr
main.adb:8:11: warning: value not in range of type "Standard.Positive"
main.adb:8:11: warning: "Constraint_Error" will be raised at run time

```

No code is generated for the user-defined check expressed via pragma Assert but the language-defined check is emitted. In this case the range constraint on X excludes zero and negative numbers, but $X * 5 = 50$, $X - 99 = -49$. As a result, the check for the last assignment would fail, raising Constraint_Error when the program runs. These results are the expected behavior for the default switch settings.

But now let's enable user-defined checks and build it. Different compiler output will appear.

Listing 16: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   X : Positive := 10;
5 begin
6   X := X * 5;
7   pragma Assert (X > 99);
8   X := X - 99;
9   Put_Line (Integer'Image (X));
10 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
main.adb:7:19: warning: assertion will fail at run time [-gnatw.a]
main.adb:8:11: warning: value not in range of type "Standard.Positive" [enabled by_
↳default]
main.adb:8:11: warning: "Constraint_Error" will be raised at run time [enabled by_
↳default]
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```
raised ADA.ASSERTIONS.ASSERTION_ERROR : main.adb:7
```

Now we also get the compiler warning about the pragma Assert condition. When run, the failure of pragma Assert on line 7 raises the exception `Ada.Assertions.Assertion_Error`. According to the expression in the assertion, `X` is expected (incorrectly) to be above 99 after the multiplication. (The exception name in the error message, `SYSTEM.ASSERTIONS.ASSERT_FAILURE`, is a GNAT-specific alias for `Ada.Assertions.Assertion_Error`.)

It's interesting to see in the output that the compiler can detect some violations at compile-time:

```
main.adb:7:19: warning: assertion will fail at run time
main.adb:7:21: warning: condition can only be True if invalid values present
main.adb:8:11: warning: value not in range of type "Standard.Positive"
```

Generally speaking, a complete analysis is beyond the scope of compilers and they may not find all errors prior to execution, even those we might detect ourselves by inspection. More errors can be found by tools dedicated to that purpose, known as static analyzers. But even an automated static analysis tool cannot guarantee it will find all potential problems.

A much more powerful alternative is formal proof, a form of static analysis that can (when possible) give strong guarantees about the checks, for all possible conditions and all possible inputs. Proof can be applied to both language-defined and user-defined checks.

Be sure you understand that formal proof, as a form of static analysis, verifies conditions prior to execution, even prior to compilation. That earliness provides significant cost benefits. Removing bugs earlier is far less expensive than doing so later because the cost to fix bugs increases exponentially over the phases of the project life cycle, especially after deployment. Preventing bug introduction into the deployed system is the least expensive approach of all. Furthermore, cost savings during the initial development will be possible as well, for reasons specific to proof. We will revisit this topic later in this section.

Formal analysis for proof can be achieved through the SPARK subset of the Ada language combined with the **gnatprove** verification tool. SPARK is a subset encompassing most of the Ada language, except for features that preclude proof. As a disclaimer, this course is not aimed at providing a full introduction to proof and the SPARK language, but rather to present in a few examples what it is about and what it can do for us.

As it turns out, our procedure `Main` is already SPARK compliant so we can start verifying it.

Listing 17: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   X : Positive := 10;
5 begin
6   X := X * 5;
7   pragma Assert (X > 99);
8   X := X - 99;
9   Put_Line (Integer'Image (X));
10 end Main;
```

Build output

```
Compile
[Ada]          main.adb
main.adb:7:20: warning: assertion will fail at run time [-gnatw.a]
main.adb:8:12: warning: value not in range of type "Standard.Positive" [enabled by
↳default]
main.adb:8:12: warning: "Constraint_Error" will be raised at run time [enabled by
↳default]
```

(continues on next page)

(continued from previous page)

```

Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Prover output

```

Phase 1 of 2: generation of Global contracts ...
Phase 2 of 2: flow analysis and proof ...
main.adb:7:20: medium: assertion might fail
gnatprove: unproved check messages considered as errors

```

Runtime output

```
raised ADA.ASSERTIONS.ASSERTION_ERROR : main.adb:7
```

The "Prove" button invokes **gnatprove** on `main.adb`. You can ignore the parameters to the invocation. For the purpose of this demonstration, the interesting output is this message:

```
main.adb:7:19: medium: assertion might fail, cannot prove X > 99 (e.g. when X = 50)
```

gnatprove can tell that the assertion `X > 99` may have a problem. There's indeed a bug here, and **gnatprove** even gives us the counterexample (when `X` is 50). As a result the code is not proven and we know we have an error to correct.

Notice that the message says the assertion "might fail" even though clearly **gnatprove** has an example for when failure is certain. That wording is a reflection of the fact that SPARK gives strong guarantees when the assertions are proven to hold, but does not guarantee that flagged problems are indeed problems. In other words, **gnatprove** does not give false positives but false negatives are possible. The result is that if **gnatprove** does not indicate a problem for the code under analysis we can be sure there is no problem, but if **gnatprove** does indicate a problem the tool may be wrong.

5.3 Initialization and Correct Data Flow

An immediate benefit from having our code compatible with the SPARK subset is that we can ask **gnatprove** to verify initialization and correct data flow, as indicated by the absence of messages during SPARK "flow analysis." Flow analysis detects programming errors such as reading uninitialized data, problematic aliasing between formal parameters, and data races between concurrent tasks.

In addition, **gnatprove** checks unit specifications for the actual data read or written, and the flow of information from inputs to outputs. As you can imagine, this verification provides significant benefits, and it can be reached with comparatively low cost.

For example, the following illustrates an initialization failure:

Listing 18: `main.adb`

```

1 with Increment;
2 with Ada.Text_IO; use Ada.Text_IO;
3
4 procedure Main is
5   B : Integer;
6 begin
7   Increment (B);

```

(continues on next page)

(continued from previous page)

```

8   Put_Line (B'Image);
9 end Main;

```

Listing 19: increment.adb

```

1 procedure Increment (Value : in out Integer) is
2 begin
3   Value := Value + 1;
4 end Increment;

```

Prover output

```

Phase 1 of 2: generation of Global contracts ...
Phase 2 of 2: analysis of data and information flow ...

```

Granted, Increment is a silly procedure as-is, but imagine it did useful things, and, as part of that, incremented the argument. **gnatprove** tells us that the caller has not assigned a value to the argument passed to Increment.

Consider this next routine, which contains a serious coding error. Flow analysis will find it for us.

Listing 20: compute_offset.adb

```

1 with Ada.Numerics.Elementary_Functions; use Ada.Numerics.Elementary_Functions;
2
3 procedure Compute_Offset (K : Float; Z : out Integer; Flag : out Boolean) is
4   X : constant Float := Sin (K);
5 begin
6   if X < 0.0 then
7     Z := 0;
8     Flag := True;
9   elsif X > 0.0 then
10    Z := 1;
11    Flag := True;
12   else
13    Flag := False;
14   end if;
15 end Compute_Offset;

```

Prover output

```

Phase 1 of 2: generation of Global contracts ...
Phase 2 of 2: analysis of data and information flow ...
compute_offset.adb:3:38: medium: "Z" might not be initialized in "Compute_Offset"
gnatprove: unproved check messages considered as errors

```

gnatprove tells us that Z might not be initialized (assigned a value) in Compute_Offset, and indeed that is correct. Z is a mode out parameter so the routine should assign a value to it: Z is an output, after all. The fact that Compute_Offset does not do so is a significant and nasty bug. Why is it so nasty? In this case, formal parameter Z is of the scalar type Integer, and scalar parameters are always passed by copy in Ada and SPARK. That means that, when returning to the caller, an integer value is copied to the caller's argument passed to Z. But this procedure doesn't always assign the value to be copied back, and in that case an arbitrary value — whatever is on the stack — is copied to the caller's argument. The poor programmer must debug the code to find the problem, yet the effect could appear well downstream from the call to Compute_Offset. That's not only painful, it is expensive. Better to find the problem before we even compile the code.

5.4 Contract-Based Programming

So far, we've seen assertions in a routine's sequence of statements, either through implicit language-defined checks (is the index in the right range?) or explicit user-defined checks. These checks are already useful by themselves but they have an important limitation: the assertions are in the implementation, hidden from the callers of the routine. For example, a call's success or failure may depend upon certain input values but the caller doesn't have that information.

Generally speaking, Ada and SPARK put a lot of emphasis on strong, complete specifications for the sake of abstraction and analysis. Callers need not examine the implementations to determine whether the arguments passed to it are changed, for example. It is possible to go beyond that, however, to specify implementation constraints and functional requirements. We use contracts to do so.

At the language level, contracts are higher-level forms of assertions associated with specifications and declarations rather than sequences of statements. Like other assertions they can be activated or deactivated at run-time, and can be statically proven. We'll concentrate here on two kinds of contracts, both associated especially (but not exclusively) with procedures and functions:

- *Preconditions*, those Boolean conditions required to be true *prior* to a call of the corresponding subprogram
- *Postconditions*, those Boolean conditions required to be true *after* a call, as a result of the corresponding subprogram's execution

In particular, preconditions specify the initial conditions, if any, required for the called routine to correctly execute. Postconditions, on the other hand, specify what the called routine's execution must have done, at least, on normal completion. Therefore, preconditions are obligations on callers (referred to as "clients") and postconditions are obligations on implementers. By the same token, preconditions are guarantees to the implementers, and postconditions are guarantees to clients.

Contract-based programming, then, is the specification and rigorous enforcement of these obligations and guarantees. Enforcement is rigorous because it is not manual, but tool-based: dynamically at run-time with exceptions, or, with SPARK, statically, prior to build.

Preconditions are specified via the "Pre" aspect. Postconditions are specified via the "Post" aspect. Usually subprograms have separate declarations and these aspects appear with those declarations, even though they are *about* the bodies. Placement on the declarations allows the obligations and guarantees to be visible to all parties. For example:

Listing 21: mid.ads

```

1 function Mid (X, Y : Integer) return Integer with
2   Pre  => X + Y /= 0,
3   Post => Mid'Result > X;
```

The precondition on line 2 specifies that, for any given call, the sum of the values passed to parameters X and Y must not be zero. (Perhaps we're dividing by $X + Y$ in the body.) The declaration also provides a guarantee about the function call's result, via the postcondition on line 3: for any given call, the value returned will be greater than the value passed to X.

Consider a client calling this function:

Listing 22: demo.adb

```

1 with Mid;
2 with Ada.Text_IO; use Ada.Text_IO;
3
4 procedure Demo is
5   A, B, C : Integer;
6 begin
```

(continues on next page)

(continued from previous page)

```

7   A := Mid (1, 2);
8   B := Mid (1, -1);
9   C := Mid (A, B);
10  Put_Line (C'Image);
11 end Demo;

```

Prover output

```

Phase 1 of 2: generation of Global contracts ...
demo.adb:1:06: error: file "mid.ads" not found
gnatprove: error during generation of Global contracts

```

gnatprove indicates that the assignment to B (line 8) might fail because of the precondition, i.e., the sum of the inputs shouldn't be 0, yet $-1 + 1 = 0$. (We will address the other output message elsewhere.)

Let's change the argument passed to Y in the second call (line 8). Instead of -1 we will pass -2:

Listing 23: demo.adb

```

1  with Mid;
2  with Ada.Text_IO; use Ada.Text_IO;
3
4  procedure Demo is
5     A, B, C : Integer;
6  begin
7     A := Mid (1, 2);
8     B := Mid (1, -2);
9     C := Mid (A, B);
10    Put_Line (C'Image);
11 end Demo;

```

Listing 24: mid.ads

```

1  function Mid (X, Y : Integer) return Integer with
2     Pre  => X + Y /= 0,
3     Post => Mid'Result > X;

```

Prover output

```

Phase 1 of 2: generation of Global contracts ...
Phase 2 of 2: flow analysis and proof ...
warning: no bodies have been analyzed by GNATprove
enable analysis of a non-generic body using SPARK_Mode

```

The second call will no longer be flagged for the precondition. In addition, **gnatprove** will know from the postcondition that A has to be greater than 1, as does B, because in both calls 1 was passed to X. Therefore, **gnatprove** can deduce that the precondition will hold for the third call `C := Mid (A, B)`; because the sum of two numbers greater than 1 will never be zero.

Postconditions can also compare the state prior to a call with the state after a call, using the 'Old attribute. For example:

Listing 25: increment.ads

```

1  procedure Increment (Value : in out Integer) with
2     Pre  => Value < Integer'Last,
3     Post => Value = Value'Old + 1;

```

Listing 26: increment.adb

```
1 procedure Increment (Value : in out Integer) is
2 begin
3   Value := Value + 1;
4 end Increment;
```

Prover output

```
Phase 1 of 2: generation of Global contracts ...
Phase 2 of 2: flow analysis and proof ...
```

The postcondition specifies that, on return, the argument passed to the parameter `Value` will be one greater than it was immediately prior to the call (`Value'Old`).

5.5 Replacing Defensive Code

One typical benefit of contract-based programming is the removal of defensive code in subprogram implementations. For example, the *Push* operation for a stack type would need to ensure that the given stack is not already full. The body of the routine would first check that, explicitly, and perhaps raise an exception or set a status code. With preconditions we can make the requirement explicit and **gnatprove** will verify that the requirement holds at all call sites.

This reduction has a number of advantages:

- The implementation is simpler, removing validation code that is often difficult to test, makes the code more complex and leads to behaviors that are difficult to define.
- The precondition documents the conditions under which it's correct to call the subprogram, moving from an implementer responsibility to mitigate invalid input to a user responsibility to fulfill the expected interface.
- Provides the means to verify that this interface is properly respected, through code review, dynamic checking at run-time, or formal static proof.

As an example, consider a procedure `Read` that returns a component value from an array. Both the `Data` and `Index` are objects visible to the procedure so they are not formal parameters.

Listing 27: p.ads

```
1 package P is
2
3   type List is array (Integer range <>) of Character;
4
5   Data : List (1 .. 100);
6   Index : Integer := Data'First;
7
8   procedure Read (V : out Character);
9
10 end P;
```

Listing 28: p.adb

```
1 package body P is
2
3   procedure Read (V : out Character) is
4   begin
5     if Index not in Data'Range then
6       V := Character'First;
```

(continues on next page)

(continued from previous page)

```

7      return;
8  end if;
9
10     V := Data (Index);
11     Index := Index + 1;
12 end Read;
13 end P;
```

Prover output

```

Phase 1 of 2: generation of Global contracts ...
Phase 2 of 2: flow analysis and proof ...
```

In addition to procedure `Read` we would also have a way to load the array components in the first place, but we can ignore that for the purpose of this discussion.

Procedure `Read` is responsible for reading an element of the array and then incrementing the index. What should it do in case of an invalid index? In this implementation there is defensive code that returns a value arbitrarily chosen. We could also redesign the code to return a status in this case, or — better — raise an exception.

An even more robust approach would be instead to ensure that this subprogram is only called when `Index` is within the indexing boundaries of `Data`. We can express that requirement with a precondition (line 9).

Listing 29: p.ads

```

1 package P is
2
3     type List is array (Integer range <>) of Character;
4
5     Data : List (1 .. 100);
6     Index : Integer := 1;
7
8     procedure Read (V : out Character)
9         with Pre => Index in Data'Range;
10
11 end P;
```

Listing 30: p.adb

```

1 package body P is
2
3     procedure Read (V : out Character) is
4     begin
5         V := Data (Index);
6         Index := Index + 1;
7     end Read;
8
9 end P;
```

Prover output

```

Phase 1 of 2: generation of Global contracts ...
Phase 2 of 2: flow analysis and proof ...
```

Now we don't need the defensive code in the procedure body. That's safe because SPARK will attempt to prove statically that the check will not fail at the point of each call.

Assuming that procedure `Read` is intended to be the only way to get values from the array, in a real application (where the principles of software engineering apply) we would take advantage of the compile-time visibility controls that packages offer. Specifically, we would move all the variables'

declarations to the private part of the package, or even the package body, so that client code could not possibly access the array directly. Only procedure `Read` would remain visible to clients, thus remaining the only means of accessing the array. However, that change would entail others, and in this chapter we are only concerned with introducing the capabilities of SPARK. Therefore, we keep the examples as simple as possible.

5.6 Proving Absence of Run-Time Errors

Earlier we said that **gnatprove** will verify both language-defined and user-defined checks. Proving that the language-defined checks will not raise exceptions at run-time is known as proving "Absence of Run-Time Errors" or AoRTE for short. Successful proof of these checks is highly significant in itself.

One of the major resulting benefits is that we can deploy the final executable with checks disabled. That has obvious performance benefits, but it is also a safety issue. If we disable the checks we also disable the run-time library support for them, but in that case the language does not define what happens if indeed an exception is raised. Formally speaking, anything could happen. We must have good reason for thinking that exceptions cannot be raised.

This is such an important issue that proof of AoRTE can be used to comply with the objectives of certification standards in various high-integrity domains (for example, DO-178B/C in avionics, EN 50128 in railway, IEC 61508 in many safety-related industries, ECSS-Q-ST-80C in space, IEC 60880 in nuclear, IEC 62304 in medical, and ISO 26262 in automotive).

As a result, the quality of the program can be guaranteed to achieve higher levels of integrity than would be possible in other programming languages.

However, successful proof of AoRTE may require additional assertions, especially preconditions. We can see that with procedure `Increment`, the procedure that takes an `Integer` argument and increments it by one. But of course, if the incoming value of the argument is the largest possible positive value, the attempt to increment it would overflow, raising `Constraint_Error`. (As you have likely already concluded, `Constraint_Error` is the most common exception you will have to deal with.) We added a precondition to allow only the integer values up to, but not including, the largest positive value:

Listing 31: increment.ads

```
1 procedure Increment (Value : in out Integer) with
2   Pre  => Value < Integer'Last,
3   Post => Value = Value'Old + 1;
```

Listing 32: increment.adb

```
1 procedure Increment (Value : in out Integer) is
2 begin
3   Value := Value + 1;
4 end Increment;
```

Prover output

```
Phase 1 of 2: generation of Global contracts ...
Phase 2 of 2: flow analysis and proof ...
```

Prove it, then comment-out the precondition and try proving it again. Not only will **gnatprove** tell us what is wrong, it will suggest a solution as well.

Without the precondition the check it provides would have to be implemented as defensive code in the body. One or the other is critical here, but note that we should never need both.

5.7 Proving Abstract Properties

The postcondition on `Increment` expresses what is, in fact, a unit-level requirement. Successfully proving such requirements is another significant robustness and cost benefit. Together with the proofs for initialization and AoRTE, these proofs ensure program integrity, that is, the program executes within safe boundaries: the control flow of the program is correctly programmed and cannot be circumvented through run-time errors, and data cannot be corrupted.

We can go even further. We can use contracts to express arbitrary abstract properties when such exist. Safety and security properties, for instance, could be expressed as postconditions and then proven by **gnatprove**.

For example, imagine we have a procedure to move a train to a new position on the track, and we want to do so safely, without leading to a collision with another train. Procedure `Move`, therefore, takes two inputs: a train identifier specifying which train to move, and the intended new position. The procedure's output is a value indicating a motion command to be given to the train in order to go to that new position. If the train cannot go to that new position safely the output command is to stop the train. Otherwise the command is for the train to continue at an indicated speed:

```
type Move_Result is (Full_Speed, Slow_Down, Keep_Going, Stop);

procedure Move
  (Train      : in Train_Id;
   New_Position : in Train_Position;
   Result      : out Move_Result)
with
  Pre => Valid_Id (Train) and
        Valid_Move (Trains (Train), New_Position) and
        At_Most_One_Train_Per_Track and
        Safe_Signaling,
  Post => At_Most_One_Train_Per_Track and
        Safe_Signaling;

function At_Most_One_Train_Per_Track return Boolean;

function Safe_Signaling return Boolean;
```

The preconditions specify that, given a safe initial state and a valid move, the result of the call will also be a safe state: there will be at most one train per track section and the track signaling system will not allow any unsafe movements.

5.8 Final Comments

Make sure you understand that **gnatprove** does not attempt to prove the program correct as a whole. It attempts to prove language-defined and user-defined assertions about parts of the program, especially individual routines and calls to those routines. Furthermore, **gnatprove** proves the routines correct only to the extent that the user-defined assertions correctly and sufficiently describe and constrain the implementation of the corresponding routines.

Although we are not proving whole program correctness, as you will have seen — and done — we can prove properties that make our software far more robust and bug-free than is possible otherwise. But in addition, consider what proving the unit-level requirements for your procedures and functions would do for the cost of unit testing and system integration. The tests would pass the first time.

However, within the scope of what SPARK can do, not everything can be proven. In some cases that is because the software behavior is not amenable to expression as boolean conditions (for example, a mouse driver). In other cases the source code is beyond the capabilities of the analyzers

that actually do the mathematical proof. In these cases the combination of proof and actual test is appropriate, and still less expensive than testing alone.

There is, of course, much more to be said about what can be done with SPARK and **gnatprove**. Those topics are reserved for the [Introduction to SPARK¹⁸](#) course.

¹⁸ <https://learn.adacore.com/courses/intro-to-spark/index.html>

C TO ADA TRANSLATION PATTERNS

6.1 Naming conventions and casing considerations

One question that may arise relatively soon when converting from C to Ada is the style of source code presentation. The Ada language doesn't impose any particular style and for many reasons, it may seem attractive to keep a C-like style — for example, camel casing — to the Ada program.

However, the code in the Ada language standard, most third-party code, and the libraries provided by GNAT follow a specific style for identifiers and reserved words. Using a different style for the rest of the program leads to inconsistencies, thereby decreasing readability and confusing automatic style checkers. For those reasons, it's usually advisable to adopt the Ada style — in which each identifier starts with an upper case letter, followed by lower case letters (or digits), with an underscore separating two "distinct" words within the identifier. Acronyms within identifiers are in upper case. For example, there is a language-defined package named `Ada.Text_IO`. Reserved words are all lower case.

Following this scheme doesn't preclude adding additional, project-specific rules.

6.2 Manually interfacing C and Ada

Before even considering translating code from C to Ada, it's worthwhile to evaluate the possibility of keeping a portion of the C code intact, and only translating selected modules to Ada. This is a necessary evil when introducing Ada to an existing large C codebase, where re-writing the entire code upfront is not practical nor cost-effective.

Fortunately, Ada has a dedicated set of features for interfacing with other languages. The `Interfaces` package hierarchy and the pragmas `Convention`, `Import`, and `Export` allow you to make inter-language calls while observing proper data representation for each language.

Let's start with the following C code:

[C]

Listing 1: `call.c`

```
1  #include <stdio.h>
2
3  struct my_struct {
4      int A, B;
5  };
6
7  void call (struct my_struct *p) {
8      printf ("%d", p->A);
9  }
```

To call that function from Ada, the Ada compiler requires a description of the data structure to pass as well as a description of the function itself. To capture how the C struct `my_struct` is represented, we can use the following record along with a pragma `Convention`. The pragma directs the compiler to lay out the data in memory the way a C compiler would.

[Ada]

Listing 2: use_my_struct.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Interfaces.C;
3
4 procedure Use_My_Struct is
5
6     type my_struct is record
7         A : Interfaces.C.int;
8         B : Interfaces.C.int;
9     end record;
10    pragma Convention (C, my_struct);
11
12    V : my_struct := (A => 1, B => 2);
13 begin
14     Put_Line ("V = ("
15              & Interfaces.C.int'Image (V.A)
16              & Interfaces.C.int'Image (V.B)
17              & ")");
18 end Use_My_Struct;
```

Build output

```

Compile
  [Ada]          use_my_struct.adb
Bind
  [gprbind]      use_my_struct.bexch
  [Ada]          use_my_struct.ali
Link
  [link]         use_my_struct.adb
```

Runtime output

```
V = ( 1 2)
```

Describing a foreign subprogram call to Ada code is called *binding* and it is performed in two stages. First, an Ada subprogram specification equivalent to the C function is coded. A C function returning a value maps to an Ada function, and a void function maps to an Ada procedure. Then, rather than implementing the subprogram using Ada code, we use a pragma `Import`:

```

procedure Call (V : my_struct);
pragma Import (C, Call, "call"); -- Third argument optional
```

The `Import` pragma specifies that whenever `Call` is invoked by Ada code, it should invoke the `Call` function with the C calling convention.

And that's all that's necessary. Here's an example of a call to `Call`:

[Ada]

Listing 3: use_my_struct.adb

```

1 with Interfaces.C;
2
3 procedure Use_My_Struct is
4
```

(continues on next page)

(continued from previous page)

```

5  type my_struct is record
6      A : Interfaces.C.int;
7      B : Interfaces.C.int;
8  end record;
9  pragma Convention (C, my_struct);
10
11 procedure Call (V : my_struct);
12 pragma Import (C, Call, "call"); -- Third argument optional
13
14 V : my_struct := (A => 1, B => 2);
15 begin
16     Call (V);
17 end Use_My_Struct;

```

6.3 Building and Debugging mixed language code

The easiest way to build an application using mixed C / Ada code is to create a simple project file for **gprbuild** and specify C as an additional language. By default, when using **gprbuild** we only compile Ada source files. To compile C code files as well, we use the `Languages` attribute and specify `c` as an option, as in the following example of a project file named *default.gpr*:

```

project Default is

    for Languages use ("ada", "c");
    for Main use ("main.adb");

end Default;

```

Then, we use this project file to build the application by simply calling **gprbuild**. Alternatively, we can specify the project file on the command-line with the `-P` option — for example, `gprbuild -P default.gpr`. In both cases, **gprbuild** compiles all C source-code file found in the directory and links the corresponding object files to build the executable.

In order to include debug information, you can use `gprbuild -cargs -g`. This option adds debug information based on both C and Ada code to the executable. Alternatively, you can specify a `Builder` package in the project file and include global compilation switches for each language using the `Global_Compilation_Switches` attribute. For example:

```

project Default is

    for Languages use ("ada", "c");
    for Main use ("main.adb");

    package Builder is
        for Global_Compilation_Switches ("Ada") use ("-g");
        for Global_Compilation_Switches ("C") use ("-g");
    end Builder;

end Default;

```

In this case, you can simply run `gprbuild -P default.gpr` to build the executable.

To debug the executable, you can use programs such as **gdb** or **ddd**, which are suitable for debugging both C and Ada source-code. If you prefer a complete IDE, you may want to look into **GNAT Studio**, which supports building and debugging an application within a single environment, and remotely running applications loaded to various embedded devices. You can find more information about **gprbuild** and **GNAT Studio** in the [Introduction to GNAT Toolchain](https://learn.adacore.com/courses/GNAT_Toolchain_Intro/index.html)¹⁹ course.

¹⁹ https://learn.adacore.com/courses/GNAT_Toolchain_Intro/index.html

6.4 Automatic interfacing

It may be useful to start interfacing Ada and C by using automatic binding generators. These can be done either by invoking **gcc -fdump-ada-spec** option (to generate an Ada binding to a C header file) or **-gnatceg** option (to generate a C binding to an Ada specification file). For example:

```
gcc -c -fdump-ada-spec my_header.h
gcc -c -gnatceg spec.ads
```

The level of interfacing is very low level and typically requires either massaging (changing the generated files) or wrapping (calling the generated files from a higher level interface). For example, numbers bound from C to Ada are only standard numbers where user-defined types may be desirable. C uses a lot of by-pointer parameters which may be better replaced by other parameter modes, etc.

However, the automatic binding generator helps having a starting point which ensures compatibility of the Ada and the C code.

6.5 Using Arrays in C interfaces

It is relatively straightforward to pass an array from Ada to C. In particular, with the GNAT compiler, passing an array is equivalent to passing a pointer to its first element. Of course, as there's no notion of boundaries in C, the length of the array needs to be passed explicitly. For example:

[C]

Listing 4: p.h

```
1 void p (int * a, int length);
```

[Ada]

Listing 5: main.adb

```
1 procedure Main is
2   type Arr is array (Integer range <>) of Integer;
3
4   procedure P (V : Arr; Length : Integer);
5   pragma Import (C, P);
6
7   X : Arr (5 .. 15);
8 begin
9   P (X, X'Length);
10 end Main;
```

The other way around — that is, retrieving an array that has been creating on the C side — is more difficult. Because C doesn't explicitly carry boundaries, they need to be recreated in some way.

The first option is to actually create an Ada array without boundaries. This is the most flexible, but also the least safe option. It involves creating an array with indices over the full range of Integer without ever creating it from Ada, but instead retrieving it as an access from C. For example:

[C]

Listing 6: f.h

```
1 int * f ();
```

[Ada]

Listing 7: main.adb

```

1 procedure Main is
2   type Arr is array (Integer) of Integer;
3   type Arr_A is access all Arr;
4
5   function F return Arr_A;
6   pragma Import (C, F);
7 begin
8   null;
9 end Main;

```

Note that Arr is a constrained type (it doesn't have the range <> notation for indices). For that reason, as it would be for C, it's possible to iterate over the whole range of integer, beyond the memory actually allocated for the array.

A somewhat safer way is to overlay an Ada array over the C one. This requires having access to the length of the array. This time, let's consider two cases, one with an array and its size accessible through functions, another one on global variables. This time, as we're using an overlay, the function will be directly mapped to an Ada function returning an address:

[C]

Listing 8: fg.h

```

1 int * f_arr (void);
2 int f_size (void);
3
4 int * g_arr;
5 int g_size;

```

[Ada]

Listing 9: fg.ads

```

1 with System;
2
3 package Fg is
4
5   type Arr is array (Integer range <>) of Integer;
6
7   function F_Arr return System.Address;
8   pragma Import (C, F_Arr, "f_arr");
9
10  function F_Size return Integer;
11  pragma Import (C, F_Size, "f_size");
12
13  F : Arr (0 .. F_Size - 1) with Address => F_Arr;
14
15  G_Size : Integer;
16  pragma Import (C, G_Size, "g_size");
17
18  G_Arr : Arr (0 .. G_Size - 1);
19  pragma Import (C, G_Arr, "g_arr");
20
21 end Fg;

```

Listing 10: main.adb

```

1 with Fg;
2

```

(continues on next page)

(continued from previous page)

```
3 procedure Main is
4 begin
5     null;
6 end Main;
```

With all solutions though, importing an array from C is a relatively unsafe pattern, as there's only so much information on the array as there would be on the C side in the first place. These are good places for careful peer reviews.

6.6 By-value vs. by-reference types

When interfacing Ada and C, the rules of parameter passing are a bit different with regards to what's a reference and what's a copy. Scalar types and pointers are passed by value, whereas record and arrays are (almost) always passed by reference. However, there may be cases where the C interface also passes values and not pointers to objects. Here's a slightly modified version of a previous example to illustrate this point:

[C]

Listing 11: call.c

```
1 #include <stdio.h>
2
3 struct my_struct {
4     int A, B;
5 };
6
7 void call (struct my_struct p) {
8     printf ("%d", p.A);
9 }
```

In Ada, a type can be modified so that parameters of this type can always be passed by copy.

[Ada]

Listing 12: main.adb

```
1 with Interfaces.C;
2
3 procedure Main is
4     type my_struct is record
5         A : Interfaces.C.int;
6         B : Interfaces.C.int;
7     end record
8     with Convention => C_Pass_By_Copy;
9
10    procedure Call (V : my_struct);
11    pragma Import (C, Call, "call");
12 begin
13     null;
14 end Main;
```

Note that this cannot be done at the subprogram declaration level, so if there is a mix of by-copy and by-reference calls, two different types need to be used on the Ada side.

6.7 Naming and prefixes

Because of the absence of namespaces, any global name in C tends to be very long. And because of the absence of overloading, they can even encode type names in their type.

In Ada, the package is a namespace — two entities declared in two different packages are clearly identified and can always be specifically designated. The C names are usually a good indication of the names of the future packages and should be stripped — it is possible to use the full name if useful. For example, here's how the following declaration and call could be translated:

[C]

Listing 13: reg_interface.h

```
1 void registerInterface_Initialize (int size);
```

Listing 14: reg_interface_test.c

```
1 #include "reg_interface.h"
2
3 int main(int argc, const char * argv[])
4 {
5     registerInterface_Initialize(15);
6
7     return 0;
8 }
```

[Ada]

Listing 15: register_interface.ads

```
1 package Register_Interface is
2     procedure Initialize (Size : Integer)
3         with Import      => True,
4             Convention => C,
5             External_Name => "registerInterface_Initialize";
6
7 end Register_Interface;
```

Listing 16: main.adb

```
1 with Register_Interface;  
2  
3 procedure Main is  
4 begin  
5   Register_Interface.Initialize (15);  
6 end Main;
```

Note that in the above example, a use clause on `Register_Interface` could allow us to omit the prefix.

6.8 Pointers

The first thing to ask when translating pointers from C to Ada is: are they needed in the first place? In Ada, pointers (or access types) should only be used with complex structures that cannot be allocated at run-time — think of a linked list or a graph for example. There are many other situations that would need a pointer in C, but do not in Ada, in particular:

- Arrays, even when dynamically allocated
- Results of functions
- Passing large structures as parameters
- Access to registers
- ... others

This is not to say that pointers aren't used in these cases but, more often than not, the pointer is hidden from the user and automatically handled by the code generated by the compiler; thus avoiding possible mistakes from being made. Generally speaking, when looking at C code, it's good practice to start by analyzing how many pointers are used and to translate as many as possible into *pointerless* Ada structures.

Here are a few examples of such patterns — additional examples can be found throughout this document.

Dynamically allocated arrays can be directly allocated on the stack:

[C]

Listing 17: array_decl.c

```
1 #include <stdlib.h>  
2  
3 int main() {  
4   int *a = malloc(sizeof(int) * 10);  
5  
6   return 0;  
7 }
```

[Ada]

Listing 18: main.adb

```
1 procedure Main is  
2   type Arr is array (Integer range <>) of Integer;  
3   A : Arr (0 .. 9);  
4 begin
```

(continues on next page)

(continued from previous page)

```

5   null;
6 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
main.adb:3:04: warning: variable "A" is never read and never assigned [-gnatwv]
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

It's even possible to create a such an array within a structure, provided that the size of the array is known when instantiating this object, using a type discriminant:

[C]

Listing 19: array_decl.c

```

1  #include <stdlib.h>
2
3  typedef struct {
4      int * a;
5  } S;
6
7  int main(int argc, const char * argv[])
8  {
9      S v;
10
11     v.a = malloc(sizeof(int) * 10);
12
13     return 0;
14 }

```

[Ada]

Listing 20: main.adb

```

1  procedure Main is
2      type Arr is array (Integer range <>) of Integer;
3
4      type S (Last : Integer) is record
5          A : Arr (0 .. Last);
6      end record;
7
8      V : S (9);
9  begin
10     null;
11 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
main.adb:8:04: warning: variable "V" is never read and never assigned [-gnatwv]
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

With regards to parameter passing, usage mode (input / output) should be preferred to implementation mode (by copy or by reference). The Ada compiler will automatically pass a reference when needed. This works also for smaller objects, so that the compiler will copy in and out when needed. One of the advantages of this approach is that it clarifies the nature of the object: in particular, it differentiates between arrays and scalars. For example:

[C]

Listing 21: p.h

```
1 void p (int * a, int * b);
```

[Ada]

Listing 22: array_types.ads

```
1 package Array_Types is
2     type Arr is array (Integer range <>) of Integer;
3
4     procedure P (A : in out Integer; B : in out Arr);
5 end Array_Types;
```

Most of the time, access to registers end up in some specific structures being mapped onto a specific location in memory. In Ada, this can be achieved through an Address clause associated to a variable, for example:

[C]

Listing 23: test_c.c

```
1 int main(int argc, const char * argv[])
2 {
3     int * r = (int *)0xFFFF00A0;
4
5     return 0;
6 }
```

[Ada]

Listing 24: test.adb

```
1 with System;
2
3 procedure Test is
4     R : Integer with Address => System'To_Address (16#FFFF00A0#);
5 begin
6     null;
7 end Test;
```

Build output

Compile	
[Ada]	test.adb
Bind	
[grbind]	test.bexch
[Ada]	test.ali
Link	
[link]	test.adb

These are some of the most common misuse of pointers in Ada. Previous sections of the document deal with specifically using access types if absolutely necessary.

6.9 Bitwise Operations

Bitwise operations such as masks and shifts in Ada should be relatively rarely needed, and, when translating C code, it's good practice to consider alternatives. In a lot of cases, these operations are used to insert several pieces of data into a larger structure. In Ada, this can be done by describing the structure layout at the type level through representation clauses, and then accessing this structure as any other.

Consider the case of using a C primitive type as a container for single bit boolean flags. In C, this would be done through masks, e.g.:

[C]

Listing 25: flags.c

```

1  #define FLAG_1 0b0001
2  #define FLAG_2 0b0010
3  #define FLAG_3 0b0100
4  #define FLAG_4 0b1000
5
6  int main(int argc, const char * argv[])
7  {
8      int value = 0;
9
10     value |= FLAG_2 | FLAG_4;
11
12     return 0;
13 }
```

In Ada, the above can be represented through a Boolean array of enumerate values:

[Ada]

Listing 26: main.adb

```

1  procedure Main is
2      type Values is (Flag_1, Flag_2, Flag_3, Flag_4);
3      type Value_Array is array (Values) of Boolean
4          with Pack;
5
6      Value : Value_Array :=
7          (Flag_2 => True,
8           Flag_4 => True,
9           others => False);
10  begin
11      null;
12  end Main;
```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Note the Pack directive for the array, which requests that the array takes as little space as possible.

It is also possible to map records on memory when additional control over the representation is needed or more complex data are used:

[C]

Listing 27: struct_map.c

```
1 int main(int argc, const char * argv[])
2 {
3     int value = 0;
4
5     value = (2 << 1) | 1;
6
7     return 0;
8 }
```

[Ada]

Listing 28: main.adb

```
1 procedure Main is
2     type Value_Rec is record
3         V1 : Boolean;
4         V2 : Integer range 0 .. 3;
5     end record;
6
7     for Value_Rec use record
8         V1 at 0 range 0 .. 0;
9         V2 at 0 range 1 .. 2;
10    end record;
11
12    Value : Value_Rec := (V1 => True, V2 => 2);
13 begin
14     null;
15 end Main;
```

Build output

```
Compile
[Ada]      main.adb
Bind
[gnatbind] main.bexch
[Ada]      main.ali
Link
[link]     main.adb
```

The benefit of using Ada structure instead of bitwise operations is threefold:

- The code is simpler to read / write and less error-prone
- Individual fields are named
- The compiler can run consistency checks (for example, check that the value indeed fit in the expected size).

Note that, in cases where bitwise operators are needed, Ada provides modular types with and, or and xor operators. Further shift operators can also be provided upon request through a pragma. So the above could also be literally translated to:

[C]

Listing 29: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4     type Value_Type is mod 2 ** 32;
```

(continues on next page)

(continued from previous page)

```

5  pragma Provide_Shift_Operators (Value_Type);
6
7  Value : Value_Type;
8  begin
9  Value := Shift_Left (2, 1) or 1;
10 Put_Line ("Value = " & Value_Type'Image (Value));
11 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```
Value = 5
```

6.10 Mapping Structures to Bit-Fields

In the previous section, we've seen how to perform bitwise operations. In this section, we look at how to interpret a data type as a bit-field and perform low-level operations on it.

In general, you can create a bit-field from any arbitrary data type. First, we declare a bit-field type like this:

[Ada]

```
type Bit_Field is array (Natural range <>) of Boolean with Pack;
```

As we've seen previously, the `Pack` aspect declared at the end of the type declaration indicates that the compiler should optimize for size. We must use this aspect to be able to interpret data types as a bit-field.

Then, we can use the `Size` and the `Address` attributes of an object of any type to declare a bit-field for this object. We've discussed the `Size` attribute *earlier in this course* (page 95).

The `Address` attribute indicates the address in memory of that object. For example, assuming we've declare a variable `V`, we can declare an actual bit-field object by referring to the `Address` attribute of `V` and using it in the declaration of the bit-field, as shown here:

[Ada]

```
B : Bit_Field (0 .. V'Size - 1) with Address => V'Address;
```

Note that, in this declaration, we're using the `Address` attribute of `V` for the `Address` aspect of `B`.

This technique is called overlays for serialization. Now, any operation that we perform on `B` will have a direct impact on `V`, since both are using the same memory location.

The approach that we use in this section relies on the `Address` aspect. Another approach would be to use unchecked conversions, which we'll discuss in the *next section* (page 150).

We should add the `Volatile` aspect to the declaration to cover the case when both objects can still be changed independently — they need to be volatile, otherwise one change might be missed. This is the updated declaration:

[Ada]

```
B : Bit_Field (0 .. V'Size - 1) with Address => V'Address, Volatile;
```

Using the Volatile aspect is important at high level of optimizations. You can find further details about this aspect in the section about the *Volatile and Atomic aspects* (page 90).

Another important aspect that should be added is Import. When used in the context of object declarations, it'll avoid default initialization which could overwrite the existing content while creating the overlay — see an example in the admonition below. The declaration now becomes:

```
B : Bit_Field (0 .. V'Size - 1)
  with
    Address => V'Address, Import, Volatile;
```

Let's look at a simple example:

[Ada]

Listing 30: simple_bitfield.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Simple_Bitfield is
4   type Bit_Field is array (Natural range <>) of Boolean with Pack;
5
6   V : Integer := 0;
7   B : Bit_Field (0 .. V'Size - 1)
8     with Address => V'Address, Import, Volatile;
9 begin
10  B (2) := True;
11  Put_Line ("V = " & Integer'Image (V));
12 end Simple_Bitfield;
```

Build output

Compile	
[Ada]	simple_bitfield.adb
Bind	
[gprbind]	simple_bitfield.bexch
[Ada]	simple_bitfield.ali
Link	
[link]	simple_bitfield.adb

Runtime output

```
V = 4
```

In this example, we first initialize V with zero. Then, we use the bit-field B and set the third element (B (2)) to True. This automatically sets bit #3 of V to 1. Therefore, as expected, the application displays the message V = 4, which corresponds to $2^2 = 4$.

Note that, in the declaration of the bit-field type above, we could also have used a positive range. For example:

```
type Bit_Field is array (Positive range <>) of Boolean with Pack;

B : Bit_Field (1 .. V'Size)
  with Address => V'Address, Import, Volatile;
```

The only difference in this case is that the first bit is B (1) instead of B (0).

In C, we would rely on bit-shifting and masking to set that specific bit:

[C]

Listing 31: bitfield.c

```

1  #include <stdio.h>
2
3  int main(int argc, const char * argv[])
4  {
5      int v = 0;
6
7      v = v | (1 << 2);
8
9      printf("v = %d\n", v);
10
11     return 0;
12 }

```

Runtime output

```
v = 4
```

Important

Ada has the concept of default initialization. For example, you may set the default value of record components:

[Ada]

Listing 32: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4
5      type Rec is record
6          X : Integer := 10;
7          Y : Integer := 11;
8      end record;
9
10     R : Rec;
11 begin
12     Put_Line ("R.X = " & Integer'Image (R.X));
13     Put_Line ("R.Y = " & Integer'Image (R.Y));
14 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```

R.X = 10
R.Y = 11

```

In the code above, we don't explicitly initialize the components of R, so they still have the default values 10 and 11, which are displayed by the application.

Likewise, the `Default_Value` aspect can be used to specify the default value in other kinds of type declarations. For example:

[Ada]

Listing 33: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4
5     type Percentage is range 0 .. 100
6       with Default_Value => 10;
7
8     P : Percentage;
9 begin
10    Put_Line ("P = " & Percentage'Image (P));
11 end Main;
```

Build output

```
Compile
  [Ada]      main.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb
```

Runtime output

```
P = 10
```

When declaring an object whose type has a default value, the object will automatically be initialized with the default value. In the example above, `P` is automatically initialized with 10, which is the default value of the `Percentage` type.

Some types have an implicit default value. For example, access types have a default value of `null`.

As we've just seen, when declaring objects for types with associated default values, automatic initialization will happen. This can also happen when creating an overlay with the `Address` aspect. The default value is then used to overwrite the content at the memory location indicated by the address. However, in most situations, this isn't the behavior we expect, since overlays are usually created to analyze and manipulate existing values. Let's look at an example where this happens:

[Ada]

Listing 34: p.ads

```
1 package P is
2
3     type Unsigned_8 is mod 2 ** 8 with Default_Value => 0;
4
5     type Byte_Field is array (Natural range <>) of Unsigned_8;
6
7     procedure Display_Bytes_Increment (V : in out Integer);
8 end P;
```

Listing 35: p.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body P is
```

(continues on next page)

(continued from previous page)

```

4
5  procedure Display_Bytes_Increment (V : in out Integer) is
6      BF : Byte_Field (1 .. V'Size / 8)
7          with Address => V'Address, Volatile;
8  begin
9      for B of BF loop
10         Put_Line ("Byte = " & Unsigned_8'Image (B));
11     end loop;
12     Put_Line ("Now incrementing...");
13     V := V + 1;
14 end Display_Bytes_Increment;
15
16 end P;
```

Listing 36: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  with P; use P;
4
5  procedure Main is
6      V : Integer := 10;
7  begin
8      Put_Line ("V = " & Integer'Image (V));
9      Display_Bytes_Increment (V);
10     Put_Line ("V = " & Integer'Image (V));
11 end Main;
```

Build output

```

Compile
  [Ada]          main.adb
  [Ada]          p.adb
p.adb:7:14: warning: default initialization of "Bf" may modify "V" [enabled by _
↳default]
p.adb:7:14: warning: use pragma Import for "Bf" to suppress initialization (RM B.
↳1(24)) [enabled by default]
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```

V = 10
Byte = 0
Byte = 0
Byte = 0
Byte = 0
Now incrementing...
V = 1
```

In this example, we expect `Display_Bytes_Increment` to display each byte of the `V` parameter and then increment it by one. Initially, `V` is set to 10, and the call to `Display_Bytes_Increment` should change it to 11. However, due to the default value associated to the `Unsigned_8` type — which is set to 0 — the value of `V` is overwritten in the declaration of `BF` (in `Display_Bytes_Increment`). Therefore, the value of `V` is 1 after the call to `Display_Bytes_Increment`. Of course, this is not the behavior that we originally intended.

Using the `Import` aspect solves this problem. This aspect tells the compiler to not apply default

initialization in the declaration because the object is imported. Let's look at the corrected example:

[Ada]

Listing 37: p.ads

```
1 package P is
2
3   type Unsigned_8 is mod 2 ** 8 with Default_Value => 0;
4
5   type Byte_Field is array (Natural range <>) of Unsigned_8;
6
7   procedure Display_Bytes_Increment (V : in out Integer);
8 end P;
```

Listing 38: p.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body P is
4
5   procedure Display_Bytes_Increment (V : in out Integer) is
6     BF : Byte_Field (1 .. V'Size / 8)
7       with Address => V'Address, Import, Volatile;
8   begin
9     for B of BF loop
10      Put_Line ("Byte = " & Unsigned_8'Image (B));
11    end loop;
12    Put_Line ("Now incrementing...");
13    V := V + 1;
14  end Display_Bytes_Increment;
15
16 end P;
```

Listing 39: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 with P; use P;
4
5 procedure Main is
6   V : Integer := 10;
7 begin
8   Put_Line ("V = " & Integer'Image (V));
9   Display_Bytes_Increment (V);
10  Put_Line ("V = " & Integer'Image (V));
11 end Main;
```

Build output

Compile	
[Ada]	main.adb
[Ada]	p.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
V = 10
```

(continues on next page)

(continued from previous page)

```

Byte = 10
Byte = 0
Byte = 0
Byte = 0
Now incrementing...
V = 11

```

This unwanted side-effect of the initialization by the `Default_Value` aspect that we've just seen can also happen in these cases:

- when we set a default value for components of a record type declaration,
- when we use the `Default_Component_Value` aspect for array types, or
- when we set use the `Initialize Scalars` pragma for a package.

Again, using the `Import` aspect when declaring the overlay eliminates this side-effect.

We can use this pattern for objects of more complex data types like arrays or records. For example:
[Ada]

Listing 40: int_array_bitfield.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Int_Array_Bitfield is
4   type Bit_Field is array (Natural range <>) of Boolean with Pack;
5
6   A : array (1 .. 2) of Integer := (others => 0);
7   B : Bit_Field (0 .. A'Size - 1)
8     with Address => A'Address, Import, Volatile;
9 begin
10  B (2) := True;
11  for I in A'Range loop
12    Put_Line ("A (" & Integer'Image (I)
13              & ")= " & Integer'Image (A (I)));
14  end loop;
15 end Int_Array_Bitfield;

```

Build output

```

Compile
[Ada]      int_array_bitfield.adb
Bind
[gnatbind] int_array_bitfield.bexch
[Ada]      int_array_bitfield.ali
Link
[link]     int_array_bitfield.adb

```

Runtime output

```

A ( 1)= 4
A ( 2)= 0

```

In the Ada example above, we're using the bit-field to set bit #3 of the first element of the array (A (1)). We could set bit #4 of the second element by using the size of the data type (in this case, `Integer'Size`):

[Ada]

```

B (Integer'Size + 3) := True;

```

In C, we would select the specific array position and, again, rely on bit-shifting and masking to set that specific bit:

[C]

Listing 41: bitfield_int_array.c

```
1  #include <stdio.h>
2
3  int main(int argc, const char * argv[])
4  {
5      int i;
6      int a[2] = {0, 0};
7
8      a[0] = a[0] | (1 << 2);
9
10     for (i = 0; i < 2; i++)
11     {
12         printf("a[%d] = %d\n", i, a[i]);
13     }
14
15     return 0;
16 }
```

Runtime output

```
a[0] = 4
a[1] = 0
```

Since we can use this pattern for any arbitrary data type, this allows us to easily create a subprogram to serialize data types and, for example, transmit complex data structures as a bitstream. For example:

[Ada]

Listing 42: serializer.ads

```
1  package Serializer is
2
3      type Bit_Field is array (Natural range <>) of Boolean with Pack;
4
5      procedure Transmit (B : Bit_Field);
6
7  end Serializer;
```

Listing 43: serializer.adb

```
1  with Ada.Text_IO; use Ada.Text_IO;
2
3  package body Serializer is
4
5      procedure Transmit (B : Bit_Field) is
6
7          procedure Show_Bit (V : Boolean) is
8              begin
9                  case V is
10                     when False => Put ("0");
11                     when True  => Put ("1");
12                 end case;
13             end Show_Bit;
14
15          begin
16              Put ("Bits: ");
```

(continues on next page)

(continued from previous page)

```

17         for E of B loop
18             Show_Bit (E);
19         end loop;
20         New_Line;
21     end Transmit;
22
23 end Serializer;

```

Listing 44: my_recs.ads

```
1 package My_Recs is
2
3     type Rec is record
4         V : Integer;
5         S : String (1 .. 3);
6     end record;
7
8 end My_Recs;
```

Listing 45: main.adb

```

1  with Serializer; use Serializer;
2  with My_Recs; use My_Recs;
3
4  procedure Main is
5      R : Rec := (5, "abc");
6      B : Bit_Field (0 .. R'Size - 1)
7          with Address => R'Address, Import, Volatile;
8  begin
9      Transmit (B);
10 end Main;

```

Build output

```
Compile
  [Ada]          main.adb
main.adb:9:14: warning: volatile actual passed by copy (RM C.6(19)) [enabled by default]
  [Ada]          my_recs.ads
  [Ada]          serializer.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

[illegible]

In this example, the `Transmit` procedure from `Serializer` package displays the individual bits of a bit-field. We could have used this strategy to actually transmit the information as a bitstream. In the main application, we call `Transmit` for the object `R` of record type `Rec`. Since `Transmit` has the bit-field type as a parameter, we can use it for any type, as long as we have a corresponding bit-field representation.

In C, we interpret the input pointer as an array of bytes, and then use shifting and masking to access the bits of that byte. Here, we use the `char` type because it has a size of one byte in most platforms.

[C]

Listing 46: my_recs.h

```
1 typedef struct {
2     int v;
3     char s[4];
4 } rec;
```

Listing 47: serializer.h

```
1 void transmit (void *bits, int len);
```

Listing 48: serializer.c

```
1 #include "serializer.h"
2
3 #include <stdio.h>
4 #include <assert.h>
5
6 void transmit (void *bits, int len)
7 {
8     int i, j;
9     char *c = (char *)bits;
10
11     assert(sizeof(char) == 1);
12
13     printf("Bits: ");
14     for (i = 0; i < len / (sizeof(char) * 8); i++)
15     {
16         for (j = 0; j < sizeof(char) * 8; j++)
17         {
18             printf("%d", c[i] >> j & 1);
19         }
20     }
21     printf("\n");
22 }
```

Listing 49: bitfield_serialization.c

```
1 #include <stdio.h>
2
3 #include "my_recs.h"
4 #include "serializer.h"
5
6 int main(int argc, const char * argv[])
7 {
8     rec r = {5, "abc"};
9
10    transmit(&r, sizeof(r) * 8);
11
12    return 0;
13 }
```

Runtime output

```
Bits: 1010000000000000000000000000000010000110010001101100011000000000
```

Similarly, we can write a subprogram that converts a bit-field — which may have been received as a bitstream — to a specific type. We can add a `To_Rec` subprogram to the `My_Recs` package to convert a bit-field to the `Rec` type. This can be used to convert a bitstream that we received into the actual data type representation.

As you know, we may write the `To_Rec` subprogram as a procedure or as a function. Since we need to use slightly different strategies for the implementation, the following example has both versions of `To_Rec`.

This is the updated code for the `My_Recs` package and the `Main` procedure:

[Ada]

Listing 50: `serializer.ads`

```

1 package Serializer is
2
3     type Bit_Field is array (Natural range <>) of Boolean with Pack;
4
5     procedure Transmit (B : Bit_Field);
6
7 end Serializer;
```

Listing 51: `serializer.adb`

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body Serializer is
4
5     procedure Transmit (B : Bit_Field) is
6
7         procedure Show_Bit (V : Boolean) is
8             begin
9                 case V is
10                     when False => Put ("0");
11                     when True  => Put ("1");
12                 end case;
13             end Show_Bit;
14
15         begin
16             Put ("Bits: ");
17             for E of B loop
18                 Show_Bit (E);
19             end loop;
20             New_Line;
21         end Transmit;
22
23 end Serializer;
```

Listing 52: `my_recs.ads`

```

1 with Serializer; use Serializer;
2
3 package My_Recs is
4
5     type Rec is record
6         V : Integer;
7         S : String (1 .. 3);
8     end record;
9
10    procedure To_Rec (B : Bit_Field;
11                    R : out Rec);
12
13    function To_Rec (B : Bit_Field) return Rec;
14
15    procedure Display (R : Rec);
16
17 end My_Recs;
```

Listing 53: my_recs.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  package body My_Recs is
4
5      procedure To_Rec (B :      Bit_Field;
6                       R : out Rec) is
7          B_R : Rec
8              with Address => B'Address, Import, Volatile;
9      begin
10         -- Assigning data from overlayed record B_R to output parameter R.
11         R := B_R;
12     end To_Rec;
13
14     function To_Rec (B : Bit_Field) return Rec is
15         R : Rec;
16         B_R : Rec
17             with Address => B'Address, Import, Volatile;
18     begin
19         -- Assigning data from overlayed record B_R to local record R.
20         R := B_R;
21
22         return R;
23     end To_Rec;
24
25     procedure Display (R : Rec) is
26     begin
27         Put ("(" & Integer'Image (R.V) & ", "
28             & (R.S) & ")");
29     end Display;
30
31 end My_Recs;

```

Listing 54: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2  with Serializer;  use Serializer;
3  with My_Recs;    use My_Recs;
4
5  procedure Main is
6      R1 : Rec := (5, "abc");
7      R2 : Rec := (0, "zzz");
8
9      B1 : Bit_Field (0 .. R1'Size - 1)
10         with Address => R1'Address, Import, Volatile;
11  begin
12      Put ("R2 = ");
13      Display (R2);
14      New_Line;
15
16      -- Getting Rec type using data from B1, which is a bit-field
17      -- representation of R1.
18      To_Rec (B1, R2);
19
20      -- We could use the function version of To_Rec:
21      -- R2 := To_Rec (B1);
22
23      Put_Line ("New bitstream received!");
24      Put ("R2 = ");
25      Display (R2);

```

(continues on next page)

(continued from previous page)

```

26   New_Line;
27 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
main.adb:18:12: warning: volatile actual passed by copy (RM C.6(19)) [enabled by
↳ default]
  [Ada]          my_recs.adb
  [Ada]          serializer.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```

R2 = ( 0, zzz)
New bitstream received!
R2 = ( 5, abc)

```

In both versions of `To_Rec`, we declare the record object `B_R` as an overlay of the input bit-field. In the procedure version of `To_Rec`, we then simply copy the data from `B_R` to the output parameter `R`. In the function version of `To_Rec`, however, we need to declare a local record object `R`, which we return after the assignment.

In C, we can interpret the input pointer as an array of bytes, and copy the individual bytes. For example:

[C]

Listing 55: `my_recs.h`

```

1  typedef struct {
2      int v;
3      char s[3];
4  } rec;
5
6  void to_r (void *bits, int len, rec *r);
7
8  void display_r (rec *r);

```

Listing 56: `my_recs.c`

```

1  #include "my_recs.h"
2
3  #include <stdio.h>
4  #include <assert.h>
5
6  void to_r (void *bits, int len, rec *r)
7  {
8      int i;
9      char *c1 = (char *)bits;
10     char *c2 = (char *)r;
11
12     assert(len == sizeof(rec) * 8);
13
14     for (i = 0; i < len / (sizeof(char) * 8); i++)
15     {

```

(continues on next page)

(continued from previous page)

```

16     c2[i] = c1[i];
17 }
18 }
19
20 void display_r (rec *r)
21 {
22     printf("{%d, %c%c%c}", r->v, r->s[0], r->s[1], r->s[2]);
23 }

```

Listing 57: bitfield_serialization.c

```

1  #include <stdio.h>
2  #include "my_recs.h"
3
4  int main(int argc, const char * argv[])
5  {
6      rec r1 = {5, "abc"};
7      rec r2 = {0, "zzz"};
8
9      printf("r2 = ");
10     display_r (&r2);
11     printf("\n");
12
13     to_r(&r1, sizeof(r1) * 8, &r2);
14
15     printf("New bitstream received!\n");
16     printf("r2 = ");
17     display_r (&r2);
18     printf("\n");
19
20     return 0;
21 }

```

Runtime output

```

r2 = {0, zzz}
New bitstream received!
r2 = {5, abc}

```

Here, `to_r` casts both pointer parameters to pointers to `char` to get a byte-aligned pointer. Then, it simply copies the data byte-by-byte.

6.10.1 Overlays vs. Unchecked Conversions

Unchecked conversions are another way of converting between unrelated data types. This conversion is done by instantiating the generic `Unchecked_Conversions` function for the types you want to convert. Let's look at a simple example:

[Ada]

Listing 58: simple_unchecked_conversion.adb

```

1  with Ada.Text_IO;           use Ada.Text_IO;
2  with Ada.Unchecked_Conversion;
3
4  procedure Simple_Unchecked_Conversion is
5      type State is (Off, State_1, State_2)
6          with Size => Integer'Size;
7

```

(continues on next page)

(continued from previous page)

```

8   for State use (Off => 0, State_1 => 32, State_2 => 64);
9
10  function As_Integer is new Ada.Unchecked_Conversion (Source => State,
11                                                    Target => Integer);
12
13  I : Integer;
14  begin
15    I := As_Integer (State_2);
16    Put_Line ("I = " & Integer'Image (I));
17  end Simple_Unchecked_Conversion;

```

Build output

```

Compile
  [Ada]          simple_unchecked_conversion.adb
Bind
  [gprbind]      simple_unchecked_conversion.bexch
  [Ada]          simple_unchecked_conversion.ali
Link
  [link]         simple_unchecked_conversion.adb

```

Runtime output

```
I = 64
```

In this example, `As_Integer` is an instantiation of `Unchecked_Conversion` to convert between the `State` enumeration and the `Integer` type. Note that, in order to ensure safe conversion, we're declaring `State` to have the same size as the `Integer` type we want to convert to.

This is the corresponding implementation using overlays:

[Ada]

Listing 59: simple_overlay.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Simple_Overlay is
4    type State is (Off, State_1, State_2)
5      with Size => Integer'Size;
6
7    for State use (Off => 0, State_1 => 32, State_2 => 64);
8
9    S : State;
10   I : Integer
11     with Address => S'Address, Import, Volatile;
12  begin
13    S := State_2;
14    Put_Line ("I = " & Integer'Image (I));
15  end Simple_Overlay;

```

Build output

```

Compile
  [Ada]          simple_overlay.adb
Bind
  [gprbind]      simple_overlay.bexch
  [Ada]          simple_overlay.ali
Link
  [link]         simple_overlay.adb

```

Runtime output

```
I = 64
```

Let's look at another example of converting between different numeric formats. In this case, we want to convert between a 16-bit fixed-point and a 16-bit integer data type. This is how we can do it using `Unchecked_Conversion`:

[Ada]

Listing 60: `fixed_int_unchecked_conversion.adb`

```
1 with Ada.Text_IO;           use Ada.Text_IO;
2 with Ada.Unchecked_Conversion;
3
4 procedure Fixed_Int_Unchecked_Conversion is
5   Delta_16 : constant := 1.0 / 2.0 ** (16 - 1);
6   Max_16   : constant := 2 ** 15;
7
8   type Fixed_16 is delta Delta_16 range -1.0 .. 1.0 - Delta_16
9     with Size => 16;
10  type Int_16   is range -Max_16 .. Max_16 - 1
11    with Size => 16;
12
13  function As_Int_16 is new Ada.Unchecked_Conversion (Source => Fixed_16,
14                                                    Target => Int_16);
15  function As_Fixed_16 is new Ada.Unchecked_Conversion (Source => Int_16,
16                                                    Target => Fixed_16);
17
18  I : Int_16 := 0;
19  F : Fixed_16 := 0.0;
20 begin
21  F := Fixed_16'Last;
22  I := As_Int_16 (F);
23
24  Put_Line ("F = " & Fixed_16'Image (F));
25  Put_Line ("I = " & Int_16'Image (I));
26 end Fixed_Int_Unchecked_Conversion;
```

Build output

Compile	
[Ada]	<code>fixed_int_unchecked_conversion.adb</code>
Bind	
[gprbind]	<code>fixed_int_unchecked_conversion.bexch</code>
[Ada]	<code>fixed_int_unchecked_conversion.ali</code>
Link	
[link]	<code>fixed_int_unchecked_conversion.adb</code>

Runtime output

```
F = 0.99997
I = 32767
```

Here, we instantiate `Unchecked_Conversion` for the `Int_16` and `Fixed_16` types, and we call the instantiated functions explicitly. In this case, we call `As_Int_16` to get the integer value corresponding to `Fixed_16'Last`.

This is how we can rewrite the implementation above using overlays:

[Ada]

Listing 61: fixed_int_overlay.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Fixed_Int_Overlay is
4   Delta_16 : constant := 1.0 / 2.0 ** (16 - 1);
5   Max_16   : constant := 2 ** 15;
6
7   type Fixed_16 is delta Delta_16 range -1.0 .. 1.0 - Delta_16
8     with Size => 16;
9   type Int_16   is range -Max_16 .. Max_16 - 1
10    with Size => 16;
11
12   I : Int_16 := 0;
13   F : Fixed_16
14     with Address => I'Address, Import, Volatile;
15 begin
16   F := Fixed_16'Last;
17
18   Put_Line ("F = " & Fixed_16'Image (F));
19   Put_Line ("I = " & Int_16'Image (I));
20 end Fixed_Int_Overlay;

```

Build output

```

Compile
  [Ada]          fixed_int_overlay.adb
Bind
  [gprbind]      fixed_int_overlay.bexch
  [Ada]          fixed_int_overlay.ali
Link
  [link]         fixed_int_overlay.adb

```

Runtime output

```

F =  0.99997
I =  32767

```

Here, the conversion to the integer value is implicit, so we don't need to call a conversion function.

Using `Unchecked_Conversion` has the advantage of making it clear that a conversion is happening, since the conversion is written explicitly in the code. With overlays, that conversion is automatic and therefore implicit. In that sense, using `Unchecked_Conversion` is a cleaner and safer approach. On the other hand, `Unchecked_Conversion` requires a copy, so it's less efficient than overlays, where no copy is performed — because one change in the source object is automatically reflected in the target object (and vice-versa). In the end, the choice between unchecked conversions and overlays depends on the level of performance that you want to achieve.

Also note that `Unchecked_Conversion` can only be instantiated for constrained types. In order to rewrite the examples using bit-fields that we've seen in the previous section, we cannot simply instantiate `Unchecked_Conversion` with the `Target` indicating the *unconstrained* bit-field, such as:

```

Ada.Unchecked_Conversion (Source => Integer,
                          Target => Bit_Field);

```

Instead, we have to declare a subtype for the specific range we're interested in. This is how we can rewrite one of the previous examples:

```
[Ada]
```


Listing 62: simple_bitfield_conversion.adb

```
1 with Ada.Text_IO;           use Ada.Text_IO;
2 with Ada.Unchecked_Conversion;
3
4 procedure Simple_Bitfield_Conversion is
5   type Bit_Field is array (Natural range <>) of Boolean with Pack;
6
7   V : Integer := 4;
8
9   -- Declaring subtype that takes the size of V into account.
10  --
11  subtype Integer_Bit_Field is Bit_Field (0 .. V'Size - 1);
12
13  -- NOTE: we could also use the Integer type in the declaration:
14  --
15  --   subtype Integer_Bit_Field is Bit_Field (0 .. Integer'Size - 1);
16  --
17
18  -- Using the Integer_Bit_Field subtype as the target
19  function As_Bit_Field is new
20    Ada.Unchecked_Conversion (Source => Integer,
21                              Target => Integer_Bit_Field);
22
23  B : Integer_Bit_Field;
24  begin
25    B := As_Bit_Field (V);
26
27    Put_Line ("V = " & Integer'Image (V));
28  end Simple_Bitfield_Conversion;
```

Build output

Compile	
[Ada]	simple_bitfield_conversion.adb
Bind	
[gprbind]	simple_bitfield_conversion.bexch
[Ada]	simple_bitfield_conversion.ali
Link	
[link]	simple_bitfield_conversion.adb

Runtime output

```
V = 4
```

In this example, we first declare the subtype `Integer_Bit_Field` as a bit-field with a length that fits the `V` variable we want to convert to. Then, we can use that subtype in the instantiation of `Unchecked_Conversion`.

HANDLING VARIABILITY AND RE-USABILITY

7.1 Understanding static and dynamic variability

It is common to see embedded software being used in a variety of configurations that require small changes to the code for each instance. For example, the same application may need to be portable between two different architectures (ARM and x86), or two different platforms with different set of devices available. Maybe the same application is used for two different generations of the product, so it needs to account for absence or presence of new features, or it's used for different projects which may select different components or configurations. All these cases, and many others, require variability in the software in order to ensure its reusability.

In C, variability is usually achieved through macros and function pointers, the former being tied to static variability (variability in different builds) the latter to dynamic variability (variability within the same build decided at run-time).

Ada offers many alternatives for both techniques, which aim at structuring possible variations of the software. When Ada isn't enough, the GNAT compilation system also provides a layer of capabilities, in particular selection of alternate bodies.

If you're familiar with object-oriented programming (OOP) — supported in languages such as C++ and Java —, you might also be interested in knowing that OOP is supported by Ada and can be used to implement variability. This should, however, be used with care, as OOP brings its own set of problems, such as loss of efficiency — dispatching calls can't be inlined and require one level of indirection — or loss of analyzability — the target of a dispatching call isn't known at run time. As a rule of thumb, OOP should be considered only for cases of dynamic variability, where several versions of the same object need to exist concurrently in the same application.

7.2 Handling variability & reusability statically

7.2.1 Genericity

One usage of C macros involves the creation of functions that works regardless of the type they're being called upon. For example, a swap macro may look like:

[C]

Listing 1: main.c

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 #define SWAP(t, a, b) ({\
5     t tmp = a; \
6     a = b; \
```

(continues on next page)

(continued from previous page)

```

7           b = tmp; \
8       })
9
10  int main()
11  {
12      int a = 10;
13      int b = 42;
14
15      printf("a = %d, b = %d\n", a, b);
16
17      SWAP (int, a, b);
18
19      printf("a = %d, b = %d\n", a, b);
20
21      return 0;
22  }

```

Runtime output

```

a = 10, b = 42
a = 42, b = 10

```

Ada offers a way to declare this kind of functions as a generic, that is, a function that is written after static arguments, such as a parameter:

[Ada]

Listing 2: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4
5      generic
6          type A_Type is private;
7          procedure Swap (Left, Right : in out A_Type);
8
9          procedure Swap (Left, Right : in out A_Type) is
10             Temp : constant A_Type := Left;
11         begin
12             Left := Right;
13             Right := Temp;
14         end Swap;
15
16         procedure Swap_I is new Swap (Integer);
17
18         A : Integer := 10;
19         B : Integer := 42;
20
21     begin
22         Put_Line ("A = "
23             & Integer'Image (A)
24             & ", B = "
25             & Integer'Image (B));
26
27         Swap_I (A, B);
28
29         Put_Line ("A = "
30             & Integer'Image (A)
31             & ", B = "
32             & Integer'Image (B));
33     end Main;

```

Build output

```

Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```

A = 10, B = 42
A = 42, B = 10

```

There are a few key differences between the C and the Ada version here. In C, the macro can be used directly and essentially get expanded by the preprocessor without any kind of checks. In Ada, the generic will first be checked for internal consistency. It then needs to be explicitly instantiated for a concrete type. From there, it's exactly as if there was an actual version of this `Swap` function, which is going to be called as any other function. All rules for parameter modes and control will apply to this instance.

In many respects, an Ada generic is a way to provide a safe specification and implementation of such macros, through both the validation of the generic itself and its usage.

Subprograms aren't the only entities that can be made generic. As a matter of fact, it's much more common to render an entire package generic. In this case the instantiation creates a new version of all the entities present in the generic, including global variables. For example:

[Ada]

Listing 3: gen.ads

```

1  generic
2    type T is private;
3  package Gen is
4    type C is tagged record
5      V : T;
6    end record;
7
8    G : Integer;
9  end Gen;

```

The above can be instantiated and used the following way:

Listing 4: main.adb

```

1  with Gen;
2
3  procedure Main is
4    package I1 is new Gen (Integer);
5    package I2 is new Gen (Integer);
6    subtype Str10 is String (1 .. 10);
7    package I3 is new Gen (Str10);
8  begin
9    I1.G := 0;
10   I2.G := 1;
11   I3.G := 2;
12 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
  [Ada]          gen.ads
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Here, I1.G, I2.G and I3.G are three distinct variables.

So far, we've only looked at generics with one kind of parameter: a so-called private type. There's actually much more that can be described in this section, such as variables, subprograms or package instantiations with certain properties. For example, the following provides a sort algorithm for any kind of structurally compatible array type:

[Ada]

Listing 5: sort.ads

```

1 generic
2   type Component is private;
3   type Index is (<);
4   with function "<" (Left, Right : Component) return Boolean;
5   type Array_Type is array (Index range <>) of Component;
6   procedure Sort (A : in out Array_Type);

```

The declaration above states that we need a type (Component), a discrete type (Index), a comparison subprogram (" $<$ "), and an array definition (Array_Type). Given these, it's possible to write an algorithm that can sort any Array_Type. Note the usage of the with reserved word in front of the function name: it exists to differentiate between the generic parameter and the beginning of the generic subprogram.

Here is a non-exhaustive overview of the kind of constraints that can be put on types:

```

type T is private; -- T is a constrained type, such as Integer
type T (<>) is private; -- T can be an unconstrained type e.g. String
type T is tagged private; -- T is a tagged type
type T is new T2 with private; -- T is an extension of T2
type T is (<); -- T is a discrete type
type T is range <>; -- T is an integer type
type T is digits <>; -- T is a floating point type
type T is access T2; -- T is an access type to T2

```

For a more complete list please reference the [Generic Formal Types Appendix](#)²⁰.

7.2.2 Simple derivation

Let's take a case where a codebase needs to handle small variations of a given device, or maybe different generations of a device, depending on the platform it's running on. In this example, we're assuming that each platform will lead to a different binary, so the code can statically resolve which set of services are available. However, we want an easy way to implement a new device based on a previous one, saying "this new device is the same as this previous device, with these new services and these changes in existing services".

We can implement such patterns using Ada's simple derivation — as opposed to tagged derivation, which is OOP-related and discussed in a later section.

Let's start from the following example:

²⁰ <https://learn.adacore.com/courses/intro-to-ada/chapters/appendices.html#appendix-a-generic-formal-types>

[Ada]

Listing 6: drivers_1.ads

```

1 package Drivers_1 is
2
3     type Device_1 is null record;
4     procedure Startup (Device : Device_1);
5     procedure Send (Device : Device_1; Data : Integer);
6     procedure Send_Fast (Device : Device_1; Data : Integer);
7     procedure Receive (Device : Device_1; Data : out Integer);
8
9 end Drivers_1;
```

Listing 7: drivers_1.adb

```

1 package body Drivers_1 is
2
3     -- NOTE: unimplemented procedures: Startup, Send, Send_Fast
4     --       mock-up implementation: Receive
5
6     procedure Startup (Device : Device_1) is null;
7
8     procedure Send (Device : Device_1; Data : Integer) is null;
9
10    procedure Send_Fast (Device : Device_1; Data : Integer) is null;
11
12    procedure Receive (Device : Device_1; Data : out Integer) is
13    begin
14        Data := 42;
15    end Receive;
16
17 end Drivers_1;
```

In the above example, Device_1 is an empty record type. It may also have some fields if required, or be a different type such as a scalar. Then the four procedures Startup, Send, Send_Fast and Receive are primitives of this type. A primitive is essentially a subprogram that has a parameter or return type directly referencing this type and declared in the same scope. At this stage, there's nothing special with this type: we're using it as we would use any other type. For example:

[Ada]

Listing 8: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Drivers_1;   use Drivers_1;
3
4 procedure Main is
5     D : Device_1;
6     I : Integer;
7 begin
8     Startup (D);
9     Send_Fast (D, 999);
10    Receive (D, I);
11    Put_Line (Integer'Image (I));
12 end Main;
```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      drivers_1.adb
```

(continues on next page)

(continued from previous page)

```
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```
42
```

Let's now assume that we need to implement a new generation of device, `Device_2`. This new device works exactly like the first one, except for the startup code that has to be done differently. We can create a new type that operates exactly like the previous one, but modifies only the behavior of `Startup`:

[Ada]

Listing 9: drivers_2.ads

```
1 with Drivers_1; use Drivers_1;
2
3 package Drivers_2 is
4
5     type Device_2 is new Device_1;
6
7     overriding
8     procedure Startup (Device : Device_2);
9
10 end Drivers_2;
```

Listing 10: drivers_2.adb

```
1 package body Drivers_2 is
2
3     overriding
4     procedure Startup (Device : Device_2) is null;
5
6 end Drivers_2;
```

Here, `Device_2` is derived from `Device_1`. It contains all the exact same properties and primitives, in particular, `Startup`, `Send`, `Send_Fast` and `Receive`. However, here, we decided to change the `Startup` function and to provide a different implementation. We override this function. The main subprogram doesn't change much, except for the fact that it now relies on a different type:

[Ada]

Listing 11: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Drivers_2;   use Drivers_2;
3
4 procedure Main is
5     D : Device_2;
6     I : Integer;
7 begin
8     Startup (D);
9     Send_Fast (D, 999);
10    Receive (D, I);
11    Put_Line (Integer'Image (I));
12 end Main;
```

Build output

```

Compile
  [Ada]          main.adb
  [Ada]          drivers_2.adb
  [Ada]          drivers_1.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```
42
```

We can continue with this approach and introduce a new generation of devices. This new device doesn't implement the `Send_Fast` service so we want to remove it from the list of available services. Furthermore, for the purpose of our example, let's assume that the hardware team went back to the `Device_1` way of implementing `Startup`. We can write this new device the following way:

[Ada]

Listing 12: drivers_3.ads

```

1  with Drivers_1; use Drivers_1;
2
3  package Drivers_3 is
4
5      type Device_3 is new Device_1;
6
7      overriding
8      procedure Startup (Device : Device_3);
9
10     procedure Send_Fast (Device : Device_3; Data : Integer)
11     is abstract;
12
13 end Drivers_3;

```

Listing 13: drivers_3.adb

```

1  package body Drivers_3 is
2
3      overriding
4      procedure Startup (Device : Device_3) is null;
5
6  end Drivers_3;

```

The `is abstract` definition makes illegal any call to a function, so calls to `Send_Fast` on `Device_3` will be flagged as being illegal. To then implement `Startup` of `Device_3` as being the same as the `Startup` of `Device_1`, we can convert the type in the implementation:

[Ada]

Listing 14: drivers_3.adb

```

1  package body Drivers_3 is
2
3      overriding
4      procedure Startup (Device : Device_3) is
5      begin
6          Drivers_1.Startup (Device_1 (Device));
7      end Startup;

```

(continues on next page)

(continued from previous page)

```
8
9 end Drivers_3;
```

Our Main now looks like:

[Ada]

Listing 15: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Drivers_3;   use Drivers_3;
3
4 procedure Main is
5     D : Device_3;
6     I : Integer;
7 begin
8     Startup (D);
9     Send_Fast (D, 999);
10    Receive (D, I);
11    Put_Line (Integer'Image (I));
12 end Main;
```

Compilation output

```
main.adb:9:04: error: cannot call abstract operation "Send_Fast" declared at
↳ drivers_3.ads:10
```

Here, the call to Send_Fast will get flagged by the compiler.

Note that the fact that the code of Main has to be changed for every implementation isn't necessarily satisfactory. We may want to go one step further, and isolate the selection of the device kind to be used for the whole application in one unique file. One way to do this is to use the same name for all types, and use a renaming to select which package to use. Here's a simplified example to illustrate that:

[Ada]

Listing 16: drivers_1.ads

```
1 package Drivers_1 is
2
3     type Transceiver is null record;
4     procedure Send (Device : Transceiver; Data : Integer);
5     procedure Receive (Device : Transceiver; Data : out Integer);
6
7 end Drivers_1;
```

Listing 17: drivers_1.adb

```
1 package body Drivers_1 is
2
3     procedure Send (Device : Transceiver; Data : Integer) is null;
4
5     procedure Receive (Device : Transceiver; Data : out Integer) is
6         pragma Unreferenced (Device);
7     begin
8         Data := 42;
9     end Receive;
10
11 end Drivers_1;
```

Listing 18: drivers_2.ads

```

1 with Drivers_1;
2
3 package Drivers_2 is
4
5     type Transceiver is new Drivers_1.Transceiver;
6     procedure Send (Device : Transceiver; Data : Integer);
7     procedure Receive (Device : Transceiver; Data : out Integer);
8
9 end Drivers_2;

```

Listing 19: drivers_2.adb

```

1 package body Drivers_2 is
2
3     procedure Send (Device : Transceiver; Data : Integer) is null;
4
5     procedure Receive (Device : Transceiver; Data : out Integer) is
6         pragma Unreferenced (Device);
7     begin
8         Data := 42;
9     end Receive;
10
11 end Drivers_2;

```

Listing 20: drivers.ads

```

1 with Drivers_1;
2
3 package Drivers renames Drivers_1;

```

Listing 21: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Drivers;      use Drivers;
3
4 procedure Main is
5     D : Transceiver;
6     I : Integer;
7 begin
8     Send (D, 999);
9     Receive (D, I);
10    Put_Line (Integer'Image (I));
11 end Main;

```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      drivers.ads
  [Ada]      drivers_1.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb

```

Runtime output

42

In the above example, the whole code can rely on `drivers.ads`, instead of relying on the specific driver. Here, `Drivers` is another name for `Driver_1`. In order to switch to `Driver_2`, the project only has to replace that one `drivers.ads` file.

In the following section, we'll go one step further and demonstrate that this selection can be done through a configuration switch selected at build time instead of a manual code modification.

7.2.3 Configuration pragma files

Configuration pragmas are a set of pragmas that modify the compilation of source-code files. You may use them to either relax or strengthen requirements. For example:

```
pragma Suppress (Overflow_Check);
```

In this example, we're suppressing the overflow check, thereby relaxing a requirement. Normally, the following program would raise a constraint error due to a failed overflow check:

[Ada]

Listing 22: p.ads

```
1 package P is
2   function Add_Max (A : Integer) return Integer;
3 end P;
```

Listing 23: p.adb

```
1 package body P is
2   function Add_Max (A : Integer) return Integer is
3   begin
4     return A + Integer'Last;
5   end Add_Max;
6 end P;
```

Listing 24: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with P;           use P;
3
4 procedure Main is
5   I : Integer := Integer'Last;
6 begin
7   I := Add_Max (I);
8   Put_Line ("I = " & Integer'Image (I));
9 end Main;
```

Build output

Compile	
[Ada]	main.adb
[Ada]	p.adb
Bind	
[grbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
raised CONSTRAINT_ERROR : p.adb:4 overflow check failed
```

When suppressing the overflow check, however, the program doesn't raise an exception, and the value that `Add_Max` returns is -2, which is a wraparound of the sum of the maximum integer values (`Integer'Last + Integer'Last`).

We could also strengthen requirements, as in this example:

```
pragma Restrictions (No_Floating_Point);
```

Here, the restriction forbids the use of floating-point types and objects. The following program would violate this restriction, so the compiler isn't able to compile the program when the restriction is used:

```
procedure Main is
  F : Float := 0.0;
  -- Declaration is not possible with No_Floating_Point restriction.
begin
  null;
end Main;
```

Restrictions are especially useful for high-integrity applications. In fact, the Ada Reference Manual has a [separate section for them](#)²¹.

When creating a project, it is practical to list all configuration pragmas in a separate file. This is called a configuration pragma file, and it usually has an `.adc` file extension. If you use **GPRbuild** for building Ada applications, you can specify the configuration pragma file in the corresponding project file. For example, here we indicate that `gnat.adc` is the configuration pragma file for our project:

```
project Default is

  for Source_Dirs use ("src");
  for Object_Dir use "obj";
  for Main use ("main.adb");

  package Compiler is
    for Local_Configuration_Pragmas use "gnat.adc";
  end Compiler;

end Default;
```

7.2.4 Configuration packages

In C, preprocessing flags are used to create blocks of code that are only compiled under certain circumstances. For example, we could have a block that is only used for debugging:

[C]

Listing 25: main.c

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 int func(int x)
5 {
6     return x % 4;
```

(continues on next page)

²¹ <http://www.ada-auth.org/standards/12rm/html/RM-H-4.html>

(continued from previous page)

```

7  }
8
9  int main()
10 {
11     int a, b;
12
13     a = 10;
14     b = func(a);
15
16     #ifdef DEBUG
17         printf("func(%d) => %d\n", a, b);
18     #endif
19
20     return 0;
21 }

```

Here, the block indicated by the `DEBUG` flag is only included in the build if we define this preprocessing flag, which is what we expect for a debug version of the build. In the release version, however, we want to keep debug information out of the build, so we don't use this flag during the build process.

Ada doesn't define a preprocessor as part of the language. Some Ada toolchains — like the GNAT toolchain — do have a preprocessor that could create code similar to the one we've just seen. When programming in Ada, however, the recommendation is to use configuration packages to select code blocks that are meant to be included in the application.

When using a configuration package, the example above can be written as:

[Ada]

Listing 26: config.ads

```

1  package Config is
2
3      Debug : constant Boolean := False;
4
5  end Config;

```

Listing 27: func.ads

```

1  function Func (X : Integer) return Integer;

```

Listing 28: func.adb

```

1  function Func (X : Integer) return Integer is
2  begin
3      return X mod 4;
4  end Func;

```

Listing 29: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2  with Config;
3  with Func;
4
5  procedure Main is
6      A, B : Integer;
7  begin
8      A := 10;
9      B := Func (A);

```

(continues on next page)

(continued from previous page)

```

10
11     if Config.Debug then
12         Put_Line ("Func(" & Integer'Image (A) & ") => "
13                 & Integer'Image (B));
14     end if;
15 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
  [Ada]          config.ads
  [Ada]          func.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

In this example, Config is a configuration package. The version of Config we're seeing here is the release version. The debug version of the Config package looks like this:

```

package Config is

    Debug : constant Boolean := True;

end Config;

```

The compiler makes sure to remove dead code. In the case of the release version, since Config.Debug is constant and set to False, the compiler is smart enough to remove the call to Put_Line from the build.

As you can see, both versions of Config are very similar to each other. The general idea is to create packages that declare the same constants, but using different values.

In C, we differentiate between the debug and release versions by selecting the appropriate preprocessing flags, but in Ada, we select the appropriate configuration package during the build process. Since the file name is usually the same (config.ads for the example above), we may want to store them in distinct directories. For the example above, we could have:

- src/debug/config.ads for the debug version, and
- src/release/config.ads for the release version.

Then, we simply select the appropriate configuration package for each version of the build by indicating the correct path to it. When using **GPRbuild**, we can select the appropriate directory where the config.ads file is located. We can use scenario variables in our project, which allow for creating different versions of a build. For example:

```

project Default is

    type Mode_Type is ("debug", "release");

    Mode : Mode_Type := external ("mode", "debug");

    for Source_Dirs use ("src", "src/" & Mode);
    for Object_Dir use "obj";
    for Main use ("main.adb");

end Default;

```

In this example, we're defining a scenario type called Mode_Type. Then, we're declaring the scenario variable Mode and using it in the Source_Dirs declaration to complete the path to the sub-

directory containing the `config.ads` file. The expression `"src/"` & `Mode` concatenates the user-specified mode to select the appropriate subdirectory.

We can then set the mode on the command-line. For example:

```
gprbuild -P default.gpr -Xmode=release
```

In addition to selecting code blocks for the build, we could also specify values that depend on the target build. For our example above, we may want to create two versions of the application, each one having a different version of a `MOD_VALUE` that is used in the implementation of `func()`. In C, we can achieve this by using preprocessing flags and defining the corresponding version in `APP_VERSION`. Then, depending on the value of `APP_VERSION`, we define the corresponding value of `MOD_VALUE`.

[C]

Listing 30: `defs.h`

```
1 #ifndef APP_VERSION
2 #define APP_VERSION    1
3 #endif
4
5 #if APP_VERSION == 1
6 #define MOD_VALUE      4
7 #endif
8
9 #if APP_VERSION == 2
10 #define MOD_VALUE      5
11 #endif
```

Listing 31: `main.c`

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 #include "defs.h"
5
6 int func(int x)
7 {
8     return x % MOD_VALUE;
9 }
10
11 int main()
12 {
13     int a, b;
14
15     a = 10;
16     b = func(a);
17
18     return 0;
19 }
```

If not defined outside, the code above will compile version #1 of the application. We can change this by specifying a value for `APP_VERSION` during the build (e.g. as a Makefile switch).

For the Ada version of this code, we can create two configuration packages for each version of the application. For example:

[Ada]

Listing 32: app_defs.ads

```

1  -- ./src/app_1/app_defs.ads
2
3  package App_Defs is
4
5      Mod_Value : constant Integer := 4;
6
7  end App_Defs;

```

Listing 33: func.ads

```

1  function Func (X : Integer) return Integer;

```

Listing 34: func.adb

```

1  with App_Defs;
2
3  function Func (X : Integer) return Integer is
4  begin
5      return X mod App_Defs.Mod_Value;
6  end Func;

```

Listing 35: main.adb

```

1  with Func;
2
3  procedure Main is
4      A, B : Integer;
5  begin
6      A := 10;
7      B := Func (A);
8  end Main;

```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      func.adb
  [Ada]      app_defs.ads
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb

```

The code above shows the version #1 of the configuration package. The corresponding implementation for version #2 looks like this:

```

-- ./src/app_2/app_defs.ads

package App_Defs is

    Mod_Value : constant Integer := 5;

end App_Defs;

```

Again, we just need to select the appropriate configuration package for each version of the build, which we can easily do when using **GPRbuild**.

7.3 Handling variability & reusability dynamically

7.3.1 Records with discriminants

In basic terms, records with discriminants are records that include "parameters" in their type definitions. This allows for adding more flexibility to the type definition. In the section about [pointers](#) (page 132), we've seen this example:

[Ada]

Listing 36: main.adb

```

1 procedure Main is
2   type Arr is array (Integer range <>) of Integer;
3
4   type S (Last : Positive) is record
5     A : Arr (0 .. Last);
6   end record;
7
8   V : S (9);
9 begin
10   null;
11 end Main;
```

Build output

```

Compile
  [Ada]          main.adb
main.adb:8:04: warning: variable "V" is never read and never assigned [-gnatvw]
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Here, Last is the discriminant for type S. When declaring the variable V as S (9), we specify the actual index of the last position of the array component A by setting the Last discriminant to 9.

We can create an equivalent implementation in C by declaring a struct with a pointer to an array:

[C]

Listing 37: main.c

```

1 #include <stdio.h>
2 #include <stdlib.h>
3
4 typedef struct {
5   int * a;
6   const int last;
7 } S;
8
9 S init_s (int last)
10 {
11   S v = { malloc (sizeof(int) * last + 1), last };
12   return v;
13 }
14
15 int main(int argc, const char * argv[])
16 {
17   S v = init_s (9);
18
```

(continues on next page)

(continued from previous page)

```

19   return 0;
20 }

```

Here, we need to explicitly allocate the array of the `S` struct via a call to `malloc()`, which allocates memory space on the heap. In the Ada version, in contrast, the array (`V.A`) is allocated on the stack and we don't need to explicitly allocate it.

Note that the information that we provide as the discriminant to the record type (in the Ada code) is constant, so we cannot assign a value to it. For example, we cannot write:

[Ada]

```
V.Last := 10;      -- COMPILATION ERROR!
```

In the C version, we declare the `last` field constant to get the same behavior.

[C]

```
v.last = 10;      // COMPILATION ERROR!
```

Note that the information provided as discriminants is visible. In the example above, we could display `Last` by writing:

[Ada]

```
Put_Line ("Last : " & Integer'Image (V.Last));
```

Also note that, even if a type is private, we can still access the information of the discriminants if they are visible in the *public* part of the type declaration. Let's rewrite the example above:

[Ada]

Listing 38: array_definition.ads

```

1  package Array_Definition is
2      type Arr is array (Integer range <>) of Integer;
3
4      type S (Last : Integer) is private;
5
6  private
7      type S (Last : Integer) is record
8          A : Arr (0 .. Last);
9      end record;
10
11 end Array_Definition;

```

Listing 39: main.adb

```

1  with Ada.Text_IO;      use Ada.Text_IO;
2  with Array_Definition; use Array_Definition;
3
4  procedure Main is
5      V : S (9);
6  begin
7      Put_Line ("Last : " & Integer'Image (V.Last));
8  end Main;

```

Build output

```

Compile
[Ada]      main.adb
main.adb:5:04: warning: variable "V" is read but never assigned [-gnatwv]

```

(continues on next page)

(continued from previous page)

```

[Ada]          array_definition.ads
Bind
  [gprbind]    main.bexch
  [Ada]        main.ali
Link
  [link]       main.adb

```

Runtime output

```
Last : 9
```

Even though the *S* type is now private, we can still display *Last* because this discriminant is visible in the *non-private* part of package *Array_Definition*.

7.3.2 Variant records

In simple terms, a variant record — a *discriminated record* in Ada terminology — is a record with discriminants that allows for changing its structure. Basically, it's a record containing a case. This is the general structure:

[Ada]

```

type Var_Rec (V : F) is record
    case V is
        when Opt_1 => F1 : Type_1;
        when Opt_2 => F2 : Type_2;
    end case;
end record;

```

Let's look at this example:

[Ada]

Listing 40: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4
5      type Float_Int (Use_Float : Boolean) is record
6          case Use_Float is
7              when True  => F : Float;
8              when False => I : Integer;
9          end case;
10         end record;
11
12         procedure Display (V : Float_Int) is
13             begin
14                 if V.Use_Float then
15                     Put_Line ("Float value: " & Float'Image (V.F));
16                 else
17                     Put_Line ("Integer value: " & Integer'Image (V.I));
18                 end if;
19             end Display;
20
21         F : constant Float_Int := (Use_Float => True,  F => 10.0);
22         I : constant Float_Int := (Use_Float => False, I => 9);
23

```

(continues on next page)

(continued from previous page)

```

24 begin
25     Display (F);
26     Display (I);
27 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```

Float value:    1.00000E+01
Integer value:  9

```

Here, we declare F containing a floating-point value, and I containing an integer value. In the Display procedure, we present the correct information to the user according to the Use_Float discriminant of the Float_Int type.

We can implement this example in C by using unions:

[C]

Listing 41: main.c

```

1  #include <stdio.h>
2  #include <stdlib.h>
3
4  typedef struct {
5      int use_float;
6      union {
7          float f;
8          int i;
9      };
10 } float_int;
11
12 float_int init_float (float f)
13 {
14     float_int v;
15
16     v.use_float = 1;
17     v.f         = f;
18     return v;
19 }
20
21 float_int init_int (int i)
22 {
23     float_int v;
24
25     v.use_float = 0;
26     v.i         = i;
27     return v;
28 }
29
30 void display (float_int v)
31 {

```

(continues on next page)

(continued from previous page)

```

32     if (v.use_float) {
33         printf("Float value   : %f\n", v.f);
34     }
35     else {
36         printf("Integer value : %d\n", v.i);
37     }
38 }
39
40 int main(int argc, const char * argv[])
41 {
42     float_int f = init_float (10.0);
43     float_int i = init_int (9);
44
45     display (f);
46     display (i);
47
48     return 0;
49 }

```

Runtime output

```

Float value   : 10.000000
Integer value : 9

```

Similar to the Ada code, we declare `f` containing a floating-point value, and `i` containing an integer value. One difference is that we use the `init_float()` and `init_int()` functions to initialize the `float_int` struct. These functions initialize the correct field of the union and set the `use_float` field accordingly.

7.3.2.1 Variant records and unions

There is, however, a difference in accessibility between variant records in Ada and unions in C. In C, we're allowed to access any field of the union regardless of the initialization:

[C]

```

float_int v = init_float (10.0);

printf("Integer value : %d\n", v.i);

```

This feature is useful to create overlays. In this specific example, however, the information displayed to the user doesn't make sense, since the union was initialized with a floating-point value (`v.f`) and, by accessing the integer field (`v.i`), we're displaying it as if it was an integer value.

In Ada, accessing the wrong component would raise an exception at run-time ("discriminant check failed"), since the component is checked before being accessed:

[Ada]

```

V : constant Float_Int := (Use_Float => True,  F => 10.0);
begin
  Put_Line ("Integer value: " & Integer'Image (V.I));
  --                                     ^ Constraint_Error is raised!

```

Using this method prevents wrong information being used in other parts of the program.

To get the same behavior in Ada as we do in C, we need to explicitly use the `Unchecked_Union` aspect in the type declaration. This is the modified example:

[Ada]

Listing 42: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4
5     type Float_Int_Union (Use_Float : Boolean) is record
6         case Use_Float is
7             when True  => F : Float;
8             when False => I : Integer;
9         end case;
10    end record
11    with Unchecked_Union;
12
13    V : constant Float_Int_Union := (Use_Float => True,  F => 10.0);
14
15 begin
16     Put_Line ("Integer value: " & Integer'Image (V.I));
17 end Main;

```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
Integer value: 1092616192
```

Now, we can display the integer component (V.I) even though we initialized the floating-point component (V.F). As expected, the information displayed by the test application in this case doesn't make sense.

Note that, when using the Unchecked_Union aspect in the declaration of a variant record, the reference discriminant is not available anymore, since it isn't stored as part of the record. Therefore, we cannot access the Use_Float discriminant as in the following code:

[Ada]

```

V : constant Float_Int_Union := (Use_Float => True,  F => 10.0);
begin
  if V.Use_Float then          -- COMPILATION ERROR!
    -- Do something...
  end if;

```

Unchecked unions are particularly useful in Ada when creating bindings for C code.

7.3.2.2 Optional components

We can also use variant records to specify optional components of a record. For example:

[Ada]

Listing 43: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   type Arr is array (Integer range <>) of Integer;
5
6   type Extra_Info is (No, Yes);
7
8   type S_Var (Last : Integer; Has_Extra_Info : Extra_Info) is record
9     A : Arr (0 .. Last);
10
11     case Has_Extra_Info is
12       when No    => null;
13       when Yes   => B : Arr (0 .. Last);
14     end case;
15   end record;
16
17   V1 : S_Var (Last => 9, Has_Extra_Info => Yes);
18   V2 : S_Var (Last => 9, Has_Extra_Info => No);
19 begin
20   Put_Line ("Size of V1 is: " & Integer'Image (V1'Size));
21   Put_Line ("Size of V2 is: " & Integer'Image (V2'Size));
22 end Main;
```

Build output

```
Compile
  [Ada]          main.adb
main.adb:17:04: warning: variable "V1" is read but never assigned [-gnatwv]
main.adb:18:04: warning: variable "V2" is read but never assigned [-gnatwv]
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb
```

Runtime output

```
Size of V1 is:  704
Size of V2 is:  384
```

Here, in the declaration of `S_Var`, we don't have any component in case `Has_Extra_Info` is false. The component is simply set to `null` in this case.

When running the example above, we see that the size of `V1` is greater than the size of `V2` due to the extra `B` component — which is only included when `Has_Extra_Info` is true.

7.3.2.3 Optional output information

We can use optional components to prevent subprograms from generating invalid information that could be misused by the caller. Consider the following example:

[C]

Listing 44: main.c

```

1  #include <stdio.h>
2  #include <stdlib.h>
3
4  float calculate (float f1,
5                  float f2,
6                  int *success)
7  {
8      if (f1 < f2) {
9          *success = 1;
10         return f2 - f1;
11     }
12     else {
13         *success = 0;
14         return 0.0;
15     }
16 }
17
18 void display (float v,
19              int success)
20 {
21     if (success) {
22         printf("Value = %f\n", v);
23     }
24     else {
25         printf("Calculation error!\n");
26     }
27 }
28
29 int main(int argc, const char * argv[])
30 {
31     float f;
32     int success;
33
34     f = calculate (1.0, 0.5, &success);
35     display (f, success);
36
37     f = calculate (0.5, 1.0, &success);
38     display (f, success);
39
40     return 0;
41 }

```

Runtime output

```

Calculation error!
Value = 0.500000

```

In this code, we're using the output parameter `success` of the `calculate()` function to indicate whether the calculation was successful or not. This approach has a major problem: there's no way to prevent that the invalid value returned by `calculate()` in case of an error is misused in another computation. For example:

[C]


```
int main(int argc, const char * argv[])
{
    float f;
    int success;

    f = calculate (1.0, 0.5, &success);

    f = f * 0.25;    // Using f in another computation even though
                    // calculate() returned a dummy value due to error!
                    // We should have evaluated "success", but we didn't.

    return 0;
}
```

We cannot prevent access to the returned value or, at least, force the caller to evaluate success before using the returned value.

This is the corresponding code in Ada:

[Ada]

Listing 45: main.adb

```
1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4
5      function Calculate (F1, F2 : Float;
6                          Success : out Boolean) return Float is
7
8          begin
9              if F1 < F2 then
10                 Success := True;
11                 return F2 - F1;
12             else
13                 Success := False;
14                 return 0.0;
15             end if;
16         end Calculate;
17
18         procedure Display (V : Float; Success : Boolean) is
19             begin
20                 if Success then
21                     Put_Line ("Value = " & Float'Image (V));
22                 else
23                     Put_Line ("Calculation error!");
24                 end if;
25             end Display;
26
27             F      : Float;
28             Success : Boolean;
29         begin
30             F := Calculate (1.0, 0.5, Success);
31             Display (F, Success);
32
33             F := Calculate (0.5, 1.0, Success);
34             Display (F, Success);
35         end Main;
```

Build output

```
Compile
[Ada]      main.adb
```

(continues on next page)

(continued from previous page)

```

Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```

Calculation error!
Value =  5.00000E-01

```

The Ada code above suffers from the same drawbacks as the C code. Again, there's no way to prevent misuse of the invalid value returned by `Calculate` in case of errors.

However, in Ada, we can use variant records to make the component unavailable and therefore prevent misuse of this information. Let's rewrite the original example and *wrap* the returned value in a variant record:

[Ada]

Listing 46: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4
5      type Opt_Float (Success : Boolean) is record
6          case Success is
7              when False => null;
8              when True  => F : Float;
9          end case;
10     end record;
11
12     function Calculate (F1, F2 : Float) return Opt_Float is
13     begin
14         if F1 < F2 then
15             return (Success => True, F => F2 - F1);
16         else
17             return (Success => False);
18         end if;
19     end Calculate;
20
21     procedure Display (V : Opt_Float) is
22     begin
23         if V.Success then
24             Put_Line ("Value = " & Float'Image (V.F));
25         else
26             Put_Line ("Calculation error!");
27         end if;
28     end Display;
29
30     begin
31         Display (Calculate (1.0, 0.5));
32         Display (Calculate (0.5, 1.0));
33     end Main;

```

Build output

```

Compile
  [Ada]      main.adb
Bind

```

(continues on next page)

(continued from previous page)

```
[gprbind]    main.bexch
[Ada]        main.ali
Link
[link]       main.adb
```

Runtime output

```
Calculation error!
Value = 5.00000E-01
```

In this example, we can determine whether the calculation was successful or not by evaluating the Success component of the `Opt_Float`. If the calculation wasn't successful, we won't be able to access the F component of the `Opt_Float`. As mentioned before, trying to access the component in this case would raise an exception. Therefore, in case of errors, we can ensure that no information is misused after the call to `Calculate`.

7.3.3 Object orientation

In the [previous section](#) (page 170), we've seen that we can add variability to records by using discriminants. Another approach is to use *tagged* records, which are the base for object-oriented programming in Ada.

7.3.3.1 Type extension

A tagged record type is declared by adding the `tagged` keyword. For example:

[Ada]

Listing 47: main.adb

```
1  procedure Main is
2
3      type Rec is record
4          V : Integer;
5      end record;
6
7      type Tagged_Rec is tagged record
8          V : Integer;
9      end record;
10
11     R1 : Rec;
12     R2 : Tagged_Rec;
13
14     pragma Unreferenced (R1, R2);
15 begin
16     R1 := (V => 0);
17     R2 := (V => 0);
18 end Main;
```

Build output

```
Compile
[Ada]    main.adb
Bind
[gprbind] main.bexch
[Ada]    main.ali
Link
[link]   main.adb
```

In this simple example, there isn't much difference between the `Rec` and `Tagged_Rec` type. However, tagged types can be derived and extended. For example:

[Ada]

Listing 48: main.adb

```

1  procedure Main is
2
3      type Rec is record
4          V : Integer;
5      end record;
6
7      -- We cannot declare this:
8      --
9      -- type Ext_Rec is new Rec with record
10     --     V : Integer;
11     -- end record;
12
13     type Tagged_Rec is tagged record
14         V : Integer;
15     end record;
16
17     -- But we can declare this:
18     --
19     type Ext_Tagged_Rec is new Tagged_Rec with record
20         V2 : Integer;
21     end record;
22
23     R1 : Rec;
24     R2 : Tagged_Rec;
25     R3 : Ext_Tagged_Rec;
26
27     pragma Unreferenced (R1, R2, R3);
28 begin
29     R1 := (V => 0);
30     R2 := (V => 0);
31     R3 := (V => 0, V2 => 0);
32 end Main;

```

Build output

Compile	
[Ada]	main.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

As indicated in the example, a type derived from an untagged type cannot have an extension. The compiler indicates this error if you uncomment the declaration of the `Ext_Rec` type above. In contrast, we can extend a tagged type, as we did in the declaration of `Ext_Tagged_Rec`. In this case, `Ext_Tagged_Rec` has all the components of the `Tagged_Rec` type (`V`, in this case) plus the additional components from its own type declaration (`V2`, in this case).

7.3.3.2 Overriding subprograms

Previously, we've seen that subprograms can be overridden. For example, if we had implemented a `Reset` and a `Display` procedure for the `Rec` type that we declared above, these procedures would be available for an `Ext_Rec` type derived from `Rec`. Also, we could override these procedures for the `Ext_Rec` type. In Ada, we don't need object-oriented programming features to do that: simple (untagged) records can be used to derive types, inherit operations and override them. However, in applications where the actual subprogram to be called is determined dynamically at run-time, we need dispatching calls. In this case, we must use tagged types to implement this.

7.3.3.3 Comparing untagged and tagged types

Let's discuss the similarities and differences between untagged and tagged types based on this example:

[Ada]

Listing 49: p.ads

```
1 package P is
2
3     type Rec is record
4         V : Integer;
5     end record;
6
7     procedure Display (R : Rec);
8     procedure Reset (R : out Rec);
9
10    type New_Rec is new Rec;
11
12    overriding procedure Display (R : New_Rec);
13    not overriding procedure New_Op (R : in out New_Rec);
14
15    type Tagged_Rec is tagged record
16        V : Integer;
17    end record;
18
19    procedure Display (R : Tagged_Rec);
20    procedure Reset (R : out Tagged_Rec);
21
22    type Ext_Tagged_Rec is new Tagged_Rec with record
23        V2 : Integer;
24    end record;
25
26    overriding procedure Display (R : Ext_Tagged_Rec);
27    overriding procedure Reset (R : out Ext_Tagged_Rec);
28    not overriding procedure New_Op (R : in out Ext_Tagged_Rec);
29
30 end P;
```

Listing 50: p.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body P is
4
5     procedure Display (R : Rec) is
6     begin
7         Put_Line ("TYPE: REC");
8         Put_Line ("Rec.V = " & Integer'Image (R.V));
```

(continues on next page)

(continued from previous page)

```

9      New_Line;
10 end Display;
11
12 procedure Reset (R : out Rec) is
13 begin
14     R.V := 0;
15 end Reset;
16
17 procedure Display (R : New_Rec) is
18 begin
19     Put_Line ("TYPE: NEW_REC");
20     Put_Line ("New_Rec.V = " & Integer'Image (R.V));
21     New_Line;
22 end Display;
23
24 procedure New_Op (R : in out New_Rec) is
25 begin
26     R.V := R.V + 1;
27 end New_Op;
28
29 procedure Display (R : Tagged_Rec) is
30 begin
31     -- Using External_Tag attribute to retrieve the tag as a string
32     Put_Line ("TYPE: " & Tagged_Rec'External_Tag);
33     Put_Line ("Tagged_Rec.V = " & Integer'Image (R.V));
34     New_Line;
35 end Display;
36
37 procedure Reset (R : out Tagged_Rec) is
38 begin
39     R.V := 0;
40 end Reset;
41
42 procedure Display (R : Ext_Tagged_Rec) is
43 begin
44     -- Using External_Tag attribute to retrieve the tag as a string
45     Put_Line ("TYPE: " & Ext_Tagged_Rec'External_Tag);
46     Put_Line ("Ext_Tagged_Rec.V = " & Integer'Image (R.V));
47     Put_Line ("Ext_Tagged_Rec.V2 = " & Integer'Image (R.V2));
48     New_Line;
49 end Display;
50
51 procedure Reset (R : out Ext_Tagged_Rec) is
52 begin
53     Tagged_Rec (R).Reset;
54     R.V2 := 0;
55 end Reset;
56
57 procedure New_Op (R : in out Ext_Tagged_Rec) is
58 begin
59     R.V := R.V + 1;
60 end New_Op;
61
62 end P;
```

Listing 51: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with P;           use P;
3
4 procedure Main is
```

(continues on next page)

(continued from previous page)

```

5   X_Rec          : Rec;
6   X_New_Rec      : New_Rec;
7
8   X_Tagged_Rec    : aliased Tagged_Rec;
9   X_Ext_Tagged_Rec : aliased Ext_Tagged_Rec;
10
11  X_Tagged_Rec_Array : constant array (1 .. 2) of access Tagged_Rec'Class
12                        := (X_Tagged_Rec'Access, X_Ext_Tagged_Rec'Access);
13 begin
14   --
15   --   Reset all objects
16   --
17   Reset (X_Rec);
18   Reset (X_New_Rec);
19   X_Tagged_Rec.Reset;    -- we could write "Reset (X_Tagged_Rec)" as well
20   X_Ext_Tagged_Rec.Reset;
21
22   --
23   --   Use new operations when available
24   --
25   New_Op (X_New_Rec);
26   X_Ext_Tagged_Rec.New_Op;
27
28   --
29   --   Display all objects
30   --
31   Display (X_Rec);
32   Display (X_New_Rec);
33   X_Tagged_Rec.Display;  -- we could write "Display (X_Tagged_Rec)" as well
34   X_Ext_Tagged_Rec.Display;
35
36   --
37   --   Resetting and display objects of Tagged_Rec'Class
38   --
39   Put_Line ("Operations on Tagged_Rec'Class");
40   Put_Line ("-----");
41   for E of X_Tagged_Rec_Array loop
42     E.Reset;
43     E.Display;
44   end loop;
45 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
  [Ada]          p.adb
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output

```

TYPE: REC
Rec.V = 0

TYPE: NEW_REC
New_Rec.V = 1

```

(continues on next page)

(continued from previous page)

```

TYPE: P.TAGGED_REC
Tagged_Rec.V = 0

TYPE: P.EXT_TAGGED_REC
Ext_Tagged_Rec.V = 1
Ext_Tagged_Rec.V2 = 0

Operations on Tagged_Rec'Class
-----
TYPE: P.TAGGED_REC
Tagged_Rec.V = 0

TYPE: P.EXT_TAGGED_REC
Ext_Tagged_Rec.V = 0
Ext_Tagged_Rec.V2 = 0

```

These are the similarities between untagged and tagged types:

- We can derive types and inherit operations in both cases.
 - Both `X_New_Rec` and `X_Ext_Tagged_Rec` inherit the `Display` and `Reset` procedures from their respective ancestors.
- We can override operations in both cases.
- We can implement new operations in both cases.
 - Both `X_New_Rec` and `X_Ext_Tagged_Rec` implement a procedure called `New_Op`, which is not available for their respective ancestors.

Now, let's look at the differences between untagged and tagged types:

- We can dispatch calls for a given type class.
 - This is what we do when we iterate over objects of the `Tagged_Rec` class — in the loop over `X_Tagged_Rec_Array` at the last part of the `Main` procedure.
- We can use the dot notation.
 - We can write both `E.Reset` or `Reset (E)` forms: they're equivalent.

7.3.3.4 Dispatching calls

Let's look more closely at the dispatching calls implemented above. First, we declare the `X_Tagged_Rec_Array` array and initialize it with the access to objects of both parent and derived tagged types:

[Ada]

```

X_Tagged_Rec      : aliased Tagged_Rec;
X_Ext_Tagged_Rec  : aliased Ext_Tagged_Rec;

X_Tagged_Rec_Array : constant array (1 .. 2) of access Tagged_Rec'Class
                    := (X_Tagged_Rec'Access, X_Ext_Tagged_Rec'Access);

```

Here, we use the `aliased` keyword to be able to get access to the objects (via the `'Access` attribute).

Then, we loop over this array and call the `Reset` and `Display` procedures:

[Ada]


```
for E of X_Tagged_Rec_Array loop
  E.Reset;
  E.Display;
end loop;
```

Since we're using dispatching calls, the actual procedure that is selected depends on the type of the object. For the first element (`X_Tagged_Rec_Array (1)`), this is `Tagged_Rec`, while for the second element (`X_Tagged_Rec_Array (2)`), this is `Ext_Tagged_Rec`.

Dispatching calls are only possible for a type class — for example, the `Tagged_Rec` 'Class. When the type of an object is known at compile time, the calls won't dispatch at runtime. For example, the call to the `Reset` procedure of the `X_Ext_Tagged_Rec` object (`X_Ext_Tagged_Rec.Reset`) will always take the overridden `Reset` procedure of the `Ext_Tagged_Rec` type. Similarly, if we perform a view conversion by writing `Tagged_Rec (A_Ext_Tagged_Rec).Display`, we're instructing the compiler to interpret `A_Ext_Tagged_Rec` as an object of type `Tagged_Rec`, so that the compiler selects the `Display` procedure of the `Tagged_Rec` type.

7.3.3.5 Interfaces

Another useful feature of object-oriented programming is the use of interfaces. In this case, we can define abstract operations, and implement them in the derived tagged types. We declare an interface by simply writing `type T is interface`. For example:

[Ada]

```
type My_Interface is interface;

procedure Op (Obj : My_Interface) is abstract;

-- We cannot declare actual objects of an interface:
--
-- Obj : My_Interface; -- ERROR!
```

All operations on an interface type are abstract, so we need to write `is abstract` in the signature — as we did in the declaration of `Op` above. Also, since interfaces are abstract types and don't have an actual implementation, we cannot declare objects for it.

We can derive tagged types from an interface and implement the actual operations of that interface:

[Ada]

```
type My_Derived is new My_Interface with null record;

procedure Op (Obj : My_Derived);
```

Note that we're not using the `tagged` keyword in the declaration because any type derived from an interface is automatically tagged.

Let's look at an example with an interface and two derived tagged types:

[Ada]

Listing 52: p.ads

```
1 package P is
2
3   type Display_Interface is interface;
4   procedure Display (D : Display_Interface) is abstract;
5
6   type Small_Display_Type is new Display_Interface with null record;
```

(continues on next page)

(continued from previous page)

```

7  procedure Display (D : Small_Display_Type);
8
9  type Big_Display_Type is new Display_Interface with null record;
10 procedure Display (D : Big_Display_Type);
11
12 end P;
```

Listing 53: p.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  package body P is
4
5      procedure Display (D : Small_Display_Type) is
6          pragma Unreferenced (D);
7      begin
8          Put_Line ("Using Small_Display_Type");
9      end Display;
10
11     procedure Display (D : Big_Display_Type) is
12         pragma Unreferenced (D);
13     begin
14         Put_Line ("Using Big_Display_Type");
15     end Display;
16
17 end P;
```

Listing 54: main.adb

```

1  with P; use P;
2
3  procedure Main is
4      D_Small : Small_Display_Type;
5      D_Big   : Big_Display_Type;
6
7      procedure Dispatching_Display (D : Display_Interface'Class) is
8      begin
9          D.Display;
10     end Dispatching_Display;
11
12     begin
13         Dispatching_Display (D_Small);
14         Dispatching_Display (D_Big);
15     end Main;
```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      p.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb
```

Runtime output

```

Using Small_Display_Type
Using Big_Display_Type
```

In this example, we have an interface type `Display_Interface` and two tagged types that are derived from `Display_Interface`: `Small_Display_Type` and `Big_Display_Type`.

Both types (`Small_Display_Type` and `Big_Display_Type`) implement the interface by overriding the `Display` procedure. Then, in the inner procedure `Dispatching_Display` of the `Main` procedure, we perform a dispatching call depending on the actual type of `D`.

7.3.3.6 Deriving from multiple interfaces

We may derive a type from multiple interfaces by simply writing type `Derived_T is new T1 and T2 with null record`. For example:

[Ada]

Listing 55: transceivers.ads

```
1 package Transceivers is
2
3     type Send_Interface is interface;
4
5     procedure Send (Obj : in out Send_Interface) is abstract;
6
7     type Receive_Interface is interface;
8
9     procedure Receive (Obj : in out Receive_Interface) is abstract;
10
11     type Transceiver is new Send_Interface and Receive_Interface
12         with null record;
13
14     procedure Send (D : in out Transceiver);
15     procedure Receive (D : in out Transceiver);
16
17 end Transceivers;
```

Listing 56: transceivers.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body Transceivers is
4
5     procedure Send (D : in out Transceiver) is
6         pragma Unreferenced (D);
7     begin
8         Put_Line ("Sending data...");
9     end Send;
10
11     procedure Receive (D : in out Transceiver) is
12         pragma Unreferenced (D);
13     begin
14         Put_Line ("Receiving data...");
15     end Receive;
16
17 end Transceivers;
```

Listing 57: main.adb

```
1 with Transceivers; use Transceivers;
2
3 procedure Main is
4     D : Transceiver;
5 begin
```

(continues on next page)

(continued from previous page)

```

6   D.Send;
7   D.Receive;
8 end Main;

```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      transceivers.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb

```

Runtime output

```

Sending data...
Receiving data...

```

In this example, we're declaring two interfaces (`Send_Interface` and `Receive_Interface`) and the tagged type `Transceiver` that derives from both interfaces. Since we need to implement the interfaces, we implement both `Send` and `Receive` for `Transceiver`.

7.3.3.7 Abstract tagged types

We may also declare abstract tagged types. Note that, because the type is abstract, we cannot use it to declare objects for it — this is the same as for interfaces. We can only use it to derive other types. Let's look at the abstract tagged type declared in the `Abstract_Transceivers` package:

[Ada]

Listing 58: `abstract_transceivers.ads`

```

1 with Transceivers; use Transceivers;
2
3 package Abstract_Transceivers is
4
5     type Abstract_Transceiver is abstract new Send_Interface and
6       Receive_Interface with null record;
7
8     procedure Send (D : in out Abstract_Transceiver);
9     -- We don't implement Receive for Abstract_Transceiver!
10
11 end Abstract_Transceivers;

```

Listing 59: `abstract_transceivers.adb`

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body Abstract_Transceivers is
4
5     procedure Send (D : in out Abstract_Transceiver) is
6       pragma Unreferenced (D);
7     begin
8       Put_Line ("Sending data...");
9     end Send;
10
11 end Abstract_Transceivers;

```

Listing 60: main.adb

```
1 with Abstract_Transceivers; use Abstract_Transceivers;
2
3 procedure Main is
4   D : Abstract_Transceiver;
5 begin
6   D.Send;
7   D.Receive;
8 end Main;
```

Compilation output

```
main.adb:4:09: error: type of object cannot be abstract
main.adb:7:06: error: call to abstract procedure must be dispatching
```

In this example, we declare the abstract tagged type `Abstract_Transceiver`. Here, we're only partially implementing the interfaces from which this type is derived: we're implementing `Send`, but we're skipping the implementation of `Receive`. Therefore, `Receive` is an abstract operation of `Abstract_Transceiver`. Since any tagged type that has abstract operations is abstract, we must indicate this by adding the `abstract` keyword in type declaration.

Also, when compiling this example, we get an error because we're trying to declare an object of `Abstract_Transceiver` (in the `Main` procedure), which is not possible. Naturally, if we derive another type from `Abstract_Transceiver` and implement `Receive` as well, then we can declare objects of this derived type. This is what we do in the `Full_Transceivers` below:

[Ada]

Listing 61: full_transceivers.ads

```
1 with Abstract_Transceivers; use Abstract_Transceivers;
2
3 package Full_Transceivers is
4
5   type Full_Transceiver is new Abstract_Transceiver with null record;
6   procedure Receive (D : in out Full_Transceiver);
7
8 end Full_Transceivers;
```

Listing 62: full_transceivers.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body Full_Transceivers is
4
5   procedure Receive (D : in out Full_Transceiver) is
6     pragma Unreferenced (D);
7   begin
8     Put_Line ("Receiving data...");
9   end Receive;
10
11 end Full_Transceivers;
```

Listing 63: main.adb

```
1 with Full_Transceivers; use Full_Transceivers;
2
3 procedure Main is
4   D : Full_Transceiver;
5 begin
```

(continues on next page)

(continued from previous page)

```

6   D.Send;
7   D.Receive;
8 end Main;

```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      full_transceivers.adb
  [Ada]      abstract_transceivers.adb
  [Ada]      transceivers.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb

```

Runtime output

```

Sending data...
Receiving data...

```

Here, we implement the Receive procedure for the Full_Transceiver. Therefore, the type doesn't have any abstract operation, so we can use it to declare objects.

7.3.3.8 From simple derivation to OOP

In the [section about simple derivation](#) (page 158), we've seen an example where the actual selection was done at *implementation* time by renaming one of the packages:

```
[Ada]
```

```

with Drivers_1;

package Drivers renames Drivers_1;

```

Although this approach is useful in many cases, there might be situations where we need to select the actual driver dynamically at runtime. Let's look at how we could rewrite that example using interfaces, tagged types and dispatching calls:

```
[Ada]
```

Listing 64: drivers_base.ads

```

1 package Drivers_Base is
2
3   type Transceiver is interface;
4
5   procedure Send (Device : Transceiver; Data : Integer) is abstract;
6   procedure Receive (Device : Transceiver; Data : out Integer) is abstract;
7   procedure Display (Device : Transceiver) is abstract;
8
9 end Drivers_Base;

```

Listing 65: drivers_1.ads

```

1 with Drivers_Base;
2
3 package Drivers_1 is

```

(continues on next page)

(continued from previous page)

```

4  type Transceiver is new Drivers_Base.Transceiver with null record;
5
6  procedure Send (Device : Transceiver; Data : Integer);
7  procedure Receive (Device : Transceiver; Data : out Integer);
8  procedure Display (Device : Transceiver);
9
10
11 end Drivers_1;

```

Listing 66: drivers_1.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  package body Drivers_1 is
4
5      procedure Send (Device : Transceiver; Data : Integer) is null;
6
7      procedure Receive (Device : Transceiver; Data : out Integer) is
8          pragma Unreferenced (Device);
9      begin
10         Data := 42;
11     end Receive;
12
13     procedure Display (Device : Transceiver) is
14         pragma Unreferenced (Device);
15     begin
16         Put_Line ("Using Drivers_1");
17     end Display;
18
19 end Drivers_1;

```

Listing 67: drivers_2.ads

```

1  with Drivers_Base;
2
3  package Drivers_2 is
4
5      type Transceiver is new Drivers_Base.Transceiver with null record;
6
7      procedure Send (Device : Transceiver; Data : Integer);
8      procedure Receive (Device : Transceiver; Data : out Integer);
9      procedure Display (Device : Transceiver);
10
11 end Drivers_2;

```

Listing 68: drivers_2.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  package body Drivers_2 is
4
5      procedure Send (Device : Transceiver; Data : Integer) is null;
6
7      procedure Receive (Device : Transceiver; Data : out Integer) is
8          pragma Unreferenced (Device);
9      begin
10         Data := 7;
11     end Receive;
12
13     procedure Display (Device : Transceiver) is

```

(continues on next page)

(continued from previous page)

```

14     pragma Unreferenced (Device);
15     begin
16         Put_Line ("Using Drivers_2");
17     end Display;
18
19 end Drivers_2;

```

Listing 69: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  with Drivers_Base;
4  with Drivers_1;
5  with Drivers_2;
6
7  procedure Main is
8      D1 : aliased Drivers_1.Transceiver;
9      D2 : aliased Drivers_2.Transceiver;
10     D  : access Drivers_Base.Transceiver'Class;
11
12     I  : Integer;
13
14     type Driver_Number is range 1 .. 2;
15
16     procedure Select_Driver (N : Driver_Number) is
17     begin
18         if N = 1 then
19             D := D1'Access;
20         else
21             D := D2'Access;
22         end if;
23         D.Display;
24     end Select_Driver;
25
26     begin
27         Select_Driver (1);
28         D.Send (999);
29         D.Receive (I);
30         Put_Line (Integer'Image (I));
31
32         Select_Driver (2);
33         D.Send (999);
34         D.Receive (I);
35         Put_Line (Integer'Image (I));
36     end Main;

```

Build output

```

Compile
  [Ada]          main.adb
  [Ada]          drivers_1.adb
  [Ada]          drivers_2.adb
  [Ada]          drivers_base.ads
Bind
  [gprbind]      main.bexch
  [Ada]          main.ali
Link
  [link]         main.adb

```

Runtime output


```
Using Drivers_1
42
Using Drivers_2
7
```

In this example, we declare the `Transceiver` interface in the `Drivers_Base` package. This interface is then used to derive the tagged types `Transceiver` from both `Drivers_1` and `Drivers_2` packages.

In the `Main` procedure, we use the access to `Transceiver'Class` — from the interface declared in the `Drivers_Base` package — to declare `D`. This object `D` contains the access to the actual driver loaded at any specific time. We select the driver at runtime in the inner `Select_Driver` procedure, which initializes `D` (with the access to the selected driver). Then, any operation on `D` triggers a dispatching call to the selected driver.

7.3.3.9 Further resources

In the appendices, we have a step-by-step [hands-on overview of object-oriented programming](#) (page 225) that discusses how to translate a simple system written in C to an equivalent system in Ada using object-oriented programming.

7.3.4 Pointer to subprograms

Pointers to subprograms allow us to dynamically select an appropriate subprogram at runtime. This selection might be triggered by an external event, or simply by the user. This can be useful when multiple versions of a routine exist, and the decision about which one to use cannot be made at compilation time.

This is an example on how to declare and use pointers to functions in C:

[C]

Listing 70: main.c

```
1  #include <stdio.h>
2  #include <stdlib.h>
3
4  void show_msg_v1 (char *msg)
5  {
6      printf("Using version #1: %s\n", msg);
7  }
8
9  void show_msg_v2 (char *msg)
10 {
11     printf("Using version #2:\n %s\n", msg);
12 }
13
14 int main()
15 {
16     int selection = 1;
17     void (*current_show_msg) (char *);
18
19     switch (selection)
20     {
21         case 1: current_show_msg = &show_msg_v1;    break;
22         case 2: current_show_msg = &show_msg_v2;    break;
23         default: current_show_msg = NULL;           break;
24     }
25
```

(continues on next page)

(continued from previous page)

```

26     if (current_show_msg != NULL)
27     {
28         current_show_msg ("Hello there!");
29     }
30     else
31     {
32         printf("ERROR: no version of show_msg() selected!\n");
33     }
34
35     return 0;
36 }

```

Runtime output

```
Using version #1: Hello there!
```

The example above contains two versions of the `show_msg()` function: `show_msg_v1()` and `show_msg_v2()`. The function is selected depending on the value of `selection`, which initializes the function pointer `current_show_msg`. If there's no corresponding value, `current_show_msg` is set to `null` — alternatively, we could have selected a default version of `show_msg()` function. By calling `current_show_msg ("Hello there!")`, we're calling the function that `current_show_msg` is pointing to.

This is the corresponding implementation in Ada:

[Ada]

Listing 71: `show_subprogram_selection.adb`

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Show_Subprogram_Selection is
4
5      procedure Show_Msg_V1 (Msg : String) is
6      begin
7          Put_Line ("Using version #1: " & Msg);
8      end Show_Msg_V1;
9
10     procedure Show_Msg_V2 (Msg : String) is
11     begin
12         Put_Line ("Using version #2: ");
13         Put_Line (Msg);
14     end Show_Msg_V2;
15
16     type Show_Msg_Proc is access procedure (Msg : String);
17
18     Current_Show_Msg : Show_Msg_Proc;
19     Selection        : Natural;
20
21 begin
22     Selection := 1;
23
24     case Selection is
25     when 1 => Current_Show_Msg := Show_Msg_V1'Access;
26     when 2 => Current_Show_Msg := Show_Msg_V2'Access;
27     when others => Current_Show_Msg := null;
28     end case;
29
30     if Current_Show_Msg /= null then
31         Current_Show_Msg ("Hello there!");
32     else

```

(continues on next page)

(continued from previous page)

```

33     Put_Line ("ERROR: no version of Show_Msg selected!");
34 end if;
35
36 end Show_Subprogram_Selection;

```

Build output

```

Compile
  [Ada]          show_subprogram_selection.adb
Bind
  [gprbind]      show_subprogram_selection.bexch
  [Ada]          show_subprogram_selection.ali
Link
  [link]         show_subprogram_selection.adb

```

Runtime output

```
Using version #1: Hello there!
```

The structure of the code above is very similar to the one used in the C code. Again, we have two version of Show_Msg: Show_Msg_V1 and Show_Msg_V2. We set Current_Show_Msg according to the value of Selection. Here, we use 'Access to get access to the corresponding procedure. If no version of Show_Msg is available, we set Current_Show_Msg to null.

Pointers to subprograms are also typically used as callback functions. This approach is extensively used in systems that process events, for example. Here, we could have a two-layered system:

- A layer of the system (an event manager) triggers events depending on information from sensors.
 - For each event, callback functions can be registered.
 - The event manager calls registered callback functions when an event is triggered.
- Another layer of the system registers callback functions for specific events and decides what to do when those events are triggered.

This approach promotes information hiding and component decoupling because:

- the layer of the system responsible for managing events doesn't need to know what the callback function actually does, while
- the layer of the system that implements callback functions remains agnostic to implementation details of the event manager — for example, how events are implemented in the event manager.

Let's see an example in C where we have a process_values () function that calls a callback function (process_one) to process a list of values:

[C]

Listing 72: process_values.h

```

1 typedef int (*process_one_callback) (int);
2
3 void process_values (int          *values,
4                     int          len,
5                     process_one_callback process_one);

```

Listing 73: process_values.c

```

1 #include "process_values.h"
2

```

(continues on next page)

(continued from previous page)

```

3  #include <assert.h>
4  #include <stdio.h>
5
6  void process_values (int          *values,
7                      int          len,
8                      process_one_callback process_one)
9  {
10     int i;
11
12     assert (process_one != NULL);
13
14     for (i = 0; i < len; i++)
15     {
16         values[i] = process_one (values[i]);
17     }
18 }

```

Listing 74: main.c

```

1  #include <stdio.h>
2  #include <stdlib.h>
3
4  #include "process_values.h"
5
6  int proc_10 (int val)
7  {
8      return val + 10;
9  }
10
11 # define LEN_VALUES      5
12
13 int main()
14 {
15
16     int values[LEN_VALUES] = { 1, 2, 3, 4, 5 };
17     int i;
18
19     process_values (values, LEN_VALUES, &proc_10);
20
21     for (i = 0; i < LEN_VALUES; i++)
22     {
23         printf("Value [%d] = %d\n", i, values[i]);
24     }
25
26     return 0;
27 }

```

Runtime output

```

Value [0] = 11
Value [1] = 12
Value [2] = 13
Value [3] = 14
Value [4] = 15

```

As mentioned previously, `process_values()` doesn't have any knowledge about what `process_one()` does with the integer value it receives as a parameter. Also, we could replace `proc_10()` by another function without having to change the implementation of `process_values()`.

Note that `process_values()` calls an `assert()` for the function pointer to compare it against

null. Here, instead of checking the validity of the function pointer, we're expecting the caller of `process_values()` to provide a valid pointer.

This is the corresponding implementation in Ada:

[Ada]

Listing 75: values_processing.ads

```
1 package Values_Processing is
2
3     type Integer_Array is array (Positive range <>) of Integer;
4
5     type Process_One_Callback is not null access
6         function (Value : Integer) return Integer;
7
8     procedure Process_Values (Values      : in out Integer_Array;
9                             Process_One :      Process_One_Callback);
10
11 end Values_Processing;
```

Listing 76: values_processing.adb

```
1 package body Values_Processing is
2
3     procedure Process_Values (Values      : in out Integer_Array;
4                             Process_One :      Process_One_Callback) is
5     begin
6         for I in Values'Range loop
7             Values (I) := Process_One (Values (I));
8         end loop;
9     end Process_Values;
10
11 end Values_Processing;
```

Listing 77: proc_10.ads

```
1 function Proc_10 (Value : Integer) return Integer;
```

Listing 78: proc_10.adb

```
1 function Proc_10 (Value : Integer) return Integer is
2 begin
3     return Value + 10;
4 end Proc_10;
```

Listing 79: show_callback.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 with Values_Processing; use Values_Processing;
4 with Proc_10;
5
6 procedure Show_Callback is
7     Values : Integer_Array := (1, 2, 3, 4, 5);
8 begin
9     Process_Values (Values, Proc_10'Access);
10
11     for I in Values'Range loop
12         Put_Line ("Value ["
13                 & Positive'Image (I)
14                 & "] = "
```

(continues on next page)

(continued from previous page)

```

15         & Integer'Image (Values (I)));
16     end loop;
17 end Show_Callback;

```

Build output

```

Compile
  [Ada]          show_callback.adb
  [Ada]          proc_10.adb
  [Ada]          values_processing.adb
Bind
  [gprbind]      show_callback.bexch
  [Ada]          show_callback.ali
Link
  [link]         show_callback.adb

```

Runtime output

```

Value [ 1] = 11
Value [ 2] = 12
Value [ 3] = 13
Value [ 4] = 14
Value [ 5] = 15

```

Similar to the implementation in C, the `Process_Values` procedure receives the access to a callback routine, which is then called for each value of the `Values` array.

Note that the declaration of `Process_One_Callback` makes use of the `not null` access declaration. By using this approach, we ensure that any parameter of this type has a valid value, so we can always call the callback routine.

7.4 Design by components using dynamic libraries

In the previous sections, we have shown how to use packages to create separate components of a system. As we know, when designing a complex system, it is advisable to separate concerns into distinct units, so we can use Ada packages to represent each unit of a system. In this section, we go one step further and create separate dynamic libraries for each component, which we'll then link to the main application.

Let's suppose we have a main system (`Main_System`) and a component A (`Component_A`) that we want to use in the main system. For example:

[Ada]

Listing 80: `component_a.ads`

```

1  --
2  -- File: component_a.ads
3  --
4  package Component_A is
5
6      type Float_Array is array (Positive range <>) of Float;
7
8      function Average (Data : Float_Array) return Float;
9
10 end Component_A;

```

Listing 81: component_a.adb

```
1  --
2  --   File: component_a.adb
3  --
4  package body Component_A is
5
6      function Average (Data : Float_Array) return Float is
7          Total : Float := 0.0;
8      begin
9          for Value of Data loop
10             Total := Total + Value;
11          end loop;
12          return Total / Float (Data'Length);
13      end Average;
14
15 end Component_A;
```

Listing 82: main_system.adb

```
1  --
2  --   File: main_system.adb
3  --
4  with Ada.Text_IO; use Ada.Text_IO;
5
6  with Component_A; use Component_A;
7
8  procedure Main_System is
9      Values      : constant Float_Array := (10.0, 11.0, 12.0, 13.0);
10     Average_Value : Float;
11 begin
12     Average_Value := Average (Values);
13     Put_Line ("Average = " & Float'Image (Average_Value));
14 end Main_System;
```

Build output

```
Compile
  [Ada]      main_system.adb
  [Ada]      component_a.adb
Bind
  [gprbind]  main_system.bexch
  [Ada]      main_system.ali
Link
  [link]     main_system.adb
```

Runtime output

```
Average =  1.15000E+01
```

Note that, in the source-code example above, we're indicating the name of each file. We'll now see how to organize those files in a structure that is suitable for the GNAT build system (**GPRbuild**).

In order to discuss how to create dynamic libraries, we need to dig into some details about the build system. With GNAT, we can use project files for **GPRbuild** to easily design dynamic libraries. Let's say we use the following directory structure for the code above:

```
| - component_a
|   | component_a.gpr
|   | - src
|   |   | component_a.adb
```

(continues on next page)

(continued from previous page)

```

|      |      | component_a.ads
|- main_system
|      |      | main_system.gpr
|      |      |- src
|      |      | main_system.adb

```

Here, we have two directories: *component_a* and *main_system*. Each directory contains a project file (with the *.gpr* file extension) and a source-code directory (*src*).

In the source-code example above, we've seen the content of files *component_a.ads*, *component_a.adb* and *main_system.adb*. Now, let's discuss how to write the project file for *Component_A* (*component_a.gpr*), which will build the dynamic library for this component:

```

library project Component_A is

  for Source_Dirs use ("src");
  for Object_Dir use "obj";
  for Create_Missing_Dirs use "True";
  for Library_Name use "component_a";
  for Library_Kind use "dynamic";
  for Library_Dir use "lib";

end Component_A;

```

The project is defined as a *library project* instead of *project*. This tells **GPRbuild** to build a library instead of an executable binary. We then specify the library name using the *Library_Name* attribute, which is required, so it must appear in a library project. The next two library-related attributes are optional, but important for our use-case. We use:

- *Library_Kind* to specify that we want to create a dynamic library — by default, this attribute is set to *static*;
- *Library_Dir* to specify the directory where the library is stored.

In the project file of our main system (*main_system.gpr*), we just need to reference the project of *Component_A* using a *with* clause and indicating the correct path to that project file:

```

with "../component_a/component_a.gpr";

project Main_System is
  for Source_Dirs use ("src");
  for Object_Dir use "obj";
  for Create_Missing_Dirs use "True";
  for Main use ("main_system.adb");
end Main_System;

```

GPRbuild takes care of selecting the correct settings to link the dynamic library created for *Component_A* with the main application (*Main_System*) and build an executable.

We can use the same strategy to create a *Component_B* and dynamically link to it in the *Main_System*. We just need to create the separate structure for this component — with the appropriate Ada packages and project file — and include it in the project file of the main system using a *with* clause:

```

with "../component_a/component_a.gpr";
with "../component_b/component_b.gpr";

...

```

Again, **GPRbuild** takes care of selecting the correct settings to link both dynamic libraries together with the main application.

You can find more details and special setting for library projects in the [GPRbuild documentation](#)²².

In the GNAT toolchain

The GNAT toolchain includes a more advanced example focusing on how to load dynamic libraries at runtime. You can find it in the `share/examples/gnat/plugins` directory of the GNAT toolchain installation. As described in the README file from that directory, this example "comprises a main program which probes regularly for the existence of shared libraries in a known location. If such libraries are present, it uses them to implement features initially not present in the main program."

²² https://docs.adacore.com/gprbuild-docs/html/gprbuild_ug/gnat_project_manager.html#library-projects

PERFORMANCE CONSIDERATIONS

8.1 Overall expectations

All in all, there should not be significant performance differences between code written in Ada and code written in C, provided that they are semantically equivalent. Taking the current GNAT implementation and its GCC C counterpart for example, most of the code generation and optimization phases are shared between C and Ada — so there's not one compiler more efficient than the other. Furthermore, the two languages are fairly similar in the way they implement imperative semantics, in particular with regards to memory management or control flow. They should be equivalent on average.

When comparing the performance of C and Ada code, differences might be observed. This usually comes from the fact that, while the two piece *appear* semantically equivalent, they happen to be actually quite different; C code semantics do not implicitly apply the same run-time checks that Ada does. This section will present common ways for improving Ada code performance.

8.2 Switches and optimizations

Clever use of compilation switches might optimize the performance of an application significantly. In this section, we'll briefly look into some of the switches available in the GNAT toolchain.

8.2.1 Optimizations levels

Optimization levels can be found in many compilers for multiple languages. On the lowest level, the GNAT compiler doesn't optimize the code at all, while at the higher levels, the compiler analyses the code and optimizes it by removing unnecessary operations and making the most use of the target processor's capabilities.

By being part of GCC, GNAT offers the same `-O_` switches as GCC:

Switch	Description
<code>-O0</code>	No optimization: the generated code is completely unoptimized. This is the default optimization level.
<code>-O1</code>	Moderate optimization.
<code>-O2</code>	Full optimization.
<code>-O3</code>	Same optimization level as for <code>-O2</code> . In addition, further optimization strategies, such as aggressive automatic inlining and vectorization.

Note that the higher the level, the longer the compilation time. For fast compilation during development phase, unless you're working on benchmarking algorithms, using `-O0` is probably a good idea.

In addition to the levels presented above, GNAT also has the `-Os` switch, which allows for optimizing code and data usage.

8.2.2 Inlining

As we've seen in the previous section, automatic inlining depends on the optimization level. The highest optimization level (`-O3`), for example, performs aggressive automatic inlining. This could mean that this level inlines too much rather than not enough. As a result, the cache may become an issue and the overall performance may be worse than the one we would achieve by compiling the same code with optimization level 2 (`-O2`). Therefore, the general recommendation is to not *just* select `-O3` for the optimized version of an application, but instead compare it the optimized version built with `-O2`.

In some cases, it's better to reduce the optimization level and perform manual inlining instead of automatic inlining. We do that by using the `Inline` aspect. Let's reuse an example from a previous chapter and inline the `Average` function:

[Ada]

Listing 1: float_arrays.ads

```
1 package Float_Arrays is
2
3     type Float_Array is array (Positive range <>) of Float;
4
5     function Average (Data : Float_Array) return Float
6         with Inline;
7
8 end Float_Arrays;
```

Listing 2: float_arrays.adb

```
1 package body Float_Arrays is
2
3     function Average (Data : Float_Array) return Float is
4         Total : Float := 0.0;
5     begin
6         for Value of Data loop
7             Total := Total + Value;
8         end loop;
9         return Total / Float (Data'Length);
10    end Average;
11
12 end Float_Arrays;
```

Listing 3: compute_average.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 with Float_Arrays; use Float_Arrays;
4
5 procedure Compute_Average is
6     Values      : constant Float_Array := (10.0, 11.0, 12.0, 13.0);
7     Average_Value : Float;
8 begin
9     Average_Value := Average (Values);
10    Put_Line ("Average = " & Float'Image (Average_Value));
11 end Compute_Average;
```

Build output

```

Compile
  [Ada]          compute_average.adb
  [Ada]          float_arrays.adb
Bind
  [gprbind]      compute_average.bexch
  [Ada]          compute_average.ali
Link
  [link]         compute_average.adb

```

Runtime output

```
Average = 1.15000E+01
```

When compiling this example, GNAT will inline `Average` in the `Compute_Average` procedure.

In order to effectively use this aspect, however, we need to set the optimization level to at least `-O1` and use the `-gnatn` switch, which instructs the compiler to take the `Inline` aspect into account.

Note, however, that the `Inline` aspect is just a *recommendation* to the compiler. Sometimes, the compiler might not be able to follow this recommendation, so it won't inline the subprogram. In this case, we get a compilation warning from GNAT.

These are some examples of situations where the compiler might not be able to inline a subprogram:

- when the code is too large,
- when it's too complicated — for example, when it involves exception handling —, or
- when it contains tasks, etc.

In addition to the `Inline` aspect, we also have the `Inline_Always` aspect. In contrast to the former aspect, however, the `Inline_Always` aspect isn't primarily related to performance. Instead, it should be used when the functionality would be incorrect if inlining was not performed by the compiler. Examples of this are procedures that insert Assembly instructions that only make sense when the procedure is inlined, such as memory barriers.

Similar to the `Inline` aspect, there might be situations where a subprogram has the `Inline_Always` aspect, but the compiler is unable to inline it. In this case, we get a compilation error from GNAT.

8.3 Checks and assertions

8.3.1 Checks

Ada provides many runtime checks to ensure that the implementation is working as expected. For example, when accessing an array, we would like to make sure that we're not accessing a memory position that is not allocated for that array. This is achieved by an index check.

Another example of runtime check is the verification of valid ranges. For example, when adding two integer numbers, we would like to ensure that the result is still in the valid range — that the value is neither too large nor too small. This is achieved by a range check. Likewise, arithmetic operations shouldn't overflow or underflow. This is achieved by an overflow check.

Although runtime checks are very useful and should be used as much as possible, they can also increase the overhead of implementations at certain hot-spots. For example, checking the index of an array in a sorting algorithm may significantly decrease its performance. In those cases, suppressing the check may be an option. We can achieve this suppression by using `pragma Suppress (Index_Check)`. For example:

```
[Ada]
```

```
procedure Sort (A : in out Integer_Array) is
  pragma Suppress (Index_Check);
begin
  -- (implementation removed...)
  null;
end Sort;
```

In case of overflow checks, we can use `pragma Suppress (Overflow_Check)` to suppress them:

```
function Some_Computation (A, B : Int32) return Int32 is
  pragma Suppress (Overflow_Check);
begin
  -- (implementation removed...)
  null;
end Sort;
```

We can also deactivate overflow checks for integer types using the `-gnato` switch when compiling a source-code file with GNAT. In this case, overflow checks in the whole file are deactivated.

It is also possible to suppress all checks at once using `pragma Suppress (All_Checks)`. In addition, GNAT offers a compilation switch called `-gnatp`, which has the same effect on the whole file.

Note, however, that this kind of suppression is just a recommendation to the compiler. There's no guarantee that the compiler will actually suppress any of the checks because the compiler may not be able to do so — typically because the hardware happens to do it. For example, if the machine traps on any access via address zero, requesting the removal of null access value checks in the generated code won't prevent the checks from happening.

It is important to differentiate between required and redundant checks. Let's consider the following example in C:

[C]

Listing 4: main.c

```
1  #include <stdio.h>
2
3  int main(int argc, const char * argv[])
4  {
5      int a = 8, b = 0, res;
6
7      res = a / b;
8
9      // printing the result
10     printf("res = %d\n", res);
11
12     return 0;
13 }
```

Because C doesn't have language-defined checks, as soon as the application tries to divide a value by zero in `res = a / b`, it'll break — on Linux, for example, you may get the following error message by the operating system: `Floating point exception (core dumped)`. Therefore, we need to manually introduce a check for zero before this operation. For example:

[C]

Listing 5: main.c

```
1  #include <stdio.h>
2
3  int main(int argc, const char * argv[])
4  {
```

(continues on next page)

(continued from previous page)

```

5      int a = 8, b = 0, res;
6
7      if (b != 0) {
8          res = a / b;
9
10         // printing the result
11         printf("res = %d\n", res);
12     }
13     else
14     {
15         // printing error message
16         printf("Error: cannot calculate value (division by zero)\n");
17     }
18
19     return 0;
20 }

```

Runtime output

```
Error: cannot calculate value (division by zero)
```

This is the corresponding code in Ada:

[Ada]

Listing 6: show_division_by_zero.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Show_Division_By_Zero is
4      A   : Integer := 8;
5      B   : Integer := 0;
6      Res : Integer;
7  begin
8      Res := A / B;
9
10     Put_Line ("Res = " & Integer'Image (Res));
11 end Show_Division_By_Zero;

```

Build output

```

Compile
  [Ada]          show_division_by_zero.adb
show_division_by_zero.adb:8:13: warning: division by zero [enabled by default]
show_division_by_zero.adb:8:13: warning: "Constraint_Error" will be raised at run_
time [enabled by default]
Bind
  [gprbind]      show_division_by_zero.bexch
  [Ada]          show_division_by_zero.ali
Link
  [link]         show_division_by_zero.adb

```

Runtime output

```
raised CONSTRAINT_ERROR : show_division_by_zero.adb:8 divide by zero
```

Similar to the first version of the C code, we're not explicitly checking for a potential division by zero here. In Ada, however, this check is *automatically inserted* by the language itself. When running the application above, an exception is raised when the application tries to divide the value in A by zero. We could introduce exception handling in our example, so that we get the same message as we did in the second version of the C code:

[Ada]

Listing 7: show_division_by_zero.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Division_By_Zero is
4   A   : Integer := 8;
5   B   : Integer := 0;
6   Res : Integer;
7 begin
8   Res := A / B;
9
10  Put_Line ("Res = " & Integer'Image (Res));
11 exception
12   when Constraint_Error =>
13     Put_Line ("Error: cannot calculate value (division by zero)");
14   when others =>
15     null;
16 end Show_Division_By_Zero;
```

Build output

```
Compile
  [Ada]          show_division_by_zero.adb
show_division_by_zero.adb:8:13: warning: division by zero [enabled by default]
show_division_by_zero.adb:8:13: warning: "Constraint_Error" will be raised at run_
↳time [enabled by default]
Bind
  [gprbind]      show_division_by_zero.bexch
  [Ada]          show_division_by_zero.ali
Link
  [link]         show_division_by_zero.adb
```

Runtime output

```
Error: cannot calculate value (division by zero)
```

This example demonstrates that the division check for `Res := A / B` is required and shouldn't be suppressed. In contrast, a check is redundant — and therefore not required — when we know that the condition that leads to a failure can never happen. In many cases, the compiler itself detects redundant checks and eliminates them (for higher optimization levels). Therefore, when improving the performance of your application, you should:

1. keep all checks active for most parts of the application;
2. identify the hot-spots of your application;
3. identify which checks haven't been eliminated by the optimizer on these hot-spots;
4. identify which of those checks are redundant;
5. only suppress those checks that are redundant, and keep the required ones.

8.3.2 Assertions

We've already discussed assertions in [this section of the SPARK chapter](#) (page 113). Assertions are user-defined checks that you can add to your code using the pragma `Assert`. For example:

[Ada]

```
function Some_Computation (A, B : Int32) return Int32 is
  Res : Int32;
begin
  -- (implementation removed...)

  pragma Assert (Res >= 0);

  return Res;
end Sort;
```

Assertions that are specified with pragma `Assert` are not enabled by default. You can enable them by setting the assertion policy to *check* — using pragma `Assertion_Policy (Check)` — or by using the `-gnata` switch when compiling with GNAT.

Similar to the checks discussed previously, assertions can generate significant overhead when used at hot-spots. Restricting those assertions to development (e.g. debug version) and turning them off on the release version may be an option. In this case, formal proof — as discussed in the [SPARK chapter](#) (page 107) — can help you. By formally proving that assertions will never fail at run-time, you can safely deactivate them.

8.4 Dynamic vs. static structures

Ada generally speaking provides more ways than C or C++ to write simple dynamic structures, that is to say structures that have constraints computed after variables. For example, it's quite typical to have initial values in record types:

[Ada]

```
type R is record
  F : Some_Field := Call_To_Some_Function;
end record;
```

However, the consequences of the above is that any declaration of a instance of this type without an explicit value for `V` will issue a call to `Call_To_Some_Function`. More subtle issue may arise with elaboration. For example, it's possible to write:

Listing 8: some_functions.ads

```
1 package Some_Functions is
2
3   function Some_Function_Call return Integer is (2);
4
5   function Some_Other_Function_Call return Integer is (10);
6
7 end Some_Functions;
```

Listing 9: values.ads

```
1 with Some_Functions; use Some_Functions;
2
3 package Values is
4   A_Start : Integer := Some_Function_Call;
```

(continues on next page)

(continued from previous page)

```

5   A_End    : Integer := Some_Other_Function_Call;
6 end Values;

```

Listing 10: arr_def.ads

```

1 with Values; use Values;
2
3 package Arr_Def is
4   type Arr is array (Integer range A_Start .. A_End) of Integer;
5 end Arr_Def;

```

It may indeed be appealing to be able to change the values of A_Start and A_End at startup so as to align a series of arrays dynamically. The consequence, however, is that these values will not be known statically, so any code that needs to access to boundaries of the array will need to read data from memory. While it's perfectly fine most of the time, there may be situations where performances are so critical that static values for array boundaries must be enforced.

Here's a last case which may also be surprising:

[Ada]

Listing 11: arr_def.ads

```

1 package Arr_Def is
2   type Arr is array (Integer range <>) of Integer;
3
4   type R (D1, D2 : Integer) is record
5     F1 : Arr (1 .. D1);
6     F2 : Arr (1 .. D2);
7   end record;
8 end Arr_Def;

```

In the code above, R contains two arrays, F1 and F2, respectively constrained by the discriminant D1 and D2. The consequence is, however, that to access F2, the run-time needs to know how large F1 is, which is dynamically constrained when creating an instance. Therefore, accessing to F2 requires a computation involving D1 which is slower than, let's say, two pointers in an C array that would point to two different arrays.

Generally speaking, when values are used in data structures, it's useful to always consider where they're coming from, and if their value is static (computed by the compiler) or dynamic (only known at run-time). There's nothing fundamentally wrong with dynamically constrained types, unless they appear in performance-critical pieces of the application.

8.5 Pointers vs. data copies

In the section about [pointers](#) (page 132), we mentioned that the Ada compiler will automatically pass parameters by reference when needed. Let's look into what "when needed" means. The fundamental point to understand is that the parameter types determine how the parameters are passed in and/or out. The parameter modes do not control how parameters are passed.

Specifically, the language standards specifies that scalar types are always passed by value, and that some other types are always passed by reference. It would not make sense to make a copy of a task when passing it as a parameter, for example. So parameters that can be passed reasonably by value will be, and those that must be passed by reference will be. That's the safest approach.

But the language also specifies that when the parameter is an array type or a record type, and the record/array components are all by-value types, then the compiler decides: it can pass the parameter using either mechanism. The critical case is when such a parameter is large, e.g., a large matrix. We don't want the compiler to pass it by value because that would entail a large copy,

and indeed the compiler will not do so. But if the array or record parameter is small, say the same size as an address, then it doesn't matter how it is passed and by copy is just as fast as by reference. That's why the language gives the choice to the compiler. Although the language does not mandate that large parameters be passed by reference, any reasonable compiler will do the right thing.

The modes do have an effect, but not in determining how the parameters are passed. Their effect, for parameters passed by value, is to determine how many times the value is copied. For mode `in` and mode `out` there is just one copy. For mode `in out` there will be two copies, one in each direction.

Therefore, unlike C, you don't have to use access types in Ada to get better performance when passing arrays or records to subprograms. The compiler will almost certainly do the right thing for you.

Let's look at this example:

[C]

Listing 12: main.c

```

1  #include <stdio.h>
2
3  struct Data {
4      int prev, curr;
5  };
6
7  void update(struct Data *d,
8             int v)
9  {
10     d->prev = d->curr;
11     d->curr = v;
12 }
13
14 void display(const struct Data *d)
15 {
16     printf("Prev : %d\n", d->prev);
17     printf("Curr : %d\n", d->curr);
18 }
19
20 int main(int argc, const char * argv[])
21 {
22     struct Data D1 = { 0, 1 };
23
24     update (&D1, 3);
25     display (&D1);
26
27     return 0;
28 }
```

Runtime output

```

Prev : 1
Curr : 3
```

In this C code example, we're using pointers to pass `D1` as a reference to `update` and `display`. In contrast, the equivalent code in Ada simply uses the parameter modes to specify the data flow directions. The mechanisms used to pass the values do not appear in the source code.

[Ada]

Listing 13: update_record.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Update_Record is
4
5     type Data is record
6         Prev : Integer;
7         Curr : Integer;
8     end record;
9
10    procedure Update (D : in out Data;
11                     V : Integer) is
12    begin
13        D.Prev := D.Curr;
14        D.Curr := V;
15    end Update;
16
17    procedure Display (D : Data) is
18    begin
19        Put_Line ("Prev: " & Integer'Image (D.Prev));
20        Put_Line ("Curr: " & Integer'Image (D.Curr));
21    end Display;
22
23    D1 : Data := (0, 1);
24
25    begin
26        Update (D1, 3);
27        Display (D1);
28    end Update_Record;

```

Build output

Compile	
[Ada]	update_record.adb
Bind	
[gprbind]	update_record.bexch
[Ada]	update_record.ali
Link	
[link]	update_record.adb

Runtime output

```

Prev: 1
Curr: 3

```

In the calls to `Update` and `Display`, `D1` is always be passed by reference. Because no extra copy takes place, we get a performance that is equivalent to the C version. If we had used arrays in the example above, `D1` would have been passed by reference as well:

[Ada]

Listing 14: update_array.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Update_Array is
4
5     type Data_State is (Prev, Curr);
6     type Data is array (Data_State) of Integer;
7
8     procedure Update (D : in out Data;
9                     V : Integer) is

```

(continues on next page)

(continued from previous page)

```

10  begin
11      D (Prev) := D (Curr);
12      D (Curr) := V;
13  end Update;
14
15  procedure Display (D : Data) is
16  begin
17      Put_Line ("Prev: " & Integer'Image (D (Prev)));
18      Put_Line ("Curr: " & Integer'Image (D (Curr)));
19  end Display;
20
21  D1 : Data := (0, 1);
22
23  begin
24      Update (D1, 3);
25      Display (D1);
26  end Update_Array;

```

Build output

```

Compile
  [Ada]          update_array.adb
Bind
  [gprbind]      update_array.bexch
  [Ada]          update_array.ali
Link
  [link]         update_array.adb

```

Runtime output

```

Prev:  1
Curr:  3

```

Again, no extra copy is performed in the calls to `Update` and `Display`, which gives us optimal performance when dealing with arrays and avoids the need to use access types to optimize the code.

8.5.1 Function returns

Previously, we've discussed the cost of passing complex records as arguments to subprograms. We've seen that we don't have to use explicit access type parameters to get better performance in Ada. In this section, we'll briefly discuss the cost of function returns.

In general, we can use either procedures or functions to initialize a data structure. Let's look at this example in C:

[C]

Listing 15: main.c

```

1  #include <stdio.h>
2
3  struct Data {
4      int prev, curr;
5  };
6
7  void init_data(struct Data *d)
8  {
9      d->prev = 0;

```

(continues on next page)

(continued from previous page)

```

10     d->curr = 1;
11 }
12
13 struct Data get_init_data()
14 {
15     struct Data d = { 0, 1 };
16
17     return d;
18 }
19
20 int main(int argc, const char * argv[])
21 {
22     struct Data D1;
23
24     D1 = get_init_data();
25
26     init_data(&D1);
27
28     return 0;
29 }

```

This code example contains two subprograms that initialize the Data structure:

- `init_data()`, which receives the data structure as a reference (using a pointer) and initializes it, and
- `get_init_data()`, which returns the initialized structure.

In C, we generally avoid implementing functions such as `get_init_data()` because of the extra copy that is needed for the function return.

This is the corresponding implementation in Ada:

[Ada]

Listing 16: init_record.adb

```

1  procedure Init_Record is
2
3      type Data is record
4          Prev : Integer;
5          Curr : Integer;
6      end record;
7
8      procedure Init (D : out Data) is
9      begin
10         D := (Prev => 0, Curr => 1);
11     end Init;
12
13     function Init return Data is
14         D : constant Data := (Prev => 0, Curr => 1);
15     begin
16         return D;
17     end Init;
18
19     D1 : Data;
20
21     pragma Unreferenced (D1);
22 begin
23     D1 := Init;
24
25     Init (D1);
26 end Init_Record;

```

Build output

```

Compile
  [Ada]          init_record.adb
Bind
  [gprbind]      init_record.bexch
  [Ada]          init_record.ali
Link
  [link]         init_record.adb

```

In this example, we have two versions of `Init`: one using a procedural form, and the other one using a functional form. Note that, because of Ada's support for subprogram overloading, we can use the same name for both subprograms.

The issue is that assignment of a function result entails a copy, just as if we assigned one variable to another. For example, when assigning a function result to a constant, the function result is copied into the memory for the constant. That's what is happening in the above examples for the initialized variables.

Therefore, in terms of performance, the same recommendations apply: for large types we should avoid writing functions like the `Init` function above. Instead, we should use the procedural form of `Init`. The reason is that the compiler necessarily generates a copy for the `Init` function, while the `Init` procedure uses a reference for the output parameter, so that the actual record initialization is performed in place in the caller's argument.

An exception to this is when we use functions returning values of limited types, which by definition do not allow assignment. Here, to avoid allowing something that would otherwise look suspiciously like an assignment, the compiler generates the function body so that it builds the result directly into the object being assigned. No copy takes place.

We could, for example, rewrite the example above using limited types:

[Ada]

Listing 17: `init_limited_record.adb`

```

1  procedure Init_Limited_Record is
2
3      type Data is limited record
4          Prev : Integer;
5          Curr : Integer;
6      end record;
7
8      function Init return Data is
9      begin
10         return D : Data do
11             D.Prev := 0;
12             D.Curr := 1;
13         end return;
14     end Init;
15
16     D1 : Data := Init;
17
18     pragma Unreferenced (D1);
19 begin
20     null;
21 end Init_Limited_Record;

```

Build output

```

Compile
  [Ada]          init_limited_record.adb
Bind

```

(continues on next page)

(continued from previous page)

[gprbind]	init_limited_record.bexch
[Ada]	init_limited_record.ali
Link	
[link]	init_limited_record.adb

In this example, `D1 : Data := Init;` has the same cost as the call to the procedural form — `Init (D1);` — that we've seen in the previous example. This is because the assignment is done in place.

Note that limited types require the use of the extended return statements (`return ... do ... end return`) in function implementations. Also note that, because the `Data` type is limited, we can only use the `Init` function in the declaration of `D1`; a statement in the code such as `D1 := Init;` is therefore forbidden.

ARGUMENTATION AND BUSINESS PERSPECTIVES

The technical benefits of a migration from C to Ada are usually relatively straightforward to demonstrate. Hopefully, this course provides a good basis for it. However, when faced with an actual business decision to make, additional considerations need to be taken into account, such as return on investment, perennity of the solution, tool support, etc. This section will cover a number of usual questions and provide elements of answers.

9.1 What's the expected ROI of a C to Ada transition?

Switching from one technology to another is a cost, may that be in terms of training, transition of the existing environment or acquisition of new tools. This investment needs to be matched with an expected return on investment, or ROI, to be consistent. Of course, it's incredibly difficult to provide a firm answer to how much money can be saved by transitioning, as this is highly dependent on specific project objectives and constraints. We're going to provide qualitative and quantitative arguments here, from the perspective of a project that has to reach a relatively high level of integrity, that is to say a system where the occurrence of a software failure is a relatively costly event.

From a qualitative standpoint, there are various times in the software development life cycle where defects can be found:

1. on the developer's desk
2. during component testing
3. during integration testing
4. after deployment
5. during maintenance

Numbers from studies vary greatly on the relative costs of defects found at each of these phases, but there's a clear ordering between them. For example, a defect found while developing is orders of magnitude less expensive to fix than a defect found e.g. at integration time, which may involve costly debugging sessions and slow down the entire system acceptance. The whole purpose of Ada and SPARK is to push defect detection to the developer's desk as much as possible; at least for all of these defects that can be identified at that level. While the strict act of writing software may be taking more effort because of all of the additional safeguards, this should have a significant and positive impact down the line and help to control costs overall. The exact value this may translate into is highly business dependent.

From a quantitative standpoint, two studies have been done almost 25 years apart and provide similar insights:

- Rational Software in 1995 found that the cost of developing software in Ada was overall half as much as the cost of developing software in C.
- VDC ran a study in 2018, finding that the cost savings of developing with Ada over C ranged from 6% to 38% in savings.

From a qualitative standpoint, in particular with regards to Ada and C from a formal proof perspective, an interesting presentation was made in 2017 by two researchers. They tried to apply formal proof on the same piece of code, developed in Ada/SPARK on one end and C/Frama-C on the other. Their results indicate that the Ada/SPARK technology is indeed more conducive to formal proof methodologies.

Although all of these studies have their own biases, they provide a good idea of what to expect in terms of savings once the initial investment in switching to Ada is made. This is assuming everything else is equal, in particular that the level of integrity is the same. In many situations, the migration to Ada is justified by an increase in terms of integrity expectations, in which case it's expected that development costs will rise (it's more expensive to develop better software) and Ada is viewed as a means to mitigate this rise in development costs.

That being said, the point of this argument is not to say that it's not possible to write very safe and secure software with languages different than Ada. With the right expertise, the right processes and the right tools, it's done every day. The point is that Ada overall reduces the level of processes, expertise and tools necessary and will allow to reach the same target at a lower cost.

9.2 Who is using Ada today?

Ada was initially born as a DoD project, and thus got its initial customer base in aerospace and defence (A&D). At the time these lines are written and from the perspective of AdaCore, A&D is still the largest consumer of Ada today and covers about 70% of the market. This creates a consistent and long lasting set of established users as these project last often for decades, using the same codebase migrating from platform to platform.

More recently however, there has been an emerging interest for Ada in new communities of users such as automotive, medical device, industrial automation and overall cyber-security. This can probably be explained by a rise of safety, reliability and cyber-security requirements. The market is moving relatively rapidly today and we're anticipating an increase of the Ada footprint in these domains, while still remaining a technology of choice for the development of mission critical software.

9.3 What is the future of the Ada technology?

The first piece of the answer lies in the user base of the Ada language, as seen in the previous question. Projects using Ada in the aerospace and defence domain maintain source code over decades, providing healthy funding foundation for Ada-based technologies.

AdaCore being the author of this course, it's difficult for us to be fair in our description of other Ada compilation technologies. We will leave to the readers the responsibility of forging their own opinion. If they present a credible alternative to the GNAT compiler, then this whole section can be considered as void.

Assuming GNAT is the only option available, and acknowledging that this is an argument that we're hearing from a number of Ada adopters, let's discuss the "sole source" issue.

First of all, it's worth noting that industries are using a lot of software that is provided by only one source, so while non-ideal, these situations are also quite common.

In the case of the GNAT compiler however, while AdaCore is the main maintainer, this maintenance is done as part of an open-source community. This means that nothing prevents a third party to start selling a competing set of products based on the same compiler, provided that it too adopts the open-source approach. Our job is to be more cost-effective than the alternative, and indeed for the vast part this has prevented a competing offering to emerge. However, should AdaCore disappear or switch focus, Ada users would not be prevented from carrying on using its software (there is no lock) and a third party could take over maintenance. This is not a theoretical case, this

has been done in the past either by companies looking at supporting their own version of GNAT, vendors occupying a specific niche that was left uncovered, or hobbyists developing their own builds.

With that in mind, it's clear that the "sole source" provider issue is a circumstantial — nothing is preventing other vendors from emerging if the conditions are met.

9.4 Is the Ada toolset complete?

A language by itself is of little use for the development of safety-critical software. Instead, a complete toolset is needed to accompany the development process, in particular tools for edition, testing, static analysis, etc.

AdaCore provides a number of these tools either in through its core or add-on package. These include (as of 2019):

- An IDE (GNAT Studio)
- An Eclipse plug-in (GNATbench)
- A debugger (GDB)
- A testing tool (GNATtest)
- A structural code coverage tool (GNATcoverage)
- A metric computation tool (GNATmetric)
- A coding standard checker (GNATcheck)
- Static analysis tools (CodePeer, SPARK Pro)
- A Simulink code generator (QGen)
- An Ada parser to develop custom tools (libadalang)

Ada is, however, an internationally standardized language, and many companies are providing third party solutions to complete the toolset. Overall, the language can be and is used with tools on par with their equivalent C counterparts.

9.5 Where can I find Ada or SPARK developers?

A common question from teams on the verge of selecting Ada and SPARK is how to manage the developer team growth and turnover. While Ada and SPARK are taught by a growing number of universities worldwide, it may still be challenging to hire new staff with prior Ada experience.

Fortunately, Ada's base semantics are very close to those of C/C++, so that a good embedded software developer should be able to learn it relatively easily. This course is definitely a resource available to get started. Online training material is also available, together with on-site in person training.

In general, getting an engineer operational in Ada and SPARK shouldn't take more than a few weeks worth of time.

9.6 How to introduce Ada and SPARK in an existing code base?

The most common scenario when introducing Ada and SPARK to a project or a team is to do it within a pre-existing C codebase, which can already spread over hundreds of thousands if not millions lines of code. Re-writing this software to Ada or SPARK is of course not practical and counterproductive.

Most teams select either a small piece of existing code which deserves particular attention, or new modules to develop, and concentrate on this. Developing this module or part of the application will also help in developing the coding patterns to be used for the particular project and company. This typically concentrates an effort of a few people on a few thousands lines of code. The resulting code can be linked to the rest of the C application. From there, the newly established practices and their benefit can slowly spread through the rest of the environment.

Establishing this initial core in Ada and SPARK is critical, and while learning the language isn't a particularly difficult task, applying it to its full capacity may require some expertise. One possibility to accelerate this initial process is to use AdaCore mentorship services.

CONCLUSION

Although Ada's syntax might seem peculiar to C developers at first glance, it was designed to increase readability and maintainability, rather than making it faster to write in a condensed manner — as it is often the case in C.

Especially in the embedded domain, C developers are used to working at a very low level, which includes mathematical operations on pointers, complex bit shifts, and logical bitwise operations. C is well designed for such operations because it was designed to replace Assembly language for faster, more efficient programming.

Ada can be used to describe high level semantics and architectures. The beauty of the language, however, is that it can be used all the way down to the lowest levels of the development, including embedded Assembly code or bit-level data management. However, although Ada supports bitwise operations such as masks and shifts, they should be relatively rarely needed. When translating C code to Ada, it's good practice to consider alternatives. In a lot of cases, these operations are used to insert several pieces of data into a larger structure. In Ada, this can be done by describing the structure layout at the type level through representation clauses, and then accessing this structure as any other. For example, we can interpret an arbitrary data type as a bit-field and perform low-level operations on it.

Because Ada is a strongly typed language, it doesn't define any implicit type conversions like C. If we try to compile Ada code that contains type mismatches, we'll get a compilation error. Because the compiler prevents mixing variables of different types without explicit type conversion, we can't accidentally end up in a situation where we assume something will happen implicitly when, in fact, our assumption is incorrect. In this sense, Ada's type system encourages programmers to think about data at a high level of abstraction. Ada supports overlays and unchecked conversions as a way of converting between unrelated data type, which are typically used for interfacing with low-level elements such as registers.

In Ada, arrays aren't interchangeable with operations on pointers like in C. Also, array types are considered first-class citizens and have dedicated semantics such as the availability of the array's boundaries at run-time. Therefore, unhandled array overflows are impossible unless checks are suppressed. Any discrete type can serve as an array index, and we can specify both the starting and ending bounds. In addition, Ada offers high-level operations for copying, slicing, and assigning values to arrays.

Although Ada supports pointers, most situations that would require a pointer in C do not in Ada. In the vast majority of the cases, indirect memory management can be hidden from the developer and thus prevent many potential errors. In C, pointers are typically used to pass references to subprograms, for example. In contrast, Ada parameter modes indicate the flow of information to the reader, leaving the means of passing that information to the compiler.

When translating pointers from C code to Ada, we need to assess whether they are needed in the first place. Ada pointers (access types) should only be used with complex structures that cannot be allocated at run-time. There are many situations that would require a pointer in C, but do not in Ada. For example, arrays — even when dynamically allocated —, results of functions, passing of large structures as parameters, access to registers, etc.

Because of the absence of namespaces, global names in C tend to be very long. Also, because of the absence of overloading, they can even encode type names in their name. In Ada, a package is

a namespace. Also, we can use the private part of a package to declare private types and private subprograms. In fact, private types are useful for preventing the users of those types from depending on the implementation details. Another use-case is the prevention of package users from accessing the package state/data arbitrarily.

Ada has a dedicated set of features for interfacing with other languages, so we can easily interface with our existing C code before translating it to Ada. Also, GNAT includes automatic binding generators. Therefore, instead of re-writing the entire C code upfront, which isn't practical or cost-effective, we can selectively translate modules from C to Ada.

When it comes to implementing concurrency and real time, Ada offers several options. Ada provides high level constructs such as tasks and protected objects to express concurrency and synchronization, which can be used when running on top of an operating system such as Linux. On more constrained systems, such as bare metal or some real-time operating systems, a subset of the Ada tasking capabilities — known as the Ravenscar and Jorvik profiles — is available. Though restricted, this subset also has nice properties, in particular the absence of deadlock, the absence of priority inversion, schedulability and very small footprint. On bare metal systems, this also essentially means that Ada comes with its own real-time kernel. The advantage of using the full Ada tasking model or the restricted profiles is to enhance portability.

Ada includes many features typically used for embedded programming:

- Built-in support for handling interrupts, so we can process interrupts by attaching a handler — as a protected procedure — to it.
- Built-in support for handling both volatile and atomic data.
- Support for register overlays, which we can use to create a structure that facilitates manipulating bits from registers.
- Support for creating data streams for serialization of arbitrary information and transmission over a communication channel, such as a serial port.
- Built-in support for fixed-point arithmetic, which is an option when our target device doesn't have a floating-point unit or the result of calculations needs to be bit-exact.

Also, Ada compilers such as GNAT have built-in support for directly mixing Ada and Assembly code.

Ada also supports contracts, which can be associated with types and variables to refine values and define valid and invalid values. The most common kind of contract is a *range constraint* — using the *range* reserved word. Ada also supports contract-based programming in the form of preconditions and postconditions. One typical benefit of contract-based programming is the removal of defensive code in subprogram implementations.

It is common to see embedded software being used in a variety of configurations that require small changes to the code for each instance. In C, variability is usually achieved through macros and function pointers, the former being tied to static variability and the latter to dynamic variability. Ada offers many alternatives for both techniques, which aim at structuring possible variations of the software. Examples of static variability in Ada are: genericity, simple derivation, configuration pragma files, and configuration packages. Examples of dynamic variability in Ada are: records with discriminants, variant records — which may include the use of unions —, object orientation, pointers to subprograms, and design by components using dynamic libraries.

There shouldn't be significant performance differences between code written in Ada and code written in C — provided that they are semantically equivalent. One reason is that the two languages are fairly similar in the way they implement imperative semantics, in particular with regards to memory management or control flow. Therefore, they should be equivalent on average. However, when a piece of code in Ada is significantly slower than its counterpart in C, this usually comes from the fact that, while the two pieces of code appear to be semantically equivalent, they happen to be actually quite different. Fortunately, there are strategies that we can use to improve the performance and make it equivalent to the C version. These are some examples:

- Clever use of compilation switches, which might optimize the performance of an application significantly.

- Suppression of checks at specific parts of the implementation.
 - Although runtime checks are very useful and should be used as much as possible, they can also increase the overhead of implementations at certain hot-spots.
- Restriction of assertions to development code.
 - For example, we may use assertions in the debug version of the code and turn them off in the release version.
 - Also, we may use formal proof to decide which assertions we turn off in the release version. By formally proving that assertions will never fail at run-time, we can safely deactivate them.

Formal proof — a form of static analysis — can give strong guarantees about checks, for all possible conditions and all possible inputs. It verifies conditions prior to execution, even prior to compilation, so we can remove bugs earlier in the development phase. This is far less expensive than doing so later because the cost to fix bugs increases exponentially over the phases of the project life cycle, especially after deployment. Preventing bug introduction into the deployed system is the least expensive approach of all.

Formal analysis for proof can be achieved through the SPARK subset of the Ada language combined with the **gnatprove** verification tool. SPARK is a subset encompassing most of the Ada language, except for features that preclude proof.

In Ada, several common programming errors that are not already detected at compile-time are detected instead at run-time, triggering *exceptions* that interrupt the normal flow of execution. However, we may be able to prove that the language-defined checks won't raise exceptions at run-time. This is known as proving *Absence of Run-Time Errors*. Successful proof of these checks is highly significant in itself. One of the major resulting benefits is that we can deploy the final executable with checks disabled.

In many situations, the migration of C code to Ada is justified by an increase in terms of integrity expectations, in which case it's expected that development costs will raise. However, Ada is a more expressive, powerful language, designed to reduce errors earlier in the life-cycle, thus reducing costs. Therefore, Ada makes it possible to write very safe and secure software at a lower cost than languages such as C.

APPENDIX A: HANDS-ON OBJECT-ORIENTED PROGRAMMING

The goal of this appendix is to present a hands-on view on how to translate a system from C to Ada and improve it with object-oriented programming.

11.1 System Overview

Let's start with an overview of a simple system that we'll implement and use below. The main system is called AB and it combines two systems A and B. System AB is not supposed to do anything useful. However, it can serve as a good model for the hands-on we're about to start.

This is a list of requirements for the individual systems A and B, and the combined system AB:

- System A:
 - The system can be activated and deactivated.
 - ✱ During activation, the system's values are reset.
 - Its current value (in floating-point) can be retrieved.
 - ✱ This value is the average of the two internal floating-point values.
 - Its current state (activated or deactivated) can be retrieved.
- System B:
 - The system can be activated and deactivated.
 - ✱ During activation, the system's value is reset.
 - Its current value (in floating-point) can be retrieved.
 - Its current state (activated or deactivated) can be retrieved.
- System AB
 - The system contains an instance of system A and an instance of system B.
 - The system can be activated and deactivated.
 - ✱ System AB activates both systems A and B during its own activation.
 - ✱ System AB deactivates both systems A and B during its own deactivation.
 - Its current value (in floating-point) can be retrieved.
 - ✱ This value is the average of the current values of systems A and B.
 - Its current state (activated or deactivated) can be retrieved.
 - ✱ AB is only considered activated when both systems A and B are activated.
 - The system's health can be checked.

- ✱ This check consists in calculating the absolute difference D between the current values of systems A and B and checking whether D is below a threshold of 0.1.

The source-code in the following section contains an implementation of these requirements.

11.2 Non Object-Oriented Approach

In this section, we look into implementations (in both C and Ada) of system AB that don't make use of object-oriented programming.

11.2.1 Starting point in C

Let's start with an implementation in C for the system described above:

[C]

Listing 1: system_a.h

```
1 typedef struct {
2     float val[2];
3     int active;
4 } A;
5
6 void A_activate (A *a);
7
8 int A_is_active (A *a);
9
10 float A_value (A *a);
11
12 void A_deactivate (A *a);
```

Listing 2: system_a.c

```
1 #include "system_a.h"
2
3 void A_activate (A *a)
4 {
5     int i;
6
7     for (i = 0; i < 2; i++)
8     {
9         a->val[i] = 0.0;
10    }
11    a->active = 1;
12 }
13
14 int A_is_active (A *a)
15 {
16     return a->active == 1;
17 }
18
19 float A_value (A *a)
20 {
21     return (a->val[0] + a->val[1]) / 2.0;
22 }
23
24 void A_deactivate (A *a)
25 {
```

(continues on next page)

(continued from previous page)

```

26     a->active = 0;
27 }

```

Listing 3: system_b.h

```

1  typedef struct {
2      float val;
3      int  active;
4  } B;
5
6  void B_activate (B *b);
7
8  int B_is_active (B *b);
9
10 float B_value (B *b);
11
12 void B_deactivate (B *b);

```

Listing 4: system_b.c

```

1  #include "system_b.h"
2
3  void B_activate (B *b)
4  {
5      b->val    = 0.0;
6      b->active = 1;
7  }
8
9  int B_is_active (B *b)
10 {
11     return b->active == 1;
12 }
13
14 float B_value (B *b)
15 {
16     return b->val;
17 }
18
19 void B_deactivate (B *b)
20 {
21     b->active = 0;
22 }

```

Listing 5: system_ab.h

```

1  #include "system_a.h"
2  #include "system_b.h"
3
4  typedef struct {
5      A a;
6      B b;
7  } AB;
8
9  void AB_activate (AB *ab);
10
11 int AB_is_active (AB *ab);
12
13 float AB_value (AB *ab);
14
15 int AB_check (AB *ab);

```

(continues on next page)

(continued from previous page)

```

16
17 void AB_deactivate (AB *ab);

```

Listing 6: system_ab.c

```

1  #include <math.h>
2  #include "system_ab.h"
3
4  void AB_activate (AB *ab)
5  {
6      A_activate (&ab->a);
7      B_activate (&ab->b);
8  }
9
10 int AB_is_active (AB *ab)
11 {
12     return A_is_active(&ab->a) && B_is_active(&ab->b);
13 }
14
15 float AB_value (AB *ab)
16 {
17     return (A_value (&ab->a) + B_value (&ab->b)) / 2;
18 }
19
20 int AB_check (AB *ab)
21 {
22     const float threshold = 0.1;
23
24     return fabs (A_value (&ab->a) - B_value (&ab->b)) < threshold;
25 }
26
27 void AB_deactivate (AB *ab)
28 {
29     A_deactivate (&ab->a);
30     B_deactivate (&ab->b);
31 }

```

Listing 7: main.c

```

1  #include <stdio.h>
2  #include "system_ab.h"
3
4  void display_active (AB *ab)
5  {
6      if (AB_is_active (ab))
7          printf ("System AB is active.\n");
8      else
9          printf ("System AB is not active.\n");
10 }
11
12 void display_check (AB *ab)
13 {
14     if (AB_check (ab))
15         printf ("System AB check: PASSED.\n");
16     else
17         printf ("System AB check: FAILED.\n");
18 }
19
20 int main()
21 {

```

(continues on next page)

(continued from previous page)

```

22     AB s;
23
24     printf ("Activating system AB...\n");
25     AB_activate (&s);
26
27     display_active (&s);
28     display_check (&s);
29
30     printf ("Deactivating system AB...\n");
31     AB_deactivate (&s);
32
33     display_active (&s);
34 }

```

Runtime output

```

Activating system AB...
System AB is active.
System AB check: PASSED.
Deactivating system AB...
System AB is not active.

```

Here, each system is implemented in a separate set of header and source-code files. For example, the API of system AB is in `system_ab.h` and its implementation in `system_ab.c`.

In the main application, we instantiate system AB and activate it. Then, we proceed to display the activation state and the result of the system's health check. Finally, we deactivate the system and display the activation state again.

11.2.2 Initial translation to Ada

The direct implementation in Ada is:

[Ada]

Listing 8: `system_a.ads`

```

1  package System_A is
2
3      type Val_Array is array (Positive range <>) of Float;
4
5      type A is record
6          Val      : Val_Array (1 .. 2);
7          Active   : Boolean;
8      end record;
9
10     procedure A_Activate (E : in out A);
11
12     function A_Is_Active (E : A) return Boolean;
13
14     function A_Value (E : A) return Float;
15
16     procedure A_Deactivate (E : in out A);
17
18 end System_A;

```

Listing 9: system_a.adb

```
1 package body System_A is
2
3     procedure A_Activate (E : in out A) is
4     begin
5         E.Val      := (others => 0.0);
6         E.Active := True;
7     end A_Activate;
8
9     function A_Is_Active (E : A) return Boolean is
10    begin
11        return E.Active;
12    end A_Is_Active;
13
14    function A_Value (E : A) return Float is
15    begin
16        return (E.Val (1) + E.Val (2)) / 2.0;
17    end A_Value;
18
19    procedure A_Deactivate (E : in out A) is
20    begin
21        E.Active := False;
22    end A_Deactivate;
23
24 end System_A;
```

Listing 10: system_b.ads

```
1 package System_B is
2
3     type B is record
4         Val      : Float;
5         Active   : Boolean;
6     end record;
7
8     procedure B_Activate (E : in out B);
9
10    function B_Is_Active (E : B) return Boolean;
11
12    function B_Value (E : B) return Float;
13
14    procedure B_Deactivate (E : in out B);
15
16 end System_B;
```

Listing 11: system_b.adb

```
1 package body System_B is
2
3     procedure B_Activate (E : in out B) is
4     begin
5         E.Val      := 0.0;
6         E.Active := True;
7     end B_Activate;
8
9     function B_Is_Active (E : B) return Boolean is
10    begin
11        return E.Active;
12    end B_Is_Active;
13
14    function B_Value (E : B) return Float is
```

(continues on next page)

(continued from previous page)

```

15  begin
16      return E.Val;
17  end B_Value;
18
19  procedure B_Deactivate (E : in out B) is
20  begin
21      E.Active := False;
22  end B_Deactivate;
23
24  end System_B;

```

Listing 12: system_ab.ads

```

1  with System_A; use System_A;
2  with System_B; use System_B;
3
4  package System_AB is
5
6      type AB is record
7          SA : A;
8          SB : B;
9      end record;
10
11     procedure AB_Activate (E : in out AB);
12
13     function AB_Is_Active (E : AB) return Boolean;
14
15     function AB_Value (E : AB) return Float;
16
17     function AB_Check (E : AB) return Boolean;
18
19     procedure AB_Deactivate (E : in out AB);
20
21  end System_AB;

```

Listing 13: system_ab.adb

```

1  package body System_AB is
2
3      procedure AB_Activate (E : in out AB) is
4      begin
5          A_Activate (E.SA);
6          B_Activate (E.SB);
7      end AB_Activate;
8
9      function AB_Is_Active (E : AB) return Boolean is
10     begin
11         return A_Is_Active (E.SA) and B_Is_Active (E.SB);
12     end AB_Is_Active;
13
14     function AB_Value (E : AB) return Float is
15     begin
16         return (A_Value (E.SA) + B_Value (E.SB)) / 2.0;
17     end AB_Value;
18
19     function AB_Check (E : AB) return Boolean is
20     Threshold : constant := 0.1;
21     begin
22         return abs (A_Value (E.SA) - B_Value (E.SB)) < Threshold;
23     end AB_Check;

```

(continues on next page)

(continued from previous page)

```

24
25 procedure AB_Deactivate (E : in out AB) is
26 begin
27     A_Deactivate (E.SA);
28     B_Deactivate (E.SB);
29 end AB_Deactivate;
30
31 end System_AB;

```

Listing 14: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 with System_AB; use System_AB;
4
5 procedure Main is
6
7     procedure Display_Active (E : AB) is
8     begin
9         if AB_Is_Active (E) then
10             Put_Line ("System AB is active");
11         else
12             Put_Line ("System AB is not active");
13         end if;
14     end Display_Active;
15
16     procedure Display_Check (E : AB) is
17     begin
18         if AB_Check (E) then
19             Put_Line ("System AB check: PASSED");
20         else
21             Put_Line ("System AB check: FAILED");
22         end if;
23     end Display_Check;
24
25     S : AB;
26 begin
27     Put_Line ("Activating system AB...");
28     AB_Activate (S);
29
30     Display_Active (S);
31     Display_Check (S);
32
33     Put_Line ("Deactivating system AB...");
34     AB_Deactivate (S);
35
36     Display_Active (S);
37 end Main;

```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      system_ab.adb
  [Ada]      system_a.adb
  [Ada]      system_b.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb

```

Runtime output

```

Activating system AB...
System AB is active
System AB check: PASSED
Deactivating system AB...
System AB is not active

```

As you can see, this is a direct translation that doesn't change much of the structure of the original C code. Here, the goal was to simply translate the system from one language to another and make sure that the behavior remains the same.

11.2.3 Improved Ada implementation

By analyzing this direct implementation, we may notice the following points:

- Packages `System_A`, `System_B` and `System_AB` are used to describe aspects of the same system. Instead of having three distinct packages, we could group them as child packages of a common parent package — let's call it `Simple`, since this system is supposed to be simple. This approach has the advantage of allowing us to later use the parent package to implement functionality that is common for all parts of the system.
- Since we have subprograms that operate on types `A`, `B` and `AB`, we should avoid exposing the record components by moving the type declarations to the private part of the corresponding packages.
- Since Ada supports subprogram overloading — as discussed in [this section from chapter 2](#) (page 62) —, we don't need to have different names for subprograms with similar functionality. For example, instead of having `A_Is_Active` and `B_Is_Active`, we can simply name these functions `Is_Active` for both types `A` and `B`.
- Some of the functions — such as `A_Is_Active` and `A_Value` — are very simple, so we could simplify them with expression functions.

This is an update to the implementation that addresses all the points above:

[Ada]

Listing 15: `simple.ads`

```

1 package Simple
2   with Pure
3 is
4 end Simple;

```

Listing 16: `simple-system_a.ads`

```

1 package Simple.System_A is
2
3   type A is private;
4
5   procedure Activate (E : in out A);
6
7   function Is_Active (E : A) return Boolean;
8
9   function Value (E : A) return Float;
10
11  procedure Finalize (E : in out A);
12
13 private
14

```

(continues on next page)

(continued from previous page)

```
15  type Val_Array is array (Positive range <>) of Float;
16
17  type A is record
18    Val    : Val_Array (1 .. 2);
19    Active : Boolean;
20  end record;
21
22  end Simple.System_A;
```

Listing 17: simple-system_a.adb

```
1  package body Simple.System_A is
2
3    procedure Activate (E : in out A) is
4    begin
5      E.Val    := (others => 0.0);
6      E.Active := True;
7    end Activate;
8
9    function Is_Active (E : A) return Boolean is
10     (E.Active);
11
12    function Value (E : A) return Float is
13    begin
14      return (E.Val (1) + E.Val (2)) / 2.0;
15    end Value;
16
17    procedure Finalize (E : in out A) is
18    begin
19      E.Active := False;
20    end Finalize;
21
22  end Simple.System_A;
```

Listing 18: simple-system_b.ads

```
1  package Simple.System_B is
2
3    type B is private;
4
5    procedure Activate (E : in out B);
6
7    function Is_Active (E : B) return Boolean;
8
9    function Value (E : B) return Float;
10
11    procedure Finalize (E : in out B);
12
13  private
14
15    type B is record
16      Val    : Float;
17      Active : Boolean;
18    end record;
19
20  end Simple.System_B;
```

Listing 19: simple-system_b.adb

```

1 package body Simple.System_B is
2
3     procedure Activate (E : in out B) is
4     begin
5         E.Val      := 0.0;
6         E.Active := True;
7     end Activate;
8
9     function Is_Active (E : B) return Boolean is
10    begin
11        return E.Active;
12    end Is_Active;
13
14    function Value (E : B) return Float is
15        (E.Val);
16
17    procedure Finalize (E : in out B) is
18    begin
19        E.Active := False;
20    end Finalize;
21
22 end Simple.System_B;

```

Listing 20: simple-system_ab.ads

```

1 with Simple.System_A; use Simple.System_A;
2 with Simple.System_B; use Simple.System_B;
3
4 package Simple.System_AB is
5
6     type AB is private;
7
8     procedure Activate (E : in out AB);
9
10    function Is_Active (E : AB) return Boolean;
11
12    function Value (E : AB) return Float;
13
14    function Check (E : AB) return Boolean;
15
16    procedure Finalize (E : in out AB);
17
18 private
19
20    type AB is record
21        SA : A;
22        SB : B;
23    end record;
24
25 end Simple.System_AB;

```

Listing 21: simple-system_ab.adb

```

1 package body Simple.System_AB is
2
3     procedure Activate (E : in out AB) is
4     begin
5         Activate (E.SA);
6         Activate (E.SB);

```

(continues on next page)

(continued from previous page)

```

7   end Activate;
8
9   function Is_Active (E : AB) return Boolean is
10      (Is_Active (E.SA) and Is_Active (E.SB));
11
12   function Value (E : AB) return Float is
13      ((Value (E.SA) + Value (E.SB)) / 2.0);
14
15   function Check (E : AB) return Boolean is
16      Threshold : constant := 0.1;
17   begin
18      return abs (Value (E.SA) - Value (E.SB)) < Threshold;
19   end Check;
20
21   procedure Finalize (E : in out AB) is
22   begin
23      Finalize (E.SA);
24      Finalize (E.SB);
25   end Finalize;
26
27 end Simple.System_AB;

```

Listing 22: main.adb

```

1  with Ada.Text_IO;      use Ada.Text_IO;
2
3  with Simple.System_AB; use Simple.System_AB;
4
5  procedure Main is
6
7      procedure Display_Active (E : AB) is
8      begin
9          if Is_Active (E) then
10             Put_Line ("System AB is active");
11          else
12             Put_Line ("System AB is not active");
13          end if;
14      end Display_Active;
15
16      procedure Display_Check (E : AB) is
17      begin
18          if Check (E) then
19             Put_Line ("System AB check: PASSED");
20          else
21             Put_Line ("System AB check: FAILED");
22          end if;
23      end Display_Check;
24
25      S : AB;
26  begin
27      Put_Line ("Activating system AB...");
28      Activate (S);
29
30      Display_Active (S);
31      Display_Check (S);
32
33      Put_Line ("Deactivating system AB...");
34      Finalize (S);
35
36      Display_Active (S);
37  end Main;

```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      simple.ads
  [Ada]      simple-system_ab.adb
  [Ada]      simple-system_a.adb
  [Ada]      simple-system_b.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb

```

Runtime output

```

Activating system AB...
System AB is active
System AB check: PASSED
Deactivating system AB...
System AB is not active

```

11.3 First Object-Oriented Approach

Until now, we haven't used any of the object-oriented programming features of the Ada language. So we can start by analyzing the API of systems A and B and deciding how to best abstract some of its elements using object-oriented programming.

11.3.1 Interfaces

The first thing we may notice is that we actually have two distinct sets of APIs there:

- one API for activating and deactivating the system.
- one API for retrieving the value of the system.

We can use this distinction to declare two interface types:

- `Activation_IF` for the `Activate` and `Deactivate` procedures and the `Is_Active` function;
- `Value_Retrieval_IF` for the `Value` function.

This is how the declaration could look like:

```

type Activation_IF is interface;

procedure Activate (E : in out Activation_IF) is abstract;
function Is_Active (E : Activation_IF) return Boolean is abstract;
procedure Deactivate (E : in out Activation_IF) is abstract;

type Value_Retrieval_IF is interface;

function Value (E : Value_Retrieval_IF) return Float is abstract;

```

Note that, because we are declaring interface types, all operations on those types must be abstract or, in the case of procedures, they can also be declared `null`. For example, we could change the declaration of the procedures above to this:

```
procedure Activate (E : in out Activation_IF) is null;  
procedure Deactivate (E : in out Activation_IF) is null;
```

When an operation is declared abstract, we must override it for the type that derives from the interface. When a procedure is declared `null`, it acts as a do-nothing default. In this case, overriding the operation is optional for the type that derives from this interface.

11.3.2 Base type

Since the original system needs both interfaces we've just described, we have to declare another type that combines those interfaces. We can do this by declaring the interface type `Sys_Base`, which serves as the base type for systems A and B. This is the declaration:

```
type Sys_Base is interface and Activation_IF and Value_Retrieval_IF;
```

Since the system activation functionality is common for both systems A and B, we could implement it as part of `Sys_Base`. That would require changing the declaration from a simple interface to an abstract record:

```
type Sys_Base is abstract new Activation_IF and Value_Retrieval_IF  
  with null record;
```

Now, we can add the Boolean component to the record (as a private component) and override the subprograms of the `Activation_IF` interface. This is the adapted declaration:

```
type Sys_Base is abstract new Activation_IF and Value_Retrieval_IF with private;  
  
  overriding procedure Activate (E : in out Sys_Base);  
  overriding function Is_Active (E : Sys_Base) return Boolean;  
  overriding procedure Deactivate (E : in out Sys_Base);  
  
private  
  
  type Sys_Base is abstract new Activation_IF and Value_Retrieval_IF with record  
    Active : Boolean;  
  end record;
```

11.3.3 Derived types

In the declaration of the `Sys_Base` type we've just seen, we're not overriding the `Value` function — from the `Value_Retrieval_IF` interface — for the `Sys_Base` type, so it remains an abstract function for `Sys_Base`. Therefore, the `Sys_Base` type itself remains abstract and needs to be explicitly declared as such.

We use this strategy to ensure that all types derived from `Sys_Base` need to implement their own version of the `Value` function. For example:

```
type A is new Sys_Base with private;  
  
  overriding function Value (E : A) return Float;
```

Here, the `A` type is derived from the `Sys_Base` and it includes its own version of the `Value` function by overriding it. Therefore, `A` is not an abstract type anymore and can be used to declare objects:

```
procedure Main is  
  Obj : A;  
  V    : Float;
```

(continues on next page)

(continued from previous page)

```
begin
  Obj.Activate;
  V := Obj.Value;
end Main;
```

Important

Note that the use of the overriding keyword in the subprogram declaration is not strictly necessary. In fact, we could leave this keyword out, and the code would still compile. However, if provided, the compiler will check whether the information is correct.

Using the overriding keyword can help to avoid bad surprises — when you *may think* that you're overriding a subprogram, but you're actually not. Similarly, you can also write `not overriding` to be explicit about subprograms that are new primitives of a derived type. For example:

```
not overriding function Check (E : AB) return Boolean;
```

We also need to declare the values that are used internally in systems A and B. For system A, this is the declaration:

```
type A is new Sys_Base with private;

overriding function Value (E : A) return Float;

private

type Val_Array is array (Positive range <>) of Float;

type A is new Sys_Base with record
  Val : Val_Array (1 .. 2);
end record;
```

11.3.4 Subprograms from parent

In the previous implementation, we've seen that the `A_Activate` and `B_Activate` procedures perform the following steps:

- initialize internal values;
- indicate that the system is active (by setting the `Active` flag to `True`).

In the implementation of the `Activate` procedure for the `Sys_Base` type, however, we're only dealing with the second step. Therefore, we need to override the `Activate` procedure and make sure that we initialize internal values as well. First, we need to declare this procedure for type A:

```
type A is new Sys_Base with private;

overriding procedure Activate (E : in out A);
```

In the implementation of `Activate`, we should call the `Activate` procedure from the parent (`Sys_Base`) to ensure that whatever was performed for the parent will be performed in the derived type as well. For example:

```
overriding procedure Activate (E : in out A) is
begin
  E.Val := (others => 0.0);
  Sys_Base (E).Activate;    -- Calling Activate for Sys_Base type:
```

(continues on next page)

(continued from previous page)

```
end;                                -- this call initializes the Active flag.
```

Here, by writing `Sys_Base (E)`, we're performing a view conversion. Basically, we're telling the compiler to view `E` not as an object of type `A`, but of type `Sys_Base`. When we do this, any operation performed on this object will be done as if it was an object of `Sys_Base` type, which includes calling the `Activate` procedure of the `Sys_Base` type.

Important

If we write `T (Obj) . Proc`, we're telling the compiler to call the `Proc` procedure of type `T` and apply it on `Obj`.

If we write `T'Class (Obj) . Proc`, however, we're telling the compiler to dispatch the call. For example, if `Obj` is of derived type `T2` and there's an overridden `Proc` procedure for type `T2`, then this procedure will be called instead of the `Proc` procedure for type `T`.

11.3.5 Type AB

While the implementation of systems `A` and `B` is almost straightforward, it gets more interesting in the case of system `AB`. Here, we have a similar API, but we don't need the activation mechanism implemented in the abstract type `Sys_Base`. Therefore, deriving from `Sys_Base` is not the best option. Instead, when declaring the `AB` type, we can simply use the same interfaces as we did for `Sys_Base`, but keep it independent from `Sys_Base`. For example:

```
type AB is new Activation_IF and Value_Retrieval_IF with private;

private

type AB is new Activation_IF and Value_Retrieval_IF with record
  SA : A;
  SB : B;
end record;
```

Naturally, we still need to override all the subprograms that are part of the `Activation_IF` and `Value_Retrieval_IF` interfaces. Also, we need to implement the additional `Check` function that was originally only available on system `AB`. Therefore, we declare these subprograms:

```
overriding procedure Activate (E : in out AB);
overriding function Is_Active (E : AB) return Boolean;
overriding procedure Deactivate (E : in out AB);

overriding function Value (E : AB) return Float;
not overriding function Check (E : AB) return Boolean;
```

11.3.6 Updated source-code

Finally, this is the complete source-code example:

[Ada]

Listing 23: simple.ads

```

1 package Simple is
2
3   type Activation_IF is interface;
4
5   procedure Activate (E : in out Activation_IF) is abstract;
6   function Is_Active (E : Activation_IF) return Boolean is abstract;
7   procedure Deactivate (E : in out Activation_IF) is abstract;
8
9   type Value_Retrieval_IF is interface;
10
11   function Value (E : Value_Retrieval_IF) return Float is abstract;
12
13   type Sys_Base is abstract new Activation_IF and Value_Retrieval_IF
14     with private;
15
16   overriding procedure Activate (E : in out Sys_Base);
17   overriding function Is_Active (E : Sys_Base) return Boolean;
18   overriding procedure Deactivate (E : in out Sys_Base);
19
20 private
21
22   type Sys_Base is abstract new Activation_IF and Value_Retrieval_IF
23     with record
24       Active : Boolean;
25     end record;
26
27 end Simple;
```

Listing 24: simple.adb

```

1 package body Simple is
2
3   overriding procedure Activate (E : in out Sys_Base) is
4   begin
5     E.Active := True;
6   end Activate;
7
8   overriding function Is_Active (E : Sys_Base) return Boolean is
9   (E.Active);
10
11   overriding procedure Deactivate (E : in out Sys_Base) is
12   begin
13     E.Active := False;
14   end Deactivate;
15
16 end Simple;
```

Listing 25: simple-system_a.ads

```

1 package Simple.System_A is
2
3   type A is new Sys_Base with private;
4
5   overriding procedure Activate (E : in out A);
```

(continues on next page)

(continued from previous page)

```

6      overriding function Value (E : A) return Float;
7
8
9  private
10
11      type Val_Array is array (Positive range <>) of Float;
12
13      type A is new Sys_Base with record
14          Val : Val_Array (1 .. 2);
15      end record;
16
17  end Simple.System_A;

```

Listing 26: simple-system_a.adb

```

1  package body Simple.System_A is
2
3      procedure Activate (E : in out A) is
4      begin
5          E.Val := (others => 0.0);
6          Sys_Base (E).Activate;
7      end Activate;
8
9      function Value (E : A) return Float is
10         pragma Assert (E.Val'Length = 2);
11     begin
12         return (E.Val (1) + E.Val (2)) / 2.0;
13     end Value;
14
15  end Simple.System_A;

```

Listing 27: simple-system_b.ads

```

1  package Simple.System_B is
2
3      type B is new Sys_Base with private;
4
5      overriding procedure Activate (E : in out B);
6
7      overriding function Value (E : B) return Float;
8
9  private
10
11      type B is new Sys_Base with record
12          Val : Float;
13      end record;
14
15  end Simple.System_B;

```

Listing 28: simple-system_b.adb

```

1  package body Simple.System_B is
2
3      procedure Activate (E : in out B) is
4      begin
5          E.Val := 0.0;
6          Sys_Base (E).Activate;
7      end Activate;
8
9      function Value (E : B) return Float is

```

(continues on next page)

(continued from previous page)

```

10     (E.Val);
11
12 end Simple.System_B;

```

Listing 29: simple-system_ab.ads

```

1  with Simple.System_A; use Simple.System_A;
2  with Simple.System_B; use Simple.System_B;
3
4  package Simple.System_AB is
5
6      type AB is new Activation_IF and Value_Retrieval_IF with private;
7
8      overriding procedure Activate (E : in out AB);
9      overriding function Is_Active (E : AB) return Boolean;
10     overriding procedure Deactivate (E : in out AB);
11
12     overriding function Value (E : AB) return Float;
13
14     not overriding function Check (E : AB) return Boolean;
15
16 private
17
18     type AB is new Activation_IF and Value_Retrieval_IF with record
19         SA : A;
20         SB : B;
21     end record;
22
23 end Simple.System_AB;

```

Listing 30: simple-system_ab.adb

```

1  package body Simple.System_AB is
2
3      procedure Activate (E : in out AB) is
4      begin
5          E.SA.Activate;
6          E.SB.Activate;
7      end Activate;
8
9      function Is_Active (E : AB) return Boolean is
10         (E.SA.Is_Active and E.SB.Is_Active);
11
12     procedure Deactivate (E : in out AB) is
13     begin
14         E.SA.Deactivate;
15         E.SB.Deactivate;
16     end Deactivate;
17
18     function Value (E : AB) return Float is
19         ((E.SA.Value + E.SB.Value) / 2.0);
20
21     function Check (E : AB) return Boolean is
22         Threshold : constant := 0.1;
23     begin
24         return abs (E.SA.Value - E.SB.Value) < Threshold;
25     end Check;
26
27 end Simple.System_AB;

```

Listing 31: main.adb

```
1 with Ada.Text_IO;      use Ada.Text_IO;
2
3 with Simple.System_AB; use Simple.System_AB;
4
5 procedure Main is
6
7     procedure Display_Active (E : AB) is
8     begin
9         if Is_Active (E) then
10             Put_Line ("System AB is active");
11         else
12             Put_Line ("System AB is not active");
13         end if;
14     end Display_Active;
15
16     procedure Display_Check (E : AB) is
17     begin
18         if Check (E) then
19             Put_Line ("System AB check: PASSED");
20         else
21             Put_Line ("System AB check: FAILED");
22         end if;
23     end Display_Check;
24
25     S : AB;
26 begin
27     Put_Line ("Activating system AB...");
28     Activate (S);
29
30     Display_Active (S);
31     Display_Check (S);
32
33     Put_Line ("Deactivating system AB...");
34     Deactivate (S);
35
36     Display_Active (S);
37 end Main;
```

Build output

```
Compile
[Ada]      main.adb
[Ada]      simple.adb
[Ada]      simple-system_ab.adb
[Ada]      simple-system_a.adb
[Ada]      simple-system_b.adb
Bind
[gprbind]  main.bexch
[Ada]      main.ali
Link
[link]     main.adb
```

Runtime output

```
Activating system AB...
System AB is active
System AB check: PASSED
Deactivating system AB...
System AB is not active
```

11.4 Further Improvements

When analyzing the complete source-code, we see that there are at least two areas that we could still improve.

11.4.1 Dispatching calls

The first issue concerns the implementation of the `Activate` procedure for types derived from `Sys_Base`. For those derived types, we're expecting that the `Activate` procedure of the parent must be called in the implementation of the overriding `Activate` procedure. For example:

```
package body Simple.System_A is

  procedure Activate (E : in out A) is
  begin
    E.Val := (others => 0.0);
    Activate (Sys_Base (E));
  end;
```

If a developer forgets to call that specific `Activate` procedure, however, the system won't work as expected. A better strategy could be the following:

- Declare a new `Activation_Reset` procedure for `Sys_Base` type.
- Make a dispatching call to the `Activation_Reset` procedure in the body of the `Activate` procedure (of the `Sys_Base` type).
- Let the derived types implement their own version of the `Activation_Reset` procedure.

This is a simplified view of the implementation using the points described above:

```
package Simple is

  type Sys_Base is abstract new Activation_IF and Value_Retrieval_IF with
    private;

  not overriding procedure Activation_Reset (E : in out Sys_Base) is abstract;

end Simple;

package body Simple is

  procedure Activate (E : in out Sys_Base) is
  begin
    -- NOTE: calling "E.Activation_Reset" does NOT dispatch!
    --       We need to use the 'Class attribute here --- not using this
    --       attribute is an error that will be caught by the compiler.
    Sys_Base'Class (E).Activation_Reset;

    E.Active := True;
  end Activate;

end Simple;

package Simple.System_A is

  type A is new Sys_Base with private;

private

  type Val_Array is array (Positive range <>) of Float;
```

(continues on next page)

(continued from previous page)

```

type A is new Sys_Base with record
  Val : Val_Array (1 .. 2);
end record;

overriding procedure Activation_Reset (E : in out A);

end Simple.System_A;

package body Simple.System_A is

  procedure Activation_Reset (E : in out A) is
  begin
    E.Val := (others => 0.0);
  end Activation_Reset;

end Simple.System_A;

```

An important detail is that, in the implementation of `Activate`, we use `Sys_Base'Class` to ensure that the call to `Activation_Reset` will dispatch. If we had just written `E.Activation_Reset` instead, then we would be calling the `Activation_Reset` procedure of `Sys_Base` itself, which is not what we actually want here. The compiler will catch the error if you don't do the conversion to the class-wide type, because it would otherwise be a statically-bound call to an abstract procedure, which is illegal at compile-time.

11.4.2 Dynamic allocation

The next area that we could improve is in the declaration of the system AB. In the previous implementation, we were explicitly describing the two components of that system, namely a component of type A and a component of type B:

```

type AB is new Activation_IF and Value_Retrieval_IF with record
  SA : A;
  SB : B;
end record;

```

Of course, this declaration matches the system requirements that we presented in the beginning. However, we could use strategies that make it easier to incorporate requirement changes later on. For example, we could hide this information about systems A and B by simply declaring an array of components of type `access Sys_Base'Class` and allocate them dynamically in the body of the package. Naturally, this approach might not be suitable for certain platforms. However, the advantage would be that, if we wanted to replace the component of type B by a new component of type C, for example, we wouldn't need to change the interface. This is how the updated declaration could look like:

```

type Sys_Base_Class_Access is access Sys_Base'Class;
type Sys_Base_Array is array (Positive range <>) of Sys_Base_Class_Access;

type AB is limited new Activation_IF and Value_Retrieval_IF with record
  S_Array : Sys_Base_Array (1 .. 2);
end record;

```

Important

Note that we're now using the `limited` keyword in the declaration of type AB. That is necessary because we want to prevent objects of type AB being copied by assignment, which would lead to two objects having the same (dynamically allocated) subsystems A and B internally. This change

requires that both `Activation_IF` and `Value_Retrieval_IF` are declared limited as well.

The body of `Activate` could then allocate those components:

```
procedure Activate (E : in out AB) is
begin
  E.S_Array := (new A, new B);
  for S of E.S_Array loop
    S.Activate;
  end loop;
end Activate;
```

And the body of `Deactivate` could deallocate them:

```
procedure Deactivate (E : in out AB) is
  procedure Free is
    new Ada.Unchecked_Deallocation (Sys_Base'Class, Sys_Base_Class_Access);
  begin
    for S of E.S_Array loop
      S.Deactivate;
      Free (S);
    end loop;
  end Deactivate;
```

11.4.3 Limited controlled types

Another approach that we could use to implement the dynamic allocation of systems A and B is to declare AB as a limited controlled type — based on the `Limited_Controlled` type of the `Ada.Finalization` package.

The `Limited_Controlled` type includes the following operations:

- `Initialize`, which is called when objects of a type derived from the `Limited_Controlled` type are being created — by declaring an object of the derived type, for example —, and
- `Finalize`, which is called when objects are being destroyed — for example, when an object gets out of scope at the end of a subprogram where it was created.

In this case, we must override those procedures, so we can use them for dynamic memory allocation. This is a simplified view of the update implementation:

```
package Simple.System_AB is
  type AB is limited new Ada.Finalization.Limited_Controlled and
    Activation_IF and Value_Retrieval_IF with private;

  overriding procedure Initialize (E : in out AB);
  overriding procedure Finalize  (E : in out AB);
end Simple.System_AB;

package body Simple.System_AB is
  overriding procedure Initialize (E : in out AB) is
  begin
    E.S_Array := (new A, new B);
  end Initialize;

  overriding procedure Finalize  (E : in out AB) is
  procedure Free is
    new Ada.Unchecked_Deallocation (Sys_Base'Class, Sys_Base_Class_Access);
  end Free;
```

(continues on next page)

(continued from previous page)

```

begin
  for S of E.S_Array loop
    Free (S);
  end loop;
end Finalize;

end Simple.System_AB;

```

11.4.4 Updated source-code

Finally, this is the complete updated source-code example:

[Ada]

Listing 32: simple.ads

```

1 package Simple is
2
3   type Activation_IF is limited interface;
4
5   procedure Activate (E : in out Activation_IF) is abstract;
6   function Is_Active (E : Activation_IF) return Boolean is abstract;
7   procedure Deactivate (E : in out Activation_IF) is abstract;
8
9   type Value_Retrieval_IF is limited interface;
10
11   function Value (E : Value_Retrieval_IF) return Float is abstract;
12
13   type Sys_Base is abstract new Activation_IF and Value_Retrieval_IF with
14     private;
15
16   overriding procedure Activate (E : in out Sys_Base);
17   overriding function Is_Active (E : Sys_Base) return Boolean;
18   overriding procedure Deactivate (E : in out Sys_Base);
19
20   not overriding procedure Activation_Reset (E : in out Sys_Base) is abstract;
21
22 private
23
24   type Sys_Base is abstract new Activation_IF and Value_Retrieval_IF with
25     record
26       Active : Boolean;
27     end record;
28
29 end Simple;

```

Listing 33: simple.adb

```

1 package body Simple is
2
3   procedure Activate (E : in out Sys_Base) is
4   begin
5     -- NOTE: calling "E.Activation_Reset" does NOT dispatch!
6     -- We need to use the 'Class attribute:
7     Sys_Base'Class (E).Activation_Reset;
8
9     E.Active := True;
10  end Activate;
11

```

(continues on next page)

(continued from previous page)

```

12  function Is_Active (E : Sys_Base) return Boolean is
13      (E.Active);
14
15  procedure Deactivate (E : in out Sys_Base) is
16  begin
17      E.Active := False;
18  end Deactivate;
19
20  end Simple;

```

Listing 34: simple-system_a.ads

```

1  package Simple.System_A is
2
3      type A is new Sys_Base with private;
4
5      overriding function Value (E : A) return Float;
6
7  private
8
9      type Val_Array is array (Positive range <>) of Float;
10
11     type A is new Sys_Base with record
12         Val : Val_Array (1 .. 2);
13     end record;
14
15     overriding procedure Activation_Reset (E : in out A);
16
17 end Simple.System_A;

```

Listing 35: simple-system_a.adb

```

1  package body Simple.System_A is
2
3      procedure Activation_Reset (E : in out A) is
4      begin
5          E.Val := (others => 0.0);
6      end Activation_Reset;
7
8      function Value (E : A) return Float is
9          pragma Assert (E.Val'Length = 2);
10     begin
11         return (E.Val (1) + E.Val (2)) / 2.0;
12     end Value;
13
14 end Simple.System_A;

```

Listing 36: simple-system_b.ads

```

1  package Simple.System_B is
2
3      type B is new Sys_Base with private;
4
5      overriding function Value (E : B) return Float;
6
7  private
8
9      type B is new Sys_Base with record
10         Val : Float;
11     end record;

```

(continues on next page)

(continued from previous page)

```

12
13     overriding procedure Activation_Reset (E : in out B);
14
15 end Simple.System_B;
```

Listing 37: simple-system_b.adb

```

1 package body Simple.System_B is
2
3     procedure Activation_Reset (E : in out B) is
4     begin
5         E.Val := 0.0;
6     end Activation_Reset;
7
8     function Value (E : B) return Float is
9         (E.Val);
10
11 end Simple.System_B;
```

Listing 38: simple-system_ab.ads

```

1 with Ada.Finalization;
2
3 package Simple.System_AB is
4
5     type AB is limited new Ada.Finalization.Limited_Controlled and
6         Activation_IF and Value_Retrieval_IF with private;
7
8     overriding procedure Activate (E : in out AB);
9     overriding function Is_Active (E : AB) return Boolean;
10    overriding procedure Deactivate (E : in out AB);
11
12    overriding function Value (E : AB) return Float;
13
14    not overriding function Check (E : AB) return Boolean;
15
16 private
17
18     type Sys_Base_Class_Access is access Sys_Base'Class;
19     type Sys_Base_Array is array (Positive range <>) of Sys_Base_Class_Access;
20
21     type AB is limited new Ada.Finalization.Limited_Controlled and
22         Activation_IF and Value_Retrieval_IF with record
23         S_Array : Sys_Base_Array (1 .. 2);
24     end record;
25
26     overriding procedure Initialize (E : in out AB);
27     overriding procedure Finalize (E : in out AB);
28
29 end Simple.System_AB;
```

Listing 39: simple-system_ab.adb

```

1 with Ada.Unchecked_Deallocation;
2
3 with Simple.System_A; use Simple.System_A;
4 with Simple.System_B; use Simple.System_B;
5
6 package body Simple.System_AB is
7
```

(continues on next page)

(continued from previous page)

```

8  overriding procedure Initialize (E : in out AB) is
9  begin
10     E.S_Array := (new A, new B);
11 end Initialize;
12
13 overriding procedure Finalize (E : in out AB) is
14     procedure Free is
15         new Ada.Unchecked_Deallocation (Sys_Base'Class, Sys_Base_Class_Access);
16     begin
17         for S of E.S_Array loop
18             Free (S);
19         end loop;
20     end Finalize;
21
22     procedure Activate (E : in out AB) is
23     begin
24         for S of E.S_Array loop
25             S.Activate;
26         end loop;
27     end Activate;
28
29     function Is_Active (E : AB) return Boolean is
30         (for all S of E.S_Array => S.Is_Active);
31
32     procedure Deactivate (E : in out AB) is
33     begin
34         for S of E.S_Array loop
35             S.Deactivate;
36         end loop;
37     end Deactivate;
38
39     function Value (E : AB) return Float is
40         ((E.S_Array (1).Value + E.S_Array (2).Value) / 2.0);
41
42     function Check (E : AB) return Boolean is
43         Threshold : constant := 0.1;
44     begin
45         return abs (E.S_Array (1).Value - E.S_Array (2).Value) < Threshold;
46     end Check;
47
48 end Simple.System_AB;

```

Listing 40: main.adb

```

1  with Ada.Text_IO;      use Ada.Text_IO;
2
3  with Simple.System_AB; use Simple.System_AB;
4
5  procedure Main is
6
7     procedure Display_Active (E : AB) is
8     begin
9         if Is_Active (E) then
10             Put_Line ("System AB is active");
11         else
12             Put_Line ("System AB is not active");
13         end if;
14     end Display_Active;
15
16     procedure Display_Check (E : AB) is
17     begin

```

(continues on next page)

(continued from previous page)

```

18     if Check (E) then
19         Put_Line ("System AB check: PASSED");
20     else
21         Put_Line ("System AB check: FAILED");
22     end if;
23 end Display_Check;
24
25 S : AB;
26 begin
27     Put_Line ("Activating system AB...");
28     Activate (S);
29
30     Display_Active (S);
31     Display_Check (S);
32
33     Put_Line ("Deactivating system AB...");
34     Deactivate (S);
35
36     Display_Active (S);
37 end Main;

```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      simple.adb
  [Ada]      simple-system_ab.adb
  [Ada]      simple-system_a.adb
  [Ada]      simple-system_b.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb

```

Runtime output

```

Activating system AB...
System AB is active
System AB check: PASSED
Deactivating system AB...
System AB is not active

```

Naturally, this is by no means the best possible implementation of system AB. By applying other software design strategies that we haven't covered here, we could most probably think of different ways to use object-oriented programming to improve this implementation. Also, in comparison to the *original implementation* (page 229), we recognize that the amount of source-code has grown. On the other hand, we now have a system that is factored nicely, and also more extensible.

BIBLIOGRAPHY

[Jorvik] A New Ravenscar-Based Profile by P. Rogers, J. Ruiz, T. Gingold and P. Bernardi, in *Reliable Software Technologies — Ada Europe 2017*, Springer-Verlag Lecture Notes in Computer Science, Number 10300.