

Introduction to Ada

Release 2021-02

Raphaël Amiard

Gustavo A. Hoffmann

May 14, 2021

CONTENTS:

1	Introduction	3
1.1	History	3
1.2	Ada today	3
1.3	Philosophy	4
1.4	SPARK	4
2	Imperative language	5
2.1	Hello world	5
2.2	Imperative language - If/Then/Else	7
2.3	Imperative language - Loops	8
2.3.1	For loops	9
2.3.2	Bare loops	11
2.3.3	While loops	12
2.4	Imperative language - Case statement	13
2.5	Imperative language - Declarative regions	14
2.6	Imperative language - conditional expressions	16
2.6.1	If expressions	16
2.6.2	Case expressions	17
3	Subprograms	19
3.1	Subprograms	19
3.1.1	Subprogram calls	20
3.1.2	Nested subprograms	21
3.1.3	Function calls	22
3.2	Parameter modes	23
3.3	Subprogram calls	23
3.3.1	In parameters	23
3.3.2	In out parameters	24
3.3.3	Out parameters	25
3.3.4	Forward declaration of subprograms	26
3.4	Renaming	26
4	Modular programming	29
4.1	Packages	29
4.2	Using a package	31
4.3	Package body	31
4.4	Child packages	33
4.4.1	Child of a child package	35
4.4.2	Multiple children	36
4.4.3	Visibility	37
4.5	Renaming	39
5	Strongly typed language	41
5.1	What is a type?	41
5.2	Integers	41

5.2.1	Operational semantics	43
5.3	Unsigned types	44
5.4	Enumerations	45
5.5	Floating-point types	46
5.5.1	Basic properties	46
5.5.2	Precision of floating-point types	47
5.5.3	Range of floating-point types	49
5.6	Strong typing	50
5.7	Derived types	53
5.8	Subtypes	54
5.8.1	Subtypes as type aliases	56
6	Records	59
6.1	Record type declaration	59
6.2	Aggregates	60
6.3	Component selection	60
6.4	Renaming	61
7	Arrays	65
7.1	Array type declaration	65
7.2	Indexing	68
7.3	Simpler array declarations	69
7.4	Range attribute	70
7.5	Unconstrained arrays	71
7.6	Predefined array type: String	73
7.7	Restrictions	75
7.8	Returning unconstrained arrays	76
7.9	Declaring arrays (2)	77
7.10	Array slices	78
7.11	Renaming	79
8	More about types	83
8.1	Aggregates: A primer	83
8.2	Overloading and qualified expressions	84
8.3	Access types (pointers)	86
8.3.1	Allocation (by type)	87
8.3.2	Dereferencing	88
8.3.3	Other features	89
8.4	Mutually recursive types	89
8.5	More about records	90
8.5.1	Dynamically sized record types	90
8.5.2	Records with discriminant	91
8.5.3	Variant records	92
8.6	Fixed-point types	94
8.6.1	Decimal fixed-point types	94
8.6.2	Fixed-point types	96
8.7	Character types	99
9	Privacy	101
9.1	Basic encapsulation	101
9.2	Abstract data types	102
9.3	Limited types	103
9.4	Child packages & privacy	104
10	Generics	109
10.1	Introduction	109
10.2	Formal type declaration	109
10.3	Formal object declaration	110
10.4	Generic body definition	110

10.5	Generic instantiation	111
10.6	Generic packages	112
10.7	Formal subprograms	114
10.8	Example: I/O instances	115
10.9	Example: ADTs	117
10.10	Example: Swap	119
10.11	Example: Reversing	122
10.12	Example: Test application	125
11	Exceptions	129
11.1	Exception declaration	129
11.2	Raising an exception	129
11.3	Handling an exception	130
11.4	Predefined exceptions	132
12	Tasking	133
12.1	Tasks	133
12.1.1	Simple task	133
12.1.2	Simple synchronization	135
12.1.3	Delay	137
12.1.4	Synchronization: rendez-vous	137
12.1.5	Select loop	139
12.1.6	Cycling tasks	140
12.2	Protected objects	144
12.2.1	Simple object	144
12.2.2	Entries	145
12.3	Task and protected types	147
12.3.1	Task types	147
12.3.2	Protected types	149
13	Design by contracts	151
13.1	Pre- and postconditions	151
13.2	Predicates	154
13.3	Type invariants	158
14	Interfacing with C	161
14.1	Multi-language project	161
14.2	Type convention	161
14.3	Foreign subprograms	162
14.3.1	Calling C subprograms in Ada	162
14.3.2	Calling Ada subprograms in C	163
14.4	Foreign variables	164
14.4.1	Using C global variables in Ada	164
14.4.2	Using Ada variables in C	165
14.5	Generating bindings	166
14.5.1	Adapting bindings	167
15	Object-oriented programming	171
15.1	Derived types	172
15.2	Tagged types	173
15.3	Classwide types	174
15.4	Dispatching operations	176
15.5	Dot notation	177
15.6	Private & Limited	179
15.7	Classwide access types	180
16	Standard library: Containers	185
16.1	Vectors	185
16.1.1	Instantiation	185

16.1.2	Initialization	186
16.1.3	Appending and prepending elements	186
16.1.4	Accessing first and last elements	188
16.1.5	Iterating	189
16.1.6	Finding and changing elements	194
16.1.7	Inserting elements	195
16.1.8	Removing elements	196
16.1.9	Other Operations	199
16.2	Sets	202
16.2.1	Initialization and iteration	202
16.2.2	Operations on elements	203
16.2.3	Other Operations	205
16.3	Indefinite maps	208
16.3.1	Hashed maps	208
16.3.2	Ordered maps	210
16.3.3	Complexity	211
17	Standard library: Dates & Times	213
17.1	Date and time handling	213
17.1.1	Delaying using date	214
17.2	Real-time	217
17.2.1	Benchmarking	218
18	Standard library: Strings	223
18.1	String operations	223
18.2	Limitation of fixed-length strings	227
18.3	Bounded strings	228
18.4	Unbounded strings	230
19	Standard library: Files and streams	233
19.1	Text I/O	233
19.2	Sequential I/O	236
19.3	Direct I/O	238
19.4	Stream I/O	240
20	Standard library: Numerics	245
20.1	Elementary Functions	245
20.2	Random Number Generation	246
20.3	Complex Types	248
20.4	Vector and Matrix Manipulation	250
21	Appendices	255
21.1	Appendix A: Generic Formal Types	255
21.1.1	Indefinite version	256
21.2	Appendix B: Containers	257

Copyright © 2018 – 2020, 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 will teach you the basics of the Ada programming language and is intended for those who already have a basic understanding of programming techniques. You will learn how to apply those techniques to programming in Ada.

This document was written by Raphaël Amiard and Gustavo A. Hoffmann, with review from Richard Kenner.

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

INTRODUCTION

1.1 History

In the 1970s the United States Department of Defense (DOD) suffered from an explosion of the number of programming languages, with different projects using different and non-standard dialects or language subsets / supersets. The DOD decided to solve this problem by issuing a request for proposals for a common, modern programming language. The winning proposal was one submitted by Jean Ichbiah from CII Honeywell-Bull.

The first Ada standard was issued in 1983; it was subsequently revised and enhanced in 1995, 2005 and 2012, with each revision bringing useful new features.

This tutorial will focus on Ada 2012 as a whole, rather than teaching different versions of the language.

1.2 Ada today

Today, Ada is heavily used in embedded real-time systems, many of which are safety critical. While Ada is and can be used as a general-purpose language, it will really shine in low-level applications:

- Embedded systems with low memory requirements (no garbage collector allowed).
- Direct interfacing with hardware.
- Soft or hard real-time systems.
- Low-level systems programming.

Specific domains seeing Ada usage include Aerospace & Defense, civil aviation, rail, and many others. These applications require a high degree of safety: a software defect is not just an annoyance, but may have severe consequences. Ada provides safety features that detect defects at an early stage — usually at compilation time or using static analysis tools. Ada can also be used to create applications in a variety of other areas, such as:

- Video game programming²
- Real-time audio³
- Kernel modules⁴

This is a non-comprehensive list that hopefully sheds light on which kind of programming Ada is good at.

² <https://github.com/AdaDoom3/AdaDoom3>

³ <http://www.electronicdesign.com/embedded-revolution/assessing-ada-language-audio-applications>

⁴ <http://www.nihamkin.com/tag/kernel.html>

In terms of modern languages, the closest in terms of targets and level of abstraction are probably C++⁵ and Rust⁶.

1.3 Philosophy

Ada's philosophy is different from most other languages. Underlying Ada's design are principles that include the following:

- Readability is more important than conciseness. Syntactically this shows through the fact that keywords are preferred to symbols, that no keyword is an abbreviation, etc.
- Very strong typing. It is very easy to introduce new types in Ada, with the benefit of preventing data usage errors.
 - It is similar to many functional languages in that regard, except that the programmer has to be much more explicit about typing in Ada, because there is almost no type inference.
- Explicit is better than implicit. Although this is a Python⁷ commandment, Ada takes it way further than any language we know of:
 - There is mostly no structural typing, and most types need to be explicitly named by the programmer.
 - As previously said, there is mostly no type inference.
 - Semantics are very well defined, and undefined behavior is limited to an absolute minimum.
 - The programmer can generally give a *lot* of information about what their program means to the compiler (and other programmers). This allows the compiler to be extremely helpful (read: strict) with the programmer.

During this course, we will explain the individual language features that are building blocks for that philosophy.

1.4 SPARK

While this class is solely about the Ada language, it is worth mentioning that another language, extremely close to and interoperable with Ada, exists: the SPARK language.

SPARK is a subset of Ada, designed so that the code written in SPARK is amenable to automatic proof. This provides a level of assurance with regard to the correctness of your code that is much higher than with a regular programming language.

There is a dedicated [course for the SPARK language](#)⁸ but keep in mind that every time we speak about the specification power of Ada during this course, it is power that you can leverage in SPARK to help proving the correctness of program properties ranging from absence of run-time errors to compliance with formally specified functional requirements.

⁵ <https://en.wikipedia.org/wiki/C%2B%2B>

⁶ <https://www.rust-lang.org/en-US/>

⁷ <https://www.python.org>

⁸ <https://learn.adacore.com/courses/intro-to-spark/index.html>

IMPERATIVE LANGUAGE

Ada is a multi-paradigm language with support for object orientation and some elements of functional programming, but its core is a simple, coherent procedural/imperative language akin to C or Pascal.

In other languages

One important distinction between Ada and a language like C is that statements and expressions are very clearly distinguished. In Ada, if you try to use an expression where a statement is required then your program will fail to compile. This rule supports a useful stylistic principle: expressions are intended to deliver values, not to have side effects. It can also prevent some programming errors, such as mistakenly using the equality operator "=" instead of the assignment operation ":=" in an assignment statement.

2.1 Hello world

Here's a very simple imperative Ada program:

Listing 1: greet.adb

```
1 with Ada.Text_IO;  
2  
3 procedure Greet is  
4 begin  
5     -- Print "Hello, World!" to the screen  
6     Ada.Text_IO.Put_Line ("Hello, World!");  
7 end Greet;
```

Build output

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

Runtime output

```
Hello, World!
```

which we'll assume is in the source file `greet.adb`.

There are several noteworthy things in the above program:

- A subprogram in Ada can be either a procedure or a function. A procedure, as illustrated above, does not return a value when called.
- `with` is used to reference external modules that are needed in the procedure. This is similar to `import` in various languages or roughly similar to `#include` in C and C++. We'll see later how they work in detail. Here, we are requesting a standard library module, the `Ada.Text_IO` package, which contains a procedure to print text on the screen: `Put_Line`.
- `Greet` is a procedure, and the main entry point for our first program. Unlike in C or C++, it can be named anything you prefer. The builder will determine the entry point. In our simple example, **gprbuild**, GNAT's builder, will use the file you passed as parameter.
- `Put_Line` is a procedure, just like `Greet`, except it is declared in the `Ada.Text_IO` module. It is the Ada equivalent of C's `printf`.
- Comments start with `--` and go to the end of the line. There is no multi-line comment syntax, that is, it is not possible to start a comment in one line and continue it in the next line. The only way to create multiple lines of comments in Ada is by using `--` on each line. For example:

```
-- We start a comment in this line...
-- and we continue on the second line...
```

In other languages

Procedures are similar to functions in C or C++ that return `void`. We'll see later how to declare functions in Ada.

Here is a minor variant of the "Hello, World" example:

Listing 2: greet.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet is
4 begin
5     -- Print "Hello, World!" to the screen
6     Put_Line ("Hello, World!");
7 end Greet;
```

Build output

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

Runtime output

```
Hello, World!
```

This version utilizes an Ada feature known as a use clause, which has the form `use package-name`. As illustrated by the call on `Put_Line`, the effect is that entities from the named package can be referenced directly, without the `package-name.` prefix.

2.2 Imperative language - If/Then/Else

This section describes Ada's `if` statement and introduces several other fundamental language facilities including integer I/O, data declarations, and subprogram parameter modes.

Ada's `if` statement is pretty unsurprising in form and function:

Listing 3: `check_positive.adb`

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Integer_Text_IO; use Ada.Integer_Text_IO;
3
4 procedure Check_Positive is
5   N : Integer;
6 begin
7   Put ("Enter an integer value: "); -- Put a String
8   Get (N); -- Read in an integer value
9   if N > 0 then
10    Put (N); -- Put an Integer
11    Put_Line (" is a positive number");
12  end if;
13 end Check_Positive;
```

The `if` statement minimally consists of the reserved word `if`, a condition (which must be a Boolean value), the reserved word `then` and a non-empty sequence of statements (the `then` part) which is executed if the condition evaluates to `True`, and a terminating `end if`.

This example declares an integer variable `N`, prompts the user for an integer, checks if the value is positive and, if so, displays the integer's value followed by the string " is a positive number". If the value is not positive, the procedure does not display any output.

The type `Integer` is a predefined signed type, and its range depends on the computer architecture. On typical current processors `Integer` is 32-bit signed.

The example illustrates some of the basic functionality for integer input-output. The relevant subprograms are in the predefined package `Ada.Integer_Text_IO` and include the `Get` procedure (which reads in a number from the keyboard) and the `Put` procedure (which displays an integer value).

Here's a slight variation on the example, which illustrates an `if` statement with an `else` part:

Listing 4: `check_positive.adb`

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Integer_Text_IO; use Ada.Integer_Text_IO;
3
4 procedure Check_Positive is
5   N : Integer;
6 begin
7   Put ("Enter an integer value: "); -- Put a String
8   Get (N); -- Reads in an integer value
9   Put (N); -- Put an Integer
10  if N > 0 then
11    Put_Line (" is a positive number");
12  else
13    Put_Line (" is not a positive number");
14  end if;
15 end Check_Positive;
```

In this example, if the input value is not positive then the program displays the value followed by the String " is not a positive number".

Our final variation illustrates an `if` statement with `elsif` sections:

Listing 5: check_direction.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Integer_Text_IO; use Ada.Integer_Text_IO;
3
4 procedure Check_Direction is
5   N : Integer;
6 begin
7   Put ("Enter an integer value: "); -- Puts a String
8   Get (N); -- Reads an Integer
9   Put (N); -- Puts an Integer
10  if N = 0 or N = 360 then
11    Put_Line (" is due east");
12  elsif N in 1 .. 89 then
13    Put_Line (" is in the northeast quadrant");
14  elsif N = 90 then
15    Put_Line (" is due north");
16  elsif N in 91 .. 179 then
17    Put_Line (" is in the northwest quadrant");
18  elsif N = 180 then
19    Put_Line (" is due west");
20  elsif N in 181 .. 269 then
21    Put_Line (" is in the southwest quadrant");
22  elsif N = 270 then
23    Put_Line (" is due south");
24  elsif N in 271 .. 359 then
25    Put_Line (" is in the southeast quadrant");
26  else
27    Put_Line (" is not in the range 0..360");
28  end if;
29 end Check_Direction;
```

This example expects the user to input an integer between 0 and 360 inclusive, and displays which quadrant or axis the value corresponds to. The `in` operator in Ada tests whether a scalar value is within a specified range and returns a Boolean result. The effect of the program should be self-explanatory; later we'll see an alternative and more efficient style to accomplish the same effect, through a case statement.

Ada's `elsif` keyword differs from C or C++, where nested `else .. if` blocks would be used instead. And another difference is the presence of the `end if` in Ada, which avoids the problem known as the "dangling else".

2.3 Imperative language - Loops

Ada has three ways of specifying loops. They differ from the C / Java / Javascript for-loop, however, with simpler syntax and semantics in line with Ada's philosophy.

2.3.1 For loops

The first kind of loop is the for loop, which allows iteration through a discrete range.

Listing 6: greet_5a.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet_5a is
4 begin
5   for I in 1 .. 5 loop
6     Put_Line ("Hello, World!" & Integer'Image (I)); -- Procedure call
7     --           ^ Procedure parameter
8   end loop;
9 end Greet_5a;
```

Build output

Compile	
[Ada]	greet_5a.adb
Bind	
[gprbind]	greet_5a.bexch
[Ada]	greet_5a.ali
Link	
[link]	greet_5a.adb

Runtime output

```

Hello, World! 1
Hello, World! 2
Hello, World! 3
Hello, World! 4
Hello, World! 5
```

A few things to note:

- 1 .. 5 is a discrete range, from 1 to 5 inclusive.
- The loop parameter I (the name is arbitrary) in the body of the loop has a value within this range.
- I is local to the loop, so you cannot refer to I outside the loop.
- Although the value of I is incremented at each iteration, from the program's perspective it is constant. An attempt to modify its value is illegal; the compiler would reject the program.
- Integer'Image is a function that takes an Integer and converts it to a String. It is an example of a language construct known as an *attribute*, indicated by the ' syntax, which will be covered in more detail later.
- The & symbol is the concatenation operator for String values
- The end loop marks the end of the loop

The "step" of the loop is limited to 1 (forward direction) and -1 (backward). To iterate backwards over a range, use the reverse keyword:

Listing 7: greet_5a_reverse.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet_5a_Reverse is
4 begin
5   for I in reverse 1 .. 5 loop
```

(continues on next page)

(continued from previous page)

```

6     Put_Line ("Hello, World!" & Integer'Image (I));
7     end loop;
8 end Greet_5a_Reverse;

```

Build output

```

Compile
  [Ada]          greet_5a_reverse.adb
Bind
  [gprbind]     greet_5a_reverse.bexch
  [Ada]         greet_5a_reverse.ali
Link
  [link]       greet_5a_reverse.adb

```

Runtime output

```

Hello, World! 5
Hello, World! 4
Hello, World! 3
Hello, World! 2
Hello, World! 1

```

The bounds of a `for` loop may be computed at run-time; they are evaluated once, before the loop body is executed. If the value of the upper bound is less than the value of the lower bound, then the loop is not executed at all. This is the case also for reverse loops. Thus no output is produced in the following example:

Listing 8: `greet_no_op.adb`

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet_No_Op is
4 begin
5   for I in reverse 5 .. 1 loop
6     Put_Line ("Hello, World!" & Integer'Image (I));
7   end loop;
8 end Greet_No_Op;

```

Build output

```

Compile
  [Ada]          greet_no_op.adb
greet_no_op.adb:5:23: warning: loop range is null, loop will not execute
Bind
  [gprbind]     greet_no_op.bexch
  [Ada]         greet_no_op.ali
Link
  [link]       greet_no_op.adb

```

The `for` loop is more general than what we illustrated here; more on that later.

2.3.2 Bare loops

The simplest loop in Ada is the bare loop, which forms the foundation of the other kinds of Ada loops.

Listing 9: greet_5b.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet_5b is
4   I : Integer := 1; -- Variable declaration
5     -- ^ Type
6     -- ^ Initial value
7 begin
8   loop
9     Put_Line ("Hello, World!" & Integer'Image (I));
10    exit when I = 5; -- Exit statement
11    -- ^ Boolean condition
12
13    -- Assignment
14    I := I + 1; -- There is no I++ short form to increment a variable
15  end loop;
16 end Greet_5b;

```

Build output

Compile	
[Ada]	greet_5b.adb
Bind	
[gprbind]	greet_5b.bexch
[Ada]	greet_5b.ali
Link	
[link]	greet_5b.adb

Runtime output

```

Hello, World! 1
Hello, World! 2
Hello, World! 3
Hello, World! 4
Hello, World! 5

```

This example has the same effect as Greet_5a shown earlier.

It illustrates several concepts:

- We have declared a variable named `I` between the `is` and the `begin`. This constitutes a *declarative region*. Ada clearly separates the declarative region from the statement part of a subprogram. A declaration can appear in a declarative region but is not allowed as a statement.
- The bare loop statement is introduced by the keyword `loop` on its own and, like every kind of loop statement, is terminated by the combination of keywords `end loop`. On its own, it is an infinite loop. You can break out of it with an `exit` statement.
- The syntax for assignment is `:=`, and the one for equality is `=`. There is no way to confuse them, because as previously noted, in Ada, statements and expressions are distinct, and expressions are not valid statements.

2.3.3 While loops

The last kind of loop in Ada is the while loop.

Listing 10: greet_5c.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet_5c is
4   I : Integer := 1;
5 begin
6   -- Condition must be a Boolean value (no Integers).
7   -- Operator "<=" returns a Boolean
8   while I <= 5 loop
9     Put_Line ("Hello, World!" & Integer'Image (I));
10
11     I := I + 1;
12   end loop;
13 end Greet_5c;
```

Build output

```
Compile
  [Ada]      greet_5c.adb
Bind
  [gprbind]  greet_5c.bexch
  [Ada]      greet_5c.ali
Link
  [link]     greet_5c.adb
```

Runtime output

```
Hello, World! 1
Hello, World! 2
Hello, World! 3
Hello, World! 4
Hello, World! 5
```

The condition is evaluated before each iteration. If the result is false, then the loop is terminated.

This program has the same effect as the previous examples.

In other languages

Note that Ada has different semantics than C-based languages with respect to the condition in a while loop. In Ada the condition has to be a Boolean value or the compiler will reject the program; the condition is not an integer that is treated as either True or False depending on whether it is non-zero or zero.

2.4 Imperative language - Case statement

Ada's case statement is similar to the C and C++ switch statement, but with some important differences.

Here's an example, a variation of a program that was shown earlier with an if statement:

Listing 11: check_direction.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Integer_Text_IO; use Ada.Integer_Text_IO;
3
4 procedure Check_Direction is
5   N : Integer;
6 begin
7   loop
8     Put ("Enter an integer value: "); -- Puts a String
9     Get (N); -- Reads an Integer
10    Put (N); -- Puts an Integer
11    case N is
12      when 0 | 360 =>
13        Put_Line (" is due east");
14      when 1 .. 89 =>
15        Put_Line (" is in the northeast quadrant");
16      when 90 =>
17        Put_Line (" is due north");
18      when 91 .. 179 =>
19        Put_Line (" is in the northwest quadrant");
20      when 180 =>
21        Put_Line (" is due west");
22      when 181 .. 269 =>
23        Put_Line (" is in the southwest quadrant");
24      when 270 =>
25        Put_Line (" is due south");
26      when 271 .. 359 =>
27        Put_Line (" is in the southeast quadrant");
28      when others =>
29        Put_Line (" Au revoir");
30        exit;
31    end case;
32  end loop;
33 end Check_Direction;
```

This program repeatedly prompts for an integer value and then, if the value is in the range 0..360, displays the associated quadrant or axis. If the value is an Integer outside this range, the loop (and the program) terminate after outputting a farewell message.

The effect of the case statement is similar to the if statement in an earlier example, but the case statement can be more efficient because it does not involve multiple range tests.

Notable points about Ada's case statement:

- The case expression (here the variable N) must be of a discrete type, i.e. either an integer type or an enumeration type. Discrete types will be covered in more detail later *discrete types* (page 41).
- Every possible value for the case expression needs to be covered by a unique branch of the case statement. This will be checked at compile time.
- A branch can specify a single value, such as 0; a range of values, such as 1 .. 89; or any combination of the two (separated by a |).
- As a special case, an optional final branch can specify others, which covers all values not included in the earlier branches.

- Execution consists of the evaluation of the case expression and then a transfer of control to the statement sequence in the unique branch that covers that value.
- When execution of the statements in the selected branch has completed, control resumes after the `end` case. Unlike C, execution does not fall through to the next branch. So Ada doesn't need (and doesn't have) a `break` statement.

2.5 Imperative language - Declarative regions

As mentioned earlier, Ada draws a clear syntactic separation between declarations, which introduce names for entities that will be used in the program, and statements, which perform the processing. The areas in the program where declarations may appear are known as declarative regions.

In any subprogram, the section between the `is` and the `begin` is a declarative region. You can have variables, constants, types, inner subprograms, and other entities there.

We've briefly mentioned variable declarations in previous subsection. Let's look at a simple example, where we declare an integer variable `X` in the declarative region and perform an initialization and an addition on it:

Listing 12: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   X : Integer;
5 begin
6   X := 0;
7   Put_Line ("The initial value of X is " & Integer'Image (X));
8
9   Put_Line ("Performing operation on X...");
10  X := X + 1;
11
12  Put_Line ("The value of X now is " & Integer'Image (X));
13 end Main;
```

Build output

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

Runtime output

```
The initial value of X is 0
Performing operation on X...
The value of X now is 1
```

Let's look at an example of a nested procedure:

Listing 13: main.adb

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

(continues on next page)

(continued from previous page)

```

4   procedure Nested is
5   begin
6       Put_Line ("Hello World");
7   end Nested;
8   begin
9       Nested;
10      -- Call to Nested
11 end Main;

```

Build output

```

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

```

Runtime output

```
Hello World
```

A declaration cannot appear as a statement. If you need to declare a local variable amidst the statements, you can introduce a new declarative region with a block statement:

Listing 14: greet.adb

```

1   with Ada.Text_IO; use Ada.Text_IO;
2
3   procedure Greet is
4   begin
5       loop
6           Put_Line ("Please enter your name: ");
7
8           declare
9               Name : String := Get_Line;
10              --           ^ Call to the Get_Line function
11           begin
12               exit when Name = "";
13               Put_Line ("Hi " & Name & "!");
14           end;
15
16           -- Name is undefined here
17       end loop;
18
19       Put_Line ("Bye!");
20 end Greet;

```

Attention: The `Get_Line` function allows you to receive input from the user, and get the result as a string. It is more or less equivalent to the `scanf` C function.

It returns a `String`, which, as we will see later, is an *Unconstrained array type* (page 71). For now we simply note that, if you wish to declare a `String` variable and do not know its size in advance, then you need to initialize the variable during its declaration.

2.6 Imperative language - conditional expressions

Ada 2012 introduced an expression analog for conditional statements (if and case).

2.6.1 If expressions

Here's an alternative version of an example we saw earlier; the if statement has been replaced by an if expression:

Listing 15: check_positive.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Integer_Text_IO; use Ada.Integer_Text_IO;
3
4 procedure Check_Positive is
5   N : Integer;
6 begin
7   Put ("Enter an integer value: "); -- Put a String
8   Get (N); -- Reads in an integer value
9   Put (N); -- Put an Integer
10  declare
11    S : String :=
12      (if N > 0 then " is a positive number"
13       else " is not a positive number");
14  begin
15    Put_Line (S);
16  end;
17 end Check_Positive;
```

The if expression evaluates to one of the two Strings depending on N, and assigns that value to the local variable S.

Ada's if expressions are similar to if statements. However, there are a few differences that stem from the fact that it is an expression:

- All branches' expressions must be of the same type
- It *must* be surrounded by parentheses if the surrounding expression does not already contain them
- An else branch is mandatory unless the expression following then has a Boolean value. In that case an else branch is optional and, if not present, defaults to else True.

Here's another example:

Listing 16: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4 begin
5   for I in 1 .. 10 loop
6     Put_Line (if I mod 2 = 0 then "Even" else "Odd");
7   end loop;
8 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

```

Odd
Even
Odd
Even
Odd
Even
Odd
Even
Odd
Even

```

This program produces 10 lines of output, alternating between "Odd" and "Even".

2.6.2 Case expressions

Analogous to if expressions, Ada also has case expressions. They work just as you would expect.

Listing 17: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4 begin
5   for I in 1 .. 10 loop
6     Put_Line (case I is
7               when 1 | 3 | 5 | 7 | 9 => "Odd",
8               when 2 | 4 | 6 | 8 | 10 => "Even");
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

```

Odd
Even
Odd
Even
Odd
Even
Odd
Even
Odd
Even

```

This program has the same effect as the preceding example.

The syntax differs from case statements, with branches separated by commas.

SUBPROGRAMS

3.1 Subprograms

So far, we have used procedures, mostly to have a main body of code to execute. Procedures are one kind of *subprogram*.

There are two kinds of subprograms in Ada, *functions* and *procedures*. The distinction between the two is that a function returns a value, and a procedure does not.

This example shows the declaration and definition of a function:

Listing 1: increment.ads

```
1 function Increment (I : Integer) return Integer;  
2 -- We declare (but don't define) a function with one  
3 -- parameter, returning an integer value
```

Listing 2: increment.adb

```
1 -- We define the Increment function  
2  
3 function Increment (I : Integer) return Integer is  
4 begin  
5     return I + 1;  
6 end Increment;
```

Subprograms in Ada can, of course, have parameters. One syntactically important note is that a subprogram which has no parameters does not have a parameter section at all, for example:

```
procedure Proc;  
  
function Func return Integer;
```

Here's another variation on the previous example:

Listing 3: increment_by.ads

```
1 function Increment_By  
2   (I      : Integer := 0;  
3    Incr   : Integer := 1) return Integer;  
4 --      ^ Default value for parameters
```

In this example, we see that parameters can have default values. When calling the subprogram, you can then omit parameters if they have a default value. Unlike C/C++, a call to a subprogram without parameters does not include parentheses.

This is the implementation of the function above:

Listing 4: increment_by.adb

```
1 function Increment_By
2   (I   : Integer := 0;
3    Incr : Integer := 1) return Integer is
4 begin
5   return I + Incr;
6 end Increment_By;
```

3.1.1 Subprogram calls

We can then call our subprogram this way:

Listing 5: show_increment.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Increment_By;
3
4 procedure Show_Increment is
5   A, B, C : Integer;
6 begin
7   C := Increment_By;
8   --      ^ Parameterless call, value of I is 0
9   --      and Incr is 1
10
11   Put_Line ("Using defaults for Increment_By is "
12             & Integer'Image (C));
13
14   A := 10;
15   B := 3;
16   C := Increment_By (A, B);
17   --      ^ Regular parameter passing
18
19   Put_Line ("Increment of " & Integer'Image (A)
20             & " with "      & Integer'Image (B)
21             & " is "       & Integer'Image (C));
22
23   A := 20;
24   B := 5;
25   C := Increment_By (I => A,
26                     Incr => B);
27   --      ^ Named parameter passing
28
29   Put_Line ("Increment of " & Integer'Image (A)
30             & " with "      & Integer'Image (B)
31             & " is "       & Integer'Image (C));
32 end Show_Increment;
```

Build output

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

Runtime output

```
Using defaults for Increment_By is 1
Increment of 10 with 3 is 13
Increment of 20 with 5 is 25
```

Ada allows you to name the parameters when you pass them, whether they have a default or not. There are some rules:

- Positional parameters come first.
- A positional parameter cannot follow a named parameter.

As a convention, people usually name parameters at the call site if the function's corresponding parameters has a default value. However, it is also perfectly acceptable to name every parameter if it makes the code clearer.

3.1.2 Nested subprograms

As briefly mentioned earlier, Ada allows you to declare one subprogram inside of another.

This is useful for two reasons:

- It lets you organize your programs in a cleaner fashion. If you need a subprogram only as a "helper" for another subprogram, then the principle of localization indicates that the helper subprogram should be declared nested.
- It allows you to share state easily in a controlled fashion, because the nested subprograms have access to the parameters, as well as any local variables, declared in the outer scope.

For the previous example, we can move the duplicated code (call to `Put_Line`) to a separate procedure. This is a shortened version with the nested `Display_Result` procedure.

Listing 6: show_increment.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Increment_By;
3
4 procedure Show_Increment is
5   A, B, C : Integer;
6
7   procedure Display_Result is
8     begin
9       Put_Line ("Increment of " & Integer'Image (A)
10                & " with " & Integer'Image (B)
11                & " is " & Integer'Image (C));
12     end Display_Result;
13
14 begin
15   A := 10;
16   B := 3;
17   C := Increment_By (A, B);
18   Display_Result;
19 end Show_Increment;
```

Build output

```

Compile
  [Ada]      show_increment.adb
  [Ada]      increment_by.adb
Bind
  [gprbind]  show_increment.bexch
  [Ada]      show_increment.ali
```

(continues on next page)

(continued from previous page)

```
Link  
[link]          show_increment.adb
```

Runtime output

```
Increment of 10 with 3 is 13
```

3.1.3 Function calls

An important feature of function calls in Ada is that the return value at a call cannot be ignored; that is, a function call cannot be used as a statement.

If you want to call a function and do not need its result, you will still need to explicitly store it in a local variable.

Listing 7: quadruple.adb

```
1 function Quadruple (I : Integer) return Integer is  
2   function Double (I : Integer) return Integer is  
3     begin  
4       return I * 2;  
5     end Double;  
6  
7   Res : Integer := Double (Double (I));  
8   --           ^ Calling the double function  
9 begin  
10  Double (I);  
11  -- ERROR: cannot use call to function "Double" as a statement  
12  
13  return Res;  
14 end Quadruple;
```

Build output

```
Compile  
[Ada]          quadruple.adb  
quadruple.adb:10:04: cannot use call to function "Double" as a statement  
quadruple.adb:10:04: return value of a function call cannot be ignored  
gprbuild: *** compilation phase failed
```

In the GNAT toolchain

In GNAT, with all warnings activated, it becomes even harder to ignore the result of a function, because unused variables will be flagged. For example, this code would not be valid:

```
function Read_Int  
  (Stream : Network_Stream; Result : out Integer) return Boolean;  
  
procedure Main is  
  Stream : Network_Stream := Get_Stream;  
  My_Int : Integer;  
  B : Boolean := Read_Int (Stream, My_Int); -- Warning here, B is never read  
begin  
  null;  
end Main;
```

You then have two solutions to silence this warning:

- Either annotate the variable with pragma Unreferenced, thus:

```
B : Boolean := Read_Int (Stream, My_Int);
pragma Unreferenced (B);
```

- Or give the variable a name that contains any of the strings discard dummy ignore junk unused (case insensitive)

3.2 Parameter modes

So far we have seen that Ada is a safety-focused language. There are many ways this is realized, but two important points are:

- Ada makes the user specify as much as possible about the behavior expected for the program, so that the compiler can warn or reject if there is an inconsistency.
- Ada provides a variety of techniques for achieving the generality and flexibility of pointers and dynamic memory management, but without the latter's drawbacks (such as memory leakage and dangling references).

Parameter modes are a feature that helps achieve the two design goals above. A subprogram parameter can be specified with a mode, which is one of the following:

in	Parameter can only be read, not written
out	Parameter can be written to, then read
in out	Parameter can be both read and written

The default mode for parameters is `in`; so far, most of the examples have been using `in` parameters.

Historically

Functions and procedures were originally more different in philosophy. Before Ada 2012, functions could only take "in" parameters.

3.3 Subprogram calls

3.3.1 In parameters

The first mode for parameters is the one we have been implicitly using so far. Parameters passed using this mode cannot be modified, so that the following program will cause an error:

Listing 8: swap.adb

```
1 procedure Swap (A, B : Integer) is
2   Tmp : Integer;
3 begin
4   Tmp := A;
5
6   -- Error: assignment to "in" mode parameter not allowed
7   A := B;
8   -- Error: assignment to "in" mode parameter not allowed
9   B := Tmp;
10 end Swap;
```

Build output

```

Compile
  [Ada]          swap.adb
swap.adb:7:04: assignment to "in" mode parameter not allowed
swap.adb:9:04: assignment to "in" mode parameter not allowed
gprbuild: *** compilation phase failed

```

The fact that this is the default mode is in itself very important. It means that a parameter will not be modified unless you explicitly specify a mode in which modification is allowed.

3.3.2 In out parameters

To correct our code above, we can use an "in out" parameter.

Listing 9: in_out_params.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure In_Out_Params is
4   procedure Swap (A, B : in out Integer) is
5     Tmp : Integer;
6     begin
7       Tmp := A;
8       A := B;
9       B := Tmp;
10    end Swap;
11
12    A : Integer := 12;
13    B : Integer := 44;
14  begin
15    Swap (A, B);
16    Put_Line (Integer'Image (A)); -- Prints 44
17  end In_Out_Params;

```

Build output

```

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

```

Runtime output

```
44
```

An in out parameter will allow read and write access to the object passed as parameter, so in the example above, we can see that A is modified after the call to Swap.

Attention: While in out parameters look a bit like references in C++, or regular parameters in Java that are passed by-reference, the Ada language standard does not mandate "by reference" passing for in out parameters except for certain categories of types as will be explained later.

In general, it is better to think of modes as higher level than by-value versus by-reference semantics. For the compiler, it means that an array passed as an in parameter might be passed by reference, because it is more efficient (which does not change anything for the user since the

parameter is not assignable). However, a parameter of a discrete type will always be passed by copy, regardless of its mode (which is more efficient on most architectures).

3.3.3 Out parameters

The "out" mode applies when the subprogram needs to write to a parameter that might be uninitialized at the point of call. Reading the value of an out parameter is permitted, but it should only be done after the subprogram has assigned a value to the parameter. Out parameters behave a bit like return values for functions. When the subprogram returns, the actual parameter (a variable) will have the value of the out parameter at the point of return.

In other languages

Ada doesn't have a tuple construct and does not allow returning multiple values from a subprogram (except by declaring a full-fledged record type). Hence, a way to return multiple values from a subprogram is to use out parameters.

For example, a procedure reading integers from the network could have one of the following specifications:

```

procedure Read_Int
  (Stream : Network_Stream; Success : out Boolean; Result : out Integer);

function Read_Int
  (Stream : Network_Stream; Result : out Integer) return Boolean;

```

While reading an out variable before writing to it should, ideally, trigger an error, imposing that as a rule would cause either inefficient run-time checks or complex compile-time rules. So from the user's perspective an out parameter acts like an uninitialized variable when the subprogram is invoked.

In the GNAT toolchain

GNAT will detect simple cases of incorrect use of out parameters. For example, the compiler will emit a warning for the following program:

Listing 10: outp.adb

```

1 procedure Outp is
2   procedure Foo (A : out Integer) is
3     B : Integer := A; -- Warning on reference to uninitialized A
4   begin
5     A := B;
6   end Foo;
7 begin
8   null;
9 end Outp;

```

Compilation output

```
outp.adb:3:22: warning: "A" may be referenced before it has a value
```

3.3.4 Forward declaration of subprograms

As we saw earlier, a subprogram can be declared without being fully defined, This is possible in general, and can be useful if you need subprograms to be mutually recursive, as in the example below:

Listing 11: mutually_recursive_subprograms.adb

```

1  procedure Mutually_Recursive_Subprograms is
2      procedure Compute_A (V : Natural);
3      -- Forward declaration of Compute_A
4
5      procedure Compute_B (V : Natural) is
6      begin
7          if V > 5 then
8              Compute_A (V - 1);
9              -- Call to Compute_A
10         end if;
11     end Compute_B;
12
13     procedure Compute_A (V : Natural) is
14     begin
15         if V > 2 then
16             Compute_B (V - 1);
17             -- Call to Compute_B
18         end if;
19     end Compute_A;
20 begin
21     Compute_A (15);
22 end Mutually_Recursive_Subprograms;

```

Build output

```

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

```

3.4 Renaming

Subprograms can be renamed by using the `renames` keyword and declaring a new name for a subprogram:

```

procedure New_Proc renames Original_Proc;

```

This can be useful, for example, to improve the readability of your application when you're using code from external sources that cannot be changed in your system. Let's look at an example:

Listing 12: a_procedure_with_very_long_name_that_cannot_be_changed.ads

```

1  procedure A_Procedure_With_Very_Long_Name_That_Cannot_Be_Changed
2      (A_Message : String);

```


Listing 13: a_procedure_with_very_long_name_that_cannot_be_changed.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure A_Procedure_With_Very_Long_Name_That_Cannot_Be_Changed
4   (A_Message : String) is
5 begin
6   Put_Line (A_Message);
7 end A_Procedure_With_Very_Long_Name_That_Cannot_Be_Changed;

```

As the wording in the name of procedure above implies, we cannot change its name. We can, however, rename it to something like `Show` in our test application and use this shorter name. Note that we also have to declare all parameters of the original subprogram — we may rename them, too, in the declaration. For example:

Listing 14: show_renaming.adb

```

1 with A_Procedure_With_Very_Long_Name_That_Cannot_Be_Changed;
2
3 procedure Show_Renaming is
4
5   procedure Show (S : String) renames
6     A_Procedure_With_Very_Long_Name_That_Cannot_Be_Changed;
7
8 begin
9   Show ("Hello World!");
10 end Show_Renaming;

```

Build output

```

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

```

Runtime output

```
Hello World!
```

Note that the original name (`A_Procedure_With_Very_Long_Name_That_Cannot_Be_Changed`) is still visible after the declaration of the `Show` procedure.

We may also rename subprograms from the standard library. For example, we may rename `Integer'Image` to `Img`:

Listing 15: show_image_renaming.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Image_Renaming is
4
5   function Img (I : Integer) return String renames Integer'Image;
6
7 begin
8   Put_Line (Img (2));
9   Put_Line (Img (3));
10 end Show_Image_Renaming;

```

Build output

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

Runtime output

```
2
3
```

Renaming also allows us to introduce default expressions that were not available in the original declaration. For example, we may specify "Hello World!" as the default for the `String` parameter of the `Show` procedure:

```
with A_Procedure_With_Very_Long_Name_That_Cannot_Be_Changed;

procedure Show_Renaming_Defaults is

  procedure Show (S : String := "Hello World!") renames
    A_Procedure_With_Very_Long_Name_That_Cannot_Be_Changed;

begin
  Show;
end Show_Renaming_Defaults;
```

MODULAR PROGRAMMING

So far, our examples have been simple standalone subprograms. Ada is helpful in that regard, since it allows arbitrary declarations in a declarative part. We were thus able to declare our types and variables in the bodies of main procedures.

However, it is easy to see that this is not going to scale up for real-world applications. We need a better way to structure our programs into modular and distinct units.

Ada encourages the separation of programs into multiple packages and sub-packages, providing many tools to a programmer on a quest for a perfectly organized code-base.

4.1 Packages

Here is an example of a package declaration in Ada:

Listing 1: week.ads

```
1 package Week is
2
3     Mon : constant String := "Monday";
4     Tue : constant String := "Tuesday";
5     Wed : constant String := "Wednesday";
6     Thu : constant String := "Thursday";
7     Fri : constant String := "Friday";
8     Sat : constant String := "Saturday";
9     Sun : constant String := "Sunday";
10
11 end Week;
```

And here is how you use it:

Listing 2: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Week;
3 -- References the Week package, and adds a dependency from Main
4 -- to Week
5
6 procedure Main is
7 begin
8     Put_Line ("First day of the week is " & Week.Mon);
9 end Main;
```

Build output

```
Compile
  [Ada]      main.adb
```

(continues on next page)

(continued from previous page)

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

Runtime output

```
First day of the week is Monday
```

Packages let you make your code modular, separating your programs into semantically significant units. Additionally the separation of a package's specification from its body (which we will see below) can reduce compilation time.

While the `with` clause indicates a dependency, you can see in the example above that you still need to prefix the referencing of entities from the `Week` package by the name of the package. (If we had included a `use Week` clause, then such a prefix would not have been necessary.)

Accessing entities from a package uses the dot notation, `A.B`, which is the same notation as the one used to access record fields.

A `with` clause can *only* appear in the prelude of a compilation unit (i.e., before the reserved word, such as `procedure`, that marks the beginning of the unit). It is not allowed anywhere else. This rule is only needed for methodological reasons: the person reading your code should be able to see immediately which units the code depends on.

In other languages

Packages look similar to, but are semantically very different from, header files in `C/C++`.

- The first and most important distinction is that packages are a language-level mechanism. This is in contrast to a `#include`'d header file, which is a functionality of the C preprocessor.
- An immediate consequence is that the `with` construct is a semantic inclusion mechanism, not a text inclusion mechanism. Hence, when you `with` a package, you are saying to the compiler "I'm depending on this semantic unit", and not "include this bunch of text in place here".
- The effect of a package thus does not vary depending on where it has been `with`ed from. Contrast this with `C/C++`, where the meaning of the included text depends on the context in which the `#include` appears.

This allows compilation/recompilation to be more efficient. It also allows tools like IDEs to have correct information about the semantics of a program. In turn, this allows better tooling in general, and code that is more analyzable, even by humans.

An important benefit of Ada `with` clauses when compared to `#include` is that it is stateless. The order of `with` and `use` clauses does not matter, and can be changed without side effects.

In the GNAT toolchain

The Ada language standard does not mandate any particular relationship between source files and packages; for example, in theory you can put all your code in one file, or use your own file naming conventions. In practice, however, an implementation will have specific rules. With GNAT, each top-level compilation unit needs to go into a separate file. In the example above, the `Week` package will be in an `.ads` file (for Ada specification), and the `Main` procedure will be in an `.adb` file (for Ada body).

4.2 Using a package

As we have seen above, the `with` clause indicates a dependency on another package. However, every reference to an entity coming from the `Week` package had to be prefixed by the full name of the package. It is possible to make every entity of a package visible directly in the current scope, using the `use` clause.

In fact, we have been using the `use` clause since almost the beginning of this tutorial.

Listing 3: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 --      ^ Make every entity of the Ada.Text_IO package
3 --      directly visible.
4 with Week;
5
6 procedure Main is
7   use Week;
8   -- Make every entity of the Week package directly visible.
9 begin
10  Put_Line ("First day of the week is " & Mon);
11 end Main;
```

Build output

```

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

Runtime output

```
First day of the week is Monday
```

As you can see in the example above:

- `Put_Line` is a subprogram that comes from the `Ada.Text_IO` package. We can reference it directly because we have used the package at the top of the `Main` unit.
- Unlike `with` clauses, a `use` clause can be placed either in the prelude, or in any declarative region. In the latter case the `use` clause will have an effect in its containing lexical scope.

4.3 Package body

In the simple example above, the `Week` package only has declarations and no body. That's not a mistake: in a package specification, which is what is illustrated above, you cannot declare bodies. Those have to be in the package body.

Listing 4: operations.ads

```

1 package Operations is
2
3   -- Declaration
4   function Increment_By
5     (I : Integer;
```

(continues on next page)

(continued from previous page)

```

6     Incr : Integer := 0) return Integer;
7
8     function Get_Increment_Value return Integer;
9
10  end Operations;

```

Listing 5: operations.adb

```

1  package body Operations is
2
3     Last_Increment : Integer := 1;
4
5     function Increment_By
6     (I      : Integer;
7      Incr  : Integer := 0) return Integer is
8  begin
9      if Incr /= 0 then
10         Last_Increment := Incr;
11     end if;
12
13     return I + Last_Increment;
14  end Increment_By;
15
16  function Get_Increment_Value return Integer is
17  begin
18     return Last_Increment;
19  end Get_Increment_Value;
20
21  end Operations;

```

Here we can see that the body of the `Increment_By` function has to be declared in the body. Coincidentally, introducing a body allows us to put the `Last_Increment` variable in the body, and make them inaccessible to the user of the `Operations` package, providing a first form of encapsulation.

This works because entities declared in the body are *only* visible in the body.

This example shows how `Last_Increment` is used indirectly:

Listing 6: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2  with Operations;
3
4  procedure Main is
5      use Operations;
6
7      I : Integer := 0;
8      R : Integer;
9
10     procedure Display_Update_Values is
11         Incr : constant Integer := Get_Increment_Value;
12     begin
13         Put_Line (Integer'Image (I)
14                 & " incremented by " & Integer'Image (Incr)
15                 & " is "           & Integer'Image (R));
16         I := R;
17     end Display_Update_Values;
18  begin
19     R := Increment_By (I);
20     Display_Update_Values;

```

(continues on next page)

(continued from previous page)

```

21   R := Increment_By (I);
22   Display_Update_Values;
23
24   R := Increment_By (I, 5);
25   Display_Update_Values;
26   R := Increment_By (I);
27   Display_Update_Values;
28
29   R := Increment_By (I, 10);
30   Display_Update_Values;
31   R := Increment_By (I);
32   Display_Update_Values;
33 end Main;

```

Build output

```

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

```

Runtime output

```

0 incremented by 1 is 1
1 incremented by 1 is 2
2 incremented by 5 is 7
7 incremented by 5 is 12
12 incremented by 10 is 22
22 incremented by 10 is 32

```

4.4 Child packages

Packages can be used to create hierarchies. We achieve this by using child packages, which extend the functionality of their parent package. One example of a child package that we've been using so far is the `Ada.Text_IO` package. Here, the parent package is called `Ada`, while the child package is called `Text_IO`. In the previous examples, we've been using the `Put_Line` procedure from the `Text_IO` child package.

Important

Ada also supports nested packages. However, since they can be more complicated to use, the recommendation is to use child packages instead. Nested packages will be covered in the advanced course.

Let's begin our discussion on child packages by taking our previous `Week` package:

Listing 7: `week.ads`

```

1 package Week is
2
3   Mon : constant String := "Monday";
4   Tue : constant String := "Tuesday";

```

(continues on next page)

(continued from previous page)

```

5   Wed : constant String := "Wednesday";
6   Thu : constant String := "Thursday";
7   Fri : constant String := "Friday";
8   Sat : constant String := "Saturday";
9   Sun : constant String := "Sunday";
10
11 end Week;
```

If we want to create a child package for Week, we may write:

Listing 8: week-child.ads

```

1 package Week.Child is
2
3     function Get_First_Of_Week return String;
4
5 end Week.Child;
```

Here, `Week` is the parent package and `Child` is the child package. This is the corresponding package body of `Week.Child`:

Listing 9: week-child.adb

```

1 package body Week.Child is
2
3     function Get_First_Of_Week return String is
4     begin
5         return Mon;
6     end Get_First_Of_Week;
7
8 end Week.Child;
```

In the implementation of the `Get_First_Of_Week` function, we can use the `Mon` string directly, even though it was declared in the parent package `Week`. We don't write `with Week` here because all elements from the specification of the `Week` package — such as `Mon`, `Tue` and so on — are visible in the child package `Week.Child`.

Now that we've completed the implementation of the `Week.Child` package, we can use elements from this child package in a subprogram by simply writing `with Week.Child`. Similarly, if we want to use these elements directly, we write `use Week.Child` in addition. For example:

Listing 10: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Week.Child; use Week.Child;
3
4 procedure Main is
5 begin
6     Put_Line ("First day of the week is " & Get_First_Of_Week);
7 end Main;
```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      week.ads
  [Ada]      week-child.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
```

(continues on next page)

(continued from previous page)

[link]	main.adb
--------	----------

Runtime output

First day of the week is Monday

4.4.1 Child of a child package

So far, we've seen a two-level package hierarchy. But the hierarchy that we can potentially create isn't limited to that. For instance, we could extend the hierarchy of the previous source-code example by declaring a `Week.Child.Grandchild` package. In this case, `Week.Child` would be the parent of the `Grandchild` package. Let's consider this implementation:

Listing 11: week-child-grandchild.ads

```

1 package Week.Child.Grandchild is
2
3     function Get_Second_Of_Week return String;
4
5 end Week.Child.Grandchild;
```

Listing 12: week-child-grandchild.adb

```

1 package body Week.Child.Grandchild is
2
3     function Get_Second_Of_Week return String is
4     begin
5         return Tue;
6     end Get_Second_Of_Week;
7
8 end Week.Child.Grandchild;
```

We can use this new `Grandchild` package in our test application in the same way as before: we can reuse the previous test application and adapt the `with` and `use`, and the function call. This is the updated code:

Listing 13: main.adb

```

1 with Ada.Text_IO;           use Ada.Text_IO;
2 with Week.Child.Grandchild; use Week.Child.Grandchild;
3
4 procedure Main is
5 begin
6     Put_Line ("Second day of the week is " & Get_Second_Of_Week);
7 end Main;
```

Build output

Compile	
[Ada]	main.adb
[Ada]	week.ads
[Ada]	week-child.adb
[Ada]	week-child-grandchild.adb
Bind	
[gprbind]	main.bexch
[Ada]	main.ali
Link	
[link]	main.adb

Runtime output

```
Second day of the week is Tuesday
```

Again, this isn't the limit for the package hierarchy. We could continue to extend the hierarchy of the previous example by implementing a `Week.Child.Grandchild.Grand_grandchild` package.

4.4.2 Multiple children

So far, we've seen a single child package of a parent package. However, a parent package can also have multiple children. We could extend the example above and implement a `Week.Child_2` package. For example:

Listing 14: week-child_2.ads

```
1 package Week.Child_2 is
2
3     function Get_Last_Of_Week return String;
4
5 end Week.Child_2;
```

Here, `Week` is still the parent package of the `Child` package, but it's also the parent of the `Child_2` package. In the same way, `Child_2` is obviously one of the child packages of `Week`.

This is the corresponding package body of `Week.Child_2`:

Listing 15: week-child_2.adb

```
1 package body Week.Child_2 is
2
3     function Get_Last_Of_Week return String is
4     begin
5         return Sun;
6     end Get_Last_Of_Week;
7
8 end Week.Child_2;
```

We can now reference both children in our test application:

Listing 16: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Week.Child; use Week.Child;
3 with Week.Child_2; use Week.Child_2;
4
5 procedure Main is
6 begin
7     Put_Line ("First day of the week is " & Get_First_Of_Week);
8     Put_Line ("Last day of the week is " & Get_Last_Of_Week);
9 end Main;
```

Build output

```
Compile
  [Ada]      main.adb
  [Ada]      week.ads
  [Ada]      week-child.adb
  [Ada]      week-child_2.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
```

(continues on next page)

(continued from previous page)

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

Runtime output

```
First day of the week is Monday
Last day of the week is Sunday
```

4.4.3 Visibility

In the previous section, we've seen that elements declared in a parent package specification are visible in the child package. This is, however, not the case for elements declared in the package body of a parent package.

Let's consider the package `Book` and its child `Additional_Operations`:

Listing 17: book.ads

```
1 package Book is
2
3   Title : constant String := "Visible for my children";
4
5   function Get_Title return String;
6
7   function Get_Author return String;
8
9 end Book;
```

Listing 18: book-additional_operations.ads

```
1 package Book.Additional_Operations is
2
3   function Get_Extended_Title return String;
4
5   function Get_Extended_Author return String;
6
7 end Book.Additional_Operations;
```

This is the body of both packages:

Listing 19: book.adb

```
1 package body Book is
2
3   Author : constant String := "Author not visible for my children";
4
5   function Get_Title return String is
6   begin
7     return Title;
8   end Get_Title;
9
10  function Get_Author return String is
11  begin
12    return Author;
13  end Get_Author;
14
15 end Book;
```

Listing 20: book-additional_operations.adb

```
1 package body Book.Additional_Operations is
2
3   function Get_Extended_Title return String is
4   begin
5     return "Book Title: " & Title;
6   end Get_Extended_Title;
7
8   function Get_Extended_Author return String is
9   begin
10    -- "Author" string declared in the body of Book package is not
11    -- visible here. Therefore, we cannot write:
12    --
13    -- return "Book Author: " & Author;
14
15    return "Book Author: Unknown";
16  end Get_Extended_Author;
17
18 end Book.Additional_Operations;
```

In the implementation of the `Get_Extended_Title`, we're using the `Title` constant from the parent package `Book`. However, as indicated in the comments of the `Get_Extended_Author` function, the `Author` string — which we declared in the body of the `Book` package — isn't visible in the `Book.Additional_Operations` package. Therefore, we cannot use it to implement the `Get_Extended_Author` function.

We can, however, use the `Get_Author` function from `Book` in the implementation of the `Get_Extended_Author` function to retrieve this string. Likewise, we can use this strategy to implement the `Get_Extended_Title` function. This is the adapted code:

Listing 21: book-additional_operations.adb

```
1 package body Book.Additional_Operations is
2
3   function Get_Extended_Title return String is
4   begin
5     return "Book Title: " & Get_Title;
6   end Get_Extended_Title;
7
8   function Get_Extended_Author return String is
9   begin
10    return "Book Author: " & Get_Author;
11  end Get_Extended_Author;
12
13 end Book.Additional_Operations;
```

This is a simple test application for the packages above:

Listing 22: main.adb

```
1 with Ada.Text_IO;           use Ada.Text_IO;
2 with Book.Additional_Operations; use Book.Additional_Operations;
3
4 procedure Main is
5 begin
6   Put_Line (Get_Extended_Title);
7   Put_Line (Get_Extended_Author);
8 end Main;
```

Build output

```

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

```

Runtime output

```

Book Title: Visible for my children
Book Author: Author not visible for my children

```

By declaring elements in the body of a package, we can implement encapsulation in Ada. Those elements will only be visible in the package body, but nowhere else. This isn't, however, the only way to achieve encapsulation in Ada: we'll discuss other approaches in the [Privacy](#) (page 101) chapter.

4.5 Renaming

Previously, we've mentioned that *subprograms can be renamed* (page 26). We can rename packages, too. Again, we use the `renames` keyword for that. The following example renames the `Ada.Text_IO` package as `T_IO`:

Listing 23: main.adb

```

1 with Ada.Text_IO;
2
3 procedure Main is
4   package TIO renames Ada.Text_IO;
5   begin
6     TIO.Put_Line ("Hello");
7   end Main;

```

Build output

```

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

```

Runtime output

```

Hello

```

We can use renaming to improve the readability of our code by using shorter package names. In the example above, we write `TIO.Put_Line` instead of the longer version (`Ada.Text_IO.Put_Line`). This approach is especially useful when we don't use packages and want to avoid that the code becomes too verbose.

Note we can also rename subprograms and objects inside packages. For instance, we could have just renamed the `Put_Line` procedure in the source-code example above:

Listing 24: main.adb

```
1 with Ada.Text_IO;  
2  
3 procedure Main is  
4   procedure Say (Something : String) renames Ada.Text_IO.Put_Line;  
5 begin  
6   Say ("Hello");  
7 end Main;
```

STRONGLY TYPED LANGUAGE

Ada is a strongly typed language. It is interestingly modern in that respect: strong static typing has been increasing in popularity in programming language design, owing to factors such as the growth of statically typed functional programming, a big push from the research community in the typing domain, and many practical languages with strong type systems.

5.1 What is a type?

In statically typed languages, a type is mainly (but not only) a *compile time* construct. It is a construct to enforce invariants about the behavior of a program. Invariants are unchangeable properties that hold for all variables of a given type. Enforcing them ensures, for example, that variables of a data type never have invalid values.

A type is used to reason about the *objects* a program manipulates (an object is a variable or a constant). The aim is to classify objects by what you can accomplish with them (i.e., the operations that are permitted), and this way you can reason about the correctness of the objects' values.

5.2 Integers

A nice feature of Ada is that you can define your own integer types, based on the requirements of your program (i.e., the range of values that makes sense). In fact, the definitional mechanism that Ada provides forms the semantic basis for the predefined integer types. There is no "magical" built-in type in that regard, which is unlike most languages, and arguably very elegant.

Listing 1: integer_type_example.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Integer_Type_Example is
4   -- Declare a signed integer type, and give the bounds
5   type My_Int is range -1 .. 20;
6     --                ^ High bound
7     --                ^ Low bound
8
9   -- Like variables, type declarations can only appear in
10  -- declarative regions
11 begin
12   for I in My_Int loop
13     Put_Line (My_Int'Image (I));
14     --                ^ 'Image attribute, converts a value to a
15     --                String
16   end loop;
17 end Integer_Type_Example;
```

Build output

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

Runtime output

```
-1
0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
```

This example illustrates the declaration of a signed integer type, and several things we can do with them.

Every type declaration in Ada starts with the `type` keyword (except for *task types* (page 147)). After the type, we can see a range that looks a lot like the ranges that we use in for loops, that defines the low and high bound of the type. Every integer in the inclusive range of the bounds is a valid value for the type.

Ada integer types

In Ada, an integer type is not specified in terms of its machine representation, but rather by its range. The compiler will then choose the most appropriate representation.

Another point to note in the above example is the `My_Int'Image (I)` expression. The `Name'Attribute (optional params)` notation is used for what is called an attribute in Ada. An attribute is a built-in operation on a type, a value, or some other program entity. It is accessed by using a `'` symbol (the ASCII apostrophe).

Ada has several types available as "built-ins"; Integer is one of them. Here is how Integer might be defined for a typical processor:

```
type Integer is range -(2 ** 31) .. +(2 ** 31 - 1);
```

`**` is the exponent operator, which means that the first valid value for Integer is -2^{31} , and the last valid value is $2^{31} - 1$.

Ada does not mandate the range of the built-in type `Integer`. An implementation for a 16-bit target would likely choose the range -2^{15} through $2^{15} - 1$.

5.2.1 Operational semantics

Unlike some other languages, Ada requires that operations on integers should be checked for overflow.

Listing 2: main.adb

```

1 procedure Main is
2   A : Integer := Integer'Last;
3   B : Integer;
4 begin
5   B := A + 5;
6   -- This operation will overflow, eg. it will
7   -- raise an exception at run time.
8 end Main;
```

Build output

```

Compile
  [Ada]          main.adb
main.adb:5:11: warning: value not in range of type "Standard.Integer"
main.adb:5:11: warning: "Constraint_Error" will be raised at run time
Bind
  [gprbind]     main.bexch
  [Ada]         main.ali
Link
  [link]       main.adb
```

Runtime output

```
raised CONSTRAINT_ERROR : main.adb:5 overflow check failed
```

There are two types of overflow checks:

- Machine-level overflow, when the result of an operation exceeds the maximum value (or is less than the minimum value) that can be represented in the storage reserved for an object of the type, and
- Type-level overflow, when the result of an operation is outside the range defined for the type.

Mainly for efficiency reasons, while machine level overflow always results in an exception, type level overflows will only be checked at specific boundaries, like assignment:

Listing 3: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   type My_Int is range 1 .. 20;
5   A : My_Int := 12;
6   B : My_Int := 15;
7   M : My_Int := (A + B) / 2;
8   -- No overflow here, overflow checks are done at
9   -- specific boundaries.
10 begin
11   for I in 1 .. M loop
12     Put_Line ("Hello, World!");
```

(continues on next page)

(continued from previous page)

```

13   end loop;
14   -- Loop body executed 13 times
15 end Main;

```

Build output

```

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

```

Runtime output

```

Hello, World!
Hello, World!
Hello, World!
Hello, World!
Hello, World!
Hello, World!
Hello, World!
Hello, World!
Hello, World!
Hello, World!
Hello, World!
Hello, World!
Hello, World!
Hello, World!
Hello, World!

```

Type level overflow will only be checked at specific points in the execution. The result, as we see above, is that you might have an operation that overflows in an intermediate computation, but no exception will be raised because the final result does not overflow.

5.3 Unsigned types

Ada also features unsigned Integer types. They're called *modular* types in Ada parlance. The reason for this designation is due to their behavior in case of overflow: They simply "wrap around", as if a modulo operation was applied.

For machine sized modular types, for example a modulus of 2^{32} , this mimics the most common implementation behavior of unsigned types. However, an advantage of Ada is that the modulus is more general:

Listing 4: main.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Main is
4     type Mod_Int is mod 2 ** 5;
5     --           ^ Range is 0 .. 31
6
7     A : Mod_Int := 20;
8     B : Mod_Int := 15;
9     M : Mod_Int := A + B;
10    -- No overflow here, M = (20 + 15) mod 32 = 3
11  begin
12    for I in 1 .. M loop

```

(continues on next page)

(continued from previous page)

```

13     Put_Line ("Hello, World!");
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

```

Hello, World!
Hello, World!
Hello, World!

```

Unlike in C/C++, since this behavior is guaranteed by the Ada specification, you can rely on it to implement portable code. Also, being able to leverage the wrapping on arbitrary bounds is very useful -- the modulus does not need to be a power of 2 -- to implement certain algorithms and data structures, such as [ring buffers](#)⁹.

5.4 Enumerations

Enumeration types are another nicety of Ada's type system. Unlike C's enums, they are *not* integers, and each new enumeration type is incompatible with other enumeration types. Enumeration types are part of the bigger family of discrete types, which makes them usable in certain situations that we will describe later but one context that we have already seen is a case statement.

Listing 5: enumeration_example.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Enumeration_Example is
4     type Days is (Monday, Tuesday, Wednesday,
5                  Thursday, Friday, Saturday, Sunday);
6     -- An enumeration type
7  begin
8     for I in Days loop
9         case I is
10            when Saturday .. Sunday =>
11                Put_Line ("Week end!");
12
13            when Monday .. Friday =>
14                Put_Line ("Hello on " & Days'Image (I));
15                -- 'Image attribute, works on enums too
16        end case;
17    end loop;
18 end Enumeration_Example;

```

Build output

⁹ https://en.m.wikipedia.org/wiki/Circular_buffer

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

Runtime output

```
Hello on MONDAY
Hello on TUESDAY
Hello on WEDNESDAY
Hello on THURSDAY
Hello on FRIDAY
Week end!
Week end!
```

Enumeration types are powerful enough that, unlike in most languages, they're used to define the standard Boolean type:

```
type Boolean is (False, True);
```

As mentioned previously, every "built-in" type in Ada is defined with facilities generally available to the user.

5.5 Floating-point types

5.5.1 Basic properties

Like most languages, Ada supports floating-point types. The most commonly used floating-point type is Float:

Listing 6: floating_point_demo.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Floating_Point_Demo is
4   A : Float := 2.5;
5 begin
6   Put_Line ("The value of A is " & Float'Image (A));
7 end Floating_Point_Demo;
```

Build output

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

Runtime output

```
The value of A is 2.50000E+00
```

The application will display 2.5 as the value of A.

The Ada language does not specify the precision (number of decimal digits in the mantissa) for `Float`; on a typical 32-bit machine the precision will be 6.

All common operations that could be expected for floating-point types are available, including absolute value and exponentiation. For example:

Listing 7: `floating_point_operations.adb`

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Floating_Point_Operations is
4   A : Float := 2.5;
5 begin
6   A := abs (A - 4.5);
7   Put_Line ("The value of A is " & Float'Image (A));
8   A := A ** 2 + 1.0;
9   Put_Line ("The value of A is " & Float'Image (A));
10 end Floating_Point_Operations;
```

Build output

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

Runtime output

```

The value of A is 2.00000E+00
The value of A is 5.00000E+00
```

The value of `A` is 2.0 after the first operation and 5.0 after the second operation.

In addition to `Float`, an Ada implementation may offer data types with higher precision such as `Long_Float` and `Long_Long_Float`. Like `Float`, the standard does not indicate the exact precision of these types: it only guarantees that the type `Long_Float`, for example, has at least the precision of `Float`. In order to guarantee that a certain precision requirement is met, we can define custom floating-point types, as we will see in the next section.

5.5.2 Precision of floating-point types

Ada allows the user to specify the precision for a floating-point type, expressed in terms of decimal digits. Operations on these custom types will then have at least the specified precision. The syntax for a simple floating-point type declaration is:

```
type T is digits <number_of_decimal_digits>;
```

The compiler will choose a floating-point representation that supports the required precision. For example:

Listing 8: `custom_floating_types.adb`

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Custom_Floating_Types is
4   type T3 is digits 3;
5   type T15 is digits 15;
6   type T18 is digits 18;
```

(continues on next page)

(continued from previous page)

```

7 begin
8   Put_Line ("T3 requires " & Integer'Image (T3'Size) & " bits");
9   Put_Line ("T15 requires " & Integer'Image (T15'Size) & " bits");
10  Put_Line ("T18 requires " & Integer'Image (T18'Size) & " bits");
11 end Custom_Floating_Types;

```

Build output

```

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

```

Runtime output

```

T3 requires 32 bits
T15 requires 64 bits
T18 requires 128 bits

```

In this example, the attribute 'Size is used to retrieve the number of bits used for the specified data type. As we can see by running this example, the compiler allocates 32 bits for T3, 64 bits for T15 and 128 bits for T18. This includes both the mantissa and the exponent.

The number of digits specified in the data type is also used in the format when displaying floating-point variables. For example:

Listing 9: display_custom_floating_types.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Display_Custom_Floating_Types is
4   type T3 is digits 3;
5   type T18 is digits 18;
6
7   C1 : constant := 1.0e-4;
8
9   A : T3 := 1.0 + C1;
10  B : T18 := 1.0 + C1;
11 begin
12   Put_Line ("The value of A is " & T3'Image (A));
13   Put_Line ("The value of B is " & T18'Image (B));
14 end Display_Custom_Floating_Types;

```

Build output

```

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

```

Runtime output

```

The value of A is 1.00E+00
The value of B is 1.000100000000000000E+00

```

As expected, the application will display the variables according to specified precision (1.00E+00 and 1.000100000000000000E+00).

5.5.3 Range of floating-point types

In addition to the precision, a range can also be specified for a floating-point type. The syntax is similar to the one used for integer data types — using the `range` keyword. This simple example creates a new floating-point type based on the type `Float`, for a normalized range between `-1.0` and `1.0`:

Listing 10: `floating_point_range.adb`

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Floating_Point_Range is
4   type T_Norm is new Float range -1.0 .. 1.0;
5   A : T_Norm;
6 begin
7   A := 1.0;
8   Put_Line ("The value of A is " & T_Norm'Image (A));
9 end Floating_Point_Range;
```

Build output

```

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

Runtime output

```
The value of A is 1.00000E+00
```

The application is responsible for ensuring that variables of this type stay within this range; otherwise an exception is raised. In this example, the exception `Constraint_Error` is raised when assigning `2.0` to the variable `A`:

Listing 11: `floating_point_range_exception.adb`

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Floating_Point_Range_Exception is
4   type T_Norm is new Float range -1.0 .. 1.0;
5   A : T_Norm;
6 begin
7   A := 2.0;
8   Put_Line ("The value of A is " & T_Norm'Image (A));
9 end Floating_Point_Range_Exception;
```

Build output

```

Compile
  [Ada]          floating_point_range_exception.adb
floating_point_range_exception.adb:7:09: warning: value not in range of type "T_
↳Norm" defined at line 4
floating_point_range_exception.adb:7:09: warning: "Constraint_Error" will be
↳raised at run time
```

(continues on next page)

(continued from previous page)

```

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

```

Runtime output

```
raised CONSTRAINT_ERROR : floating_point_range_exception.adb:7 range check failed
```

Ranges can also be specified for custom floating-point types. For example:

Listing 12: custom_range_types.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Numerics; use Ada.Numerics;
3
4 procedure Custom_Range_Types is
5   type T6_Inv_Trig is digits 6 range -Pi / 2.0 .. Pi / 2.0;
6 begin
7   null;
8 end Custom_Range_Types;

```

Build output

```

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

```

In this example, we are defining a type called `T6_Inv_Trig`, which has a range from $-\pi/2$ to $\pi/2$ with a minimum precision of 6 digits. (`Pi` is defined in the predefined package `Ada.Numerics`.)

5.6 Strong typing

As noted earlier, Ada is strongly typed. As a result, different types of the same family are incompatible with each other; a value of one type cannot be assigned to a variable from the other type. For example:

Listing 13: illegal_example.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Illegal_Example is
4   -- Declare two different floating point types
5   type Meters is new Float;
6   type Miles is new Float;
7
8   Dist_Imperial : Miles;
9
10  -- Declare a constant
11  Dist_Metric : constant Meters := 1000.0;
12 begin
13  -- Not correct: types mismatch

```

(continues on next page)

(continued from previous page)

```

14   Dist_Imperial := Dist_Metric * 621.371e-6;
15   Put_Line (Miles'Image (Dist_Imperial));
16 end Illegal_Example;

```

Build output

```

Compile
  [Ada]          illegal_example.adb
illegal_example.adb:14:33: expected type "Miles" defined at line 6
illegal_example.adb:14:33: found type "Meters" defined at line 5
gprbuild: *** compilation phase failed

```

A consequence of these rules is that, in the general case, a "mixed mode" expression like $2 * 3.0$ will trigger a compilation error. In a language like C or Python, such expressions are made valid by implicit conversions. In Ada, such conversions must be made explicit:

Listing 14: conv.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 procedure Conv is
3   type Meters is new Float;
4   type Miles is new Float;
5   Dist_Imperial : Miles;
6   Dist_Metric : constant Meters := 1000.0;
7 begin
8   Dist_Imperial := Miles (Dist_Metric) * 621.371e-6;
9   --           ^ Type conversion, from Meters to Miles
10  -- Now the code is correct
11
12   Put_Line (Miles'Image (Dist_Imperial));
13 end Conv;

```

Build output

```

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

```

Runtime output

```
6.21371E-01
```

Of course, we probably do not want to write the conversion code every time we convert from meters to miles. The idiomatic Ada way in that case would be to introduce conversion functions along with the types.

Listing 15: conv.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Conv is
4   type Meters is new Float;
5   type Miles is new Float;
6
7   -- Function declaration, like procedure but returns a value.
8   function To_Miles (M : Meters) return Miles is
9   --           ^ Return type

```

(continues on next page)

(continued from previous page)

```

10  begin
11      return Miles (M) * 621.371e-6;
12  end To_Miles;
13
14  Dist_Imperial : Miles;
15  Dist_Metric   : constant Meters := 1000.0;
16  begin
17      Dist_Imperial := To_Miles (Dist_Metric);
18      Put_Line (Miles'Image (Dist_Imperial));
19  end Conv;

```

Build output

```

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

```

Runtime output

```
6.21371E-01
```

If you write a lot of numeric code, having to explicitly provide such conversions might seem painful at first. However, this approach brings some advantages. Notably, you can rely on the absence of implicit conversions, which will in turn prevent some subtle errors.

In other languages

In C, for example, the rules for implicit conversions may not always be completely obvious. In Ada, however, the code will always do exactly what it seems to do. For example:

```

int a = 3, b = 2;
float f = a / b;

```

This code will compile fine, but the result of `f` will be 1.0 instead of 1.5, because the compiler will generate an integer division (three divided by two) that results in one. The software developer must be aware of data conversion issues and use an appropriate casting:

```

int a = 3, b = 2;
float f = (float)a / b;

```

In the corrected example, the compiler will convert both variables to their corresponding floating-point representation before performing the division. This will produce the expected result.

This example is very simple, and experienced C developers will probably notice and correct it before it creates bigger problems. However, in more complex applications where the type declaration is not always visible — e.g. when referring to elements of a `struct` — this situation might not always be evident and quickly lead to software defects that can be harder to find.

The Ada compiler, in contrast, will always reject code that mixes floating-point and integer variables without explicit conversion. The following Ada code, based on the erroneous example in C, will not compile:

Listing 16: main.adb

```

1  procedure Main is
2      A : Integer := 3;

```

(continues on next page)

(continued from previous page)

```

3   B : Integer := 2;
4   F : Float;
5   begin
6     F := A / B;
7   end Main;

```

Compilation output

```

main.adb:6:11: expected type "Standard.Float"
main.adb:6:11: found type "Standard.Integer"

```

The offending line must be changed to `F := Float (A) / Float (B);` in order to be accepted by the compiler.

- You can use Ada's strong typing to help enforce invariants in your code, as in the example above: Since Miles and Meters are two different types, you cannot mistakenly convert an instance of one to an instance of the other.

5.7 Derived types

In Ada you can create new types based on existing ones. This is very useful: you get a type that has the same properties as some existing type but is treated as a distinct type in the interest of strong typing.

Listing 17: main.adb

```

1  procedure Main is
2    -- ID card number type, incompatible with Integer.
3    type Social_Security_Number
4    is new Integer range 0 .. 999_99_9999;
5    --           ^ Since a SSN has 9 digits max, and cannot be
6    --           negative, we enforce a validity constraint.
7
8    SSN : Social_Security_Number := 555_55_5555;
9    --           ^ You can put underscores as formatting in
10   --           any number.
11
12   I   : Integer;
13
14   Invalid : Social_Security_Number := -1;
15   --           ^ This will cause a runtime error
16   --           (and a compile time warning with
17   --           GNAT)
18   begin
19     I := SSN;           -- Illegal, they have different types
20     SSN := I;          -- Likewise illegal
21     I := Integer (SSN); -- OK with explicit conversion
22     SSN := Social_Security_Number (I); -- Likewise OK
23   end Main;

```

Build output

```

Compile
  [Ada]          main.adb
main.adb:14:40: warning: value not in range of type "Social_Security_Number"
↳ defined at line 3
main.adb:14:40: warning: "Constraint_Error" will be raised at run time

```

(continues on next page)

(continued from previous page)

```
main.adb:19:09: expected type "Standard.Integer"
main.adb:19:09: found type "Social_Security_Number" defined at line 3
main.adb:20:11: expected type "Social_Security_Number" defined at line 3
main.adb:20:11: found type "Standard.Integer"
gprbuild: *** compilation phase failed
```

The type `Social_Security` is said to be a *derived type*; its *parent type* is `Integer`.

As illustrated in this example, you can refine the valid range when defining a derived scalar type (such as integer, floating-point and enumeration).

The syntax for enumerations uses the range `<range>` syntax:

Listing 18: greet.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet is
4   type Days is (Monday, Tuesday, Wednesday, Thursday,
5                Friday, Saturday, Sunday);
6
7   type Weekend_Days is new Days range Saturday .. Sunday;
8   -- New type, where only Saturday and Sunday are valid literals.
9 begin
10  null;
11 end Greet;
```

Build output

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

5.8 Subtypes

As we are starting to see, types may be used in Ada to enforce constraints on the valid range of values. However, we sometimes want to enforce constraints on some values while staying within a single type. This is where subtypes come into play. A subtype does not introduce a new type.

Listing 19: greet.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet is
4   type Days is (Monday, Tuesday, Wednesday, Thursday,
5                Friday, Saturday, Sunday);
6
7   -- Declaration of a subtype
8   subtype Weekend_Days is Days range Saturday .. Sunday;
9   -- ^ Constraint of the subtype
10
11  M : Days := Sunday;
12
13  S : Weekend_Days := M;
```

(continues on next page)

(continued from previous page)

```

14  -- No error here, Days and Weekend_Days are of the same type.
15  begin
16    for I in Days loop
17      case I is
18        -- Just like a type, a subtype can be used as a
19        -- range
20        when Weekend_Days =>
21          Put_Line ("Week end!");
22        when others =>
23          Put_Line ("Hello on " & Days'Image (I));
24      end case;
25    end loop;
26  end Greet;

```

Build output

```

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

```

Runtime output

```

Hello on MONDAY
Hello on TUESDAY
Hello on WEDNESDAY
Hello on THURSDAY
Hello on FRIDAY
Week end!
Week end!

```

Several subtypes are predefined in the standard package in Ada, and are automatically available to you:

```

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

```

While subtypes of a type are statically compatible with each other, constraints are enforced at run time: if you violate a subtype constraint, an exception will be raised.

Listing 20: greet.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Greet is
4    type Days is (Monday, Tuesday, Wednesday, Thursday,
5                 Friday, Saturday, Sunday);
6
7    subtype Weekend_Days is Days range Saturday .. Sunday;
8    Day : Days := Saturday;
9    Weekend : Weekend_Days;
10  begin
11    Weekend := Day;
12    --      ^ Correct: Same type, subtype constraints are respected
13    Weekend := Monday;
14    --      ^ Wrong value for the subtype
15    --      Compiles, but exception at runtime
16  end Greet;

```

Build output

```
Compile
  [Ada]          greet.adb
greet.adb:13:15: warning: value not in range of type "Weekend_Days" defined at
↳line 7
greet.adb:13:15: warning: "Constraint_Error" will be raised at run time
Bind
  [gprbind]     greet.bexch
  [Ada]         greet.ali
Link
  [link]        greet.adb
```

Runtime output

```
raised CONSTRAINT_ERROR : greet.adb:13 range check failed
```

5.8.1 Subtypes as type aliases

Previously, we've seen that we can create new types by declaring `type Miles is new Float`. We could also create type aliases, which generate alternative names — *aliases* — for known types. Note that type aliases are sometimes called *type synonyms*.

We achieve this in Ada by using subtypes without new constraints. In this case, however, we don't get all of the benefits of Ada's strong type checking. Let's rewrite an example using type aliases:

Listing 21: undetected_imperial_metric_error.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Undetected_Imperial_Metric_Error is
4   -- Declare two type aliases
5   subtype Meters is Float;
6   subtype Miles is Float;
7
8   Dist_Imperial : Miles;
9
10  -- Declare a constant
11  Dist_Metric : constant Meters := 100.0;
12 begin
13   -- No conversion to Miles type required:
14   Dist_Imperial := (Dist_Metric * 1609.0) / 1000.0;
15
16   -- Not correct, but undetected:
17   Dist_Imperial := Dist_Metric;
18
19   Put_Line (Miles'Image (Dist_Imperial));
20 end Undetected_Imperial_Metric_Error;
```

Build output

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

Runtime output

```
1.00000E+02
```

In the example above, the fact that both `Meters` and `Miles` are subtypes of `Float` allows us to mix variables of both types without type conversion. This, however, can lead to all sorts of programming mistakes that we'd like to avoid, as we can see in the undetected error highlighted in the code above. In that example, the error in the assignment of a value in meters to a variable meant to store values in miles remains undetected because both `Meters` and `Miles` are subtypes of `Float`. Therefore, the recommendation is to use strong typing — via type `X is new Y` — for cases such as the one above.

There are, however, many situations where type aliases are useful. For example, in an application that uses floating-point types in multiple contexts, we could use type aliases to indicate additional meaning to the types or to avoid long variable names. For example, instead of writing:

```
Paid_Amount, Due_Amount : Float;
```

We could write:

```
subtype Amount is Float;
Paid, Due : Amount;
```

In other languages

In C, for example, we can use a typedef declaration to create a type alias. For example:

```
typedef float meters;
```

This corresponds to the declaration that we've seen above using subtypes. Other programming languages include this concept in similar ways. For example:

- C++: `using meters = float;`
- Swift: `typealias Meters = Double`
- Kotlin: `typealias Meters = Double`
- Haskell: `type Meters = Float`

Note, however, that subtypes in Ada correspond to type aliases if, and only if, they don't have new constraints. Thus, if we add a new constraint to a subtype declaration, we don't have a type alias anymore. For example, the following declaration *can't* be consider a type alias of `Float`:

```
subtype Meters is Float range 0.0 .. 1_000_000.0;
```

Let's look at another example:

```
subtype Degree_Celsius is Float;
subtype Liquid_Water_Temperature is Degree_Celsius range 0.0 .. 100.0;
subtype Running_Water_Temperature is Liquid_Water_Temperature;
```

In this example, `Liquid_Water_Temperature` isn't an alias of `Degree_Celsius`, since it adds a new constraint that wasn't part of the declaration of the `Degree_Celsius`. However, we do have two type aliases here:

- `Degree_Celsius` is an alias of `Float`;
- `Running_Water_Temperature` is an alias of `Liquid_Water_Temperature`, even if `Liquid_Water_Temperature` itself has a constrained range.

RECORDS

So far, all the types we have encountered have values that are not decomposable: each instance represents a single piece of data. Now we are going to see our first class of composite types: records.

Records allow composing a value out of instances of other types. Each of those instances will be given a name. The pair consisting of a name and an instance of a specific type is called a field, or a component.

6.1 Record type declaration

Here is an example of a simple record declaration:

```
type Date is record
  -- The following declarations are components of the record
  Day   : Integer range 1 .. 31;
  Month : Months;
  Year  : Integer range 1 .. 3000; -- You can add custom constraints on fields
end record;
```

Fields look a lot like variable declarations, except that they are inside of a record definition. And as with variable declarations, you can specify additional constraints when supplying the subtype of the field.

```
type Date is record
  Day   : Integer range 1 .. 31;
  Month : Months := January;
  -- This component has a default value
  Year  : Integer range 1 .. 3000 := 2012;
  --                                     ^ Default value
end record;
```

Record components can have default values. When a variable having the record type is declared, a field with a default initialization will be automatically set to this value. The value can be any expression of the component type, and may be run-time computable.

6.2 Aggregates

```
Ada_Birthday      : Date := (10, December, 1815);
Leap_Day_2020    : Date := (Day => 29, Month => February, Year => 2020);
--                ^ By name
```

Records have a convenient notation for expressing values, illustrated above. This notation is called aggregate notation, and the literals are called aggregates. They can be used in a variety of contexts that we will see throughout the course, one of which is to initialize records.

An aggregate is a list of values separated by commas and enclosed in parentheses. It is allowed in any context where a value of the record is expected.

Values for the components can be specified positionally, as in `Ada_Birthday` example, or by name, as in `Leap_Day_2020`. A mixture of positional and named values is permitted, but you cannot use a positional notation after a named one.

6.3 Component selection

To access components of a record instance, you use an operation that is called component selection. This is achieved by using the dot notation. For example, if we declare a variable `Some_Day` of the `Date` record type mentioned above, we can access the `Year` component by writing `Some_Day.Year`.

Let's look at an example:

Listing 1: record_selection.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Record_Selection is
4
5     type Months is
6         (January, February, March, April, May, June, July,
7          August, September, October, November, December);
8
9     type Date is record
10        Day   : Integer range 1 .. 31;
11        Month : Months;
12        Year  : Integer range 1 .. 3000 := 2032;
13    end record;
14
15    procedure Display_Date (D : Date) is
16    begin
17        Put_Line ("Day:" & Integer'Image (D.Day)
18                & ", Month: " & Months'Image (D.Month)
19                & ", Year:" & Integer'Image (D.Year));
20    end Display_Date;
21
22    Some_Day : Date := (1, January, 2000);
23
24 begin
25     Display_Date (Some_Day);
26
27     Put_Line ("Changing year...");
28     Some_Day.Year := 2001;
29
30     Display_Date (Some_Day);
31 end Record_Selection;
```

Build output

```

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

```

Runtime output

```

Day: 1, Month: JANUARY, Year: 2000
Changing year...
Day: 1, Month: JANUARY, Year: 2001

```

As you can see in this example, we can use the dot notation in the expression `D.Year` or `Some_Day.Year` to access the information stored in that component, as well as to modify this information in assignments. To be more specific, when we use `D.Year` in the call to `Put_Line`, we're retrieving the information stored in that component. When we write `Some_Day.Year := 2001`, we're overwriting the information that was previously stored in the `Year` component of `Some_Day`.

6.4 Renaming

In previous chapters, we've discussed *subprogram* (page 26) and *package* (page 39) renaming. We can rename record components as well. Instead of writing the full component selection using the dot notation, we can declare an alias that allows us to access the same component. This is useful to simplify the implementation of a subprogram, for example.

We can rename record components by using the `renames` keyword in a variable declaration. For example:

```

Some_Day : Date
Y        : Integer renames Some_Day.Year;

```

Here, `Y` is an alias, so that every time we using `Y`, we are really using the `Year` component of `Some_Day`.

Let's look at a complete example:

Listing 2: dates.ads

```

1 package Dates is
2
3   type Months is
4     (January, February, March, April, May, June, July,
5      August, September, October, November, December);
6
7   type Date is record
8     Day   : Integer range 1 .. 31;
9     Month : Months;
10    Year  : Integer range 1 .. 3000 := 2032;
11  end record;
12
13  procedure Increase_Month (Some_Day : in out Date);
14
15  procedure Display_Month (Some_Day : Date);
16
17 end Dates;

```

Listing 3: dates.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body Dates is
4
5     procedure Increase_Month (Some_Day : in out Date) is
6         -- Renaming components from the Date record
7         M : Months renames Some_Day.Month;
8         Y : Integer renames Some_Day.Year;
9
10        -- Renaming function (for Months enumeration)
11        function Next (M : Months) return Months
12            renames Months'Succ;
13    begin
14        if M = December then
15            M := January;
16            Y := Y + 1;
17        else
18            M := Next (M);
19        end if;
20    end Increase_Month;
21
22    procedure Display_Month (Some_Day : Date) is
23        -- Renaming components from the Date record
24        M : Months renames Some_Day.Month;
25        Y : Integer renames Some_Day.Year;
26    begin
27        Put_Line ("Month: " & Months'Image (M)
28                & ", Year:" & Integer'Image (Y));
29    end Display_Month;
30
31 end Dates;
```

Listing 4: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Dates;      use Dates;
3
4 procedure Main is
5     D : Date := (1, January, 2000);
6 begin
7     Display_Month (D);
8
9     Put_Line ("Increasing month...");
10    Increase_Month (D);
11
12    Display_Month (D);
13 end Main;
```

Build output

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

Runtime output

```
Month: JANUARY, Year: 2000
Increasing month...
Month: FEBRUARY, Year: 2000
```

We apply renaming to two components of the `Date` record in the implementation of the `Increase_Month` procedure. Then, instead of directly using `Some_Day.Month` and `Some_Day.Year` in the next operations, we simply use the renamed versions `M` and `Y`.

Note that, in the example above, we also rename `Months'Succ` — which is the function that gives us the next month — to `Next`.

ARRAYS

Arrays provide another fundamental family of composite types in Ada.

7.1 Array type declaration

Arrays in Ada are used to define contiguous collections of elements that can be selected by indexing. Here's a simple example:

Listing 1: greet.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet is
4   type My_Int is range 0 .. 1000;
5   type Index is range 1 .. 5;
6
7   type My_Int_Array is array (Index) of My_Int;
8     --                               ^ Type of elements
9     --                               ^ Bounds of the array
10  Arr : My_Int_Array := (2, 3, 5, 7, 11);
11     --                               ^ Array literal, called aggregate in Ada
12 begin
13   for I in Index loop
14     Put (My_Int'Image (Arr (I)));
15     --                               ^ Take the Ith element
16   end loop;
17   New_Line;
18 end Greet;
```

Build output

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

Runtime output

```
2 3 5 7 11
```

The first point to note is that we specify the index type for the array, rather than its size. Here we declared an integer type named `Index` ranging from 1 to 5, so each array instance will have 5 elements, with the initial element at index 1 and the last element at index 5.

Although this example used an integer type for the index, Ada is more general: any discrete type is permitted to index an array, including *Enum types* (page 45). We will soon see what that means.

Another point to note is that querying an element of the array at a given index uses the same syntax as for function calls: that is, the array object followed by the index in parentheses.

Thus when you see an expression such as `A (B)`, whether it is a function call or an array subscript depends on what `A` refers to.

Finally, notice how we initialize the array with the `(2, 3, 5, 7, 11)` expression. This is another kind of aggregate in Ada, and is in a sense a literal expression for an array, in the same way that `3` is a literal expression for an integer. The notation is very powerful, with a number of properties that we will introduce later. A detailed overview appears in the notation of *aggregate types* (page 83).

Unrelated to arrays, the example also illustrated two procedures from `Ada.Text_IO`:

- `Put`, which displays a string without a terminating end of line
- `New_Line`, which outputs an end of line

Let's now delve into what it means to be able to use any discrete type to index into the array.

In other languages

Semantically, an array object in Ada is the entire data structure, and not simply a handle or pointer. Unlike C and C++, there is no implicit equivalence between an array and a pointer to its initial element.

Listing 2: `array_bounds_example.adb`

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Array_Bounds_Example is
4   type My_Int is range 0 .. 1000;
5   type Index is range 11 .. 15;
6   --           ^ Low bound can be any value
7   type My_Int_Array is array (Index) of My_Int;
8   Tab : My_Int_Array := (2, 3, 5, 7, 11);
9 begin
10  for I in Index loop
11    Put (My_Int'Image (Tab (I)));
12  end loop;
13  New_Line;
14 end Array_Bounds_Example;
```

Build output

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

Runtime output

```
2 3 5 7 11
```

One effect is that the bounds of an array can be any values. In the first example we constructed an array type whose first index is 1, but in the example above we declare an array type whose first index is 11.

That's perfectly fine in Ada, and moreover since we use the index type as a range to iterate over the array indices, the code using the array does not need to change.

That leads us to an important consequence with regard to code dealing with arrays. Since the bounds can vary, you should not assume / hard-code specific bounds when iterating / using arrays. That means the code above is good, because it uses the index type, but a for loop as shown below is bad practice even though it works correctly:

```
for I in 11 .. 15 loop
  Tab (I) := Tab (I) * 2;
end loop;
```

Since you can use any discrete type to index an array, enumeration types are permitted.

Listing 3: month_example.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Month_Example is
4   type Month_Duration is range 1 .. 31;
5   type Month is (Jan, Feb, Mar, Apr, May, Jun,
6                 Jul, Aug, Sep, Oct, Nov, Dec);
7
8   type My_Int_Array is array (Month) of Month_Duration;
9   --           ^ Can use an enumeration type as the
10  --           index
11
12  Tab : constant My_Int_Array :=
13  --   ^ constant is like a variable but cannot be
14  --   modified
15    (31, 28, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31);
16  -- Maps months to number of days (ignoring leap years)
17
18  Feb_Days : Month_Duration := Tab (Feb);
19  -- Number of days in February
20 begin
21   for M in Month loop
22     Put_Line
23       (Month'Image (M) & " has "
24        & Month_Duration'Image (Tab (M)) & " days.");
25     --           ^ Concatenation operator
26   end loop;
27 end Month_Example;
```

Build output

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

Runtime output

```
JAN has 31 days.
FEB has 28 days.
MAR has 31 days.
APR has 30 days.
MAY has 31 days.
JUN has 30 days.
```

(continues on next page)

(continued from previous page)

```
JUL has 31 days.
AUG has 31 days.
SEP has 30 days.
OCT has 31 days.
NOV has 30 days.
DEC has 31 days.
```

In the example above, we are:

- Creating an array type mapping months to month durations in days.
- Creating an array, and instantiating it with an aggregate mapping months to their actual durations in days.
- Iterating over the array, printing out the months, and the number of days for each.

Being able to use enumeration values as indices is very helpful in creating mappings such as shown above one, and is an often used feature in Ada.

7.2 Indexing

We have already seen the syntax for selecting elements of an array. There are however a few more points to note.

First, as is true in general in Ada, the indexing operation is strongly typed. If you use a value of the wrong type to index the array, you will get a compile-time error.

Listing 4: greet.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet is
4   type My_Int is range 0 .. 1000;
5
6   type My_Index is range 1 .. 5;
7   type Your_Index is range 1 .. 5;
8
9   type My_Int_Array is array (My_Index) of My_Int;
10  Tab : My_Int_Array := (2, 3, 5, 7, 11);
11 begin
12   for I in Your_Index loop
13     Put (My_Int'Image (Tab (I)));
14     -- ^ Compile time error
15   end loop;
16   New_Line;
17 end Greet;
```

Build output

```
Compile
  [Ada]      greet.adb
greet.adb:13:31: expected type "My_Index" defined at line 6
greet.adb:13:31: found type "Your_Index" defined at line 7
gprbuild: *** compilation phase failed
```

Second, arrays in Ada are bounds checked. This means that if you try to access an element outside of the bounds of the array, you will get a run-time error instead of accessing random memory as in unsafe languages.

Listing 5: greet.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet is
4   type My_Int is range 0 .. 1000;
5   type Index is range 1 .. 5;
6   type My_Int_Array is array (Index) of My_Int;
7   Tab : My_Int_Array := (2, 3, 5, 7, 11);
8 begin
9   for I in Index range 2 .. 6 loop
10    Put (My_Int'Image (Tab (I)));
11    --           ^ Will raise an exception when
12    --           I = 6
13  end loop;
14  New_Line;
15 end Greet;

```

Build output

```

Compile
  [Ada]          greet.adb
greet.adb:9:30: warning: static value out of range of type "Index" defined at line_
↪5
greet.adb:9:30: warning: "Constraint_Error" will be raised at run time
greet.adb:9:30: warning: suspicious loop bound out of range of loop subtype
greet.adb:9:30: warning: loop executes zero times or raises Constraint_Error
Bind
  [gprbind]     greet.bexch
  [Ada]         greet.ali
Link
  [link]       greet.adb

```

Runtime output

```

raised CONSTRAINT_ERROR : greet.adb:9 range check failed

```

7.3 Simpler array declarations

In the previous examples, we have always explicitly created an index type for the array. While this can be useful for typing and readability purposes, sometimes you simply want to express a range of values. Ada allows you to do that, too.

Listing 6: simple_array_bounds.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Simple_Array_Bounds is
4   type My_Int is range 0 .. 1000;
5   type My_Int_Array is array (1 .. 5) of My_Int;
6   --           ^ Subtype of Integer
7   Tab : My_Int_Array := (2, 3, 5, 7, 11);
8 begin
9   for I in 1 .. 5 loop
10    --           ^ Likewise
11    Put (My_Int'Image (Tab (I)));
12  end loop;

```

(continues on next page)

(continued from previous page)

```

13   New_Line;
14 end Simple_Array_Bounds;

```

Build output

```

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

```

Runtime output

```
2 3 5 7 11
```

This example defines the range of the array via the range syntax, which specifies an anonymous subtype of Integer and uses it to index the array.

This means that the type of the index is Integer. Similarly, when you use an anonymous range in a for loop as in the example above, the type of the iteration variable is also Integer, so you can use I to index Tab.

You can also use a named subtype for the bounds for an array.

7.4 Range attribute

We noted earlier that hard coding bounds when iterating over an array is a bad idea, and showed how to use the array's index type/subtype to iterate over its range in a for loop. That raises the question of how to write an iteration when the array has an anonymous range for its bounds, since there is no name to refer to the range. Ada solves that via several attributes of array objects:

Listing 7: range_example.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Range_Example is
4   type My_Int is range 0 .. 1000;
5   type My_Int_Array is array (1 .. 5) of My_Int;
6   Tab : My_Int_Array := (2, 3, 5, 7, 11);
7 begin
8   for I in Tab'Range loop
9     --      ^ Gets the range of Tab
10    Put (My_Int'Image (Tab (I)));
11  end loop;
12  New_Line;
13 end Range_Example;

```

Build output

```

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

```

Runtime output

```
2 3 5 7 11
```

If you want more fine grained control, you can use the separate attributes 'First and 'Last.

Listing 8: array_attributes_example.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Array_Attributes_Example is
4   type My_Int is range 0 .. 1000;
5   type My_Int_Array is array (1 .. 5) of My_Int;
6   Tab : My_Int_Array := (2, 3, 5, 7, 11);
7 begin
8   for I in Tab'First .. Tab'Last - 1 loop
9     -- ^ Iterate on every index except the last
10    Put (My_Int'Image (Tab (I)));
11  end loop;
12  New_Line;
13 end Array_Attributes_Example;
```

Build output

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

Runtime output

```
2 3 5 7
```

The 'Range, 'First and 'Last attributes in these examples could also have been applied to the array type name, and not just the array instances.

Although not illustrated in the above examples, another useful attribute for an array instance A is A'Length, which is the number of elements that A contains.

It is legal and sometimes useful to have a "null array", which contains no elements. To get this effect, define an index range whose upper bound is less than the lower bound.

7.5 Unconstrained arrays

Let's now consider one of the most powerful aspects of Ada's array facility.

Every array type we have defined so far has a fixed size: every instance of this type will have the same bounds and therefore the same number of elements and the same size.

However, Ada also allows you to declare array types whose bounds are not fixed: in that case, the bounds will need to be provided when creating instances of the type.

Listing 9: unconstrained_array_example.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Unconstrained_Array_Example is
4   type Days is (Monday, Tuesday, Wednesday,
```

(continues on next page)

(continued from previous page)

```

5           Thursday, Friday, Saturday, Sunday);
6
7     type Workload_Type is array (Days range <>) of Natural;
8     -- Indefinite array type
9     --           ^ Bounds are of type Days,
10    --           but not known
11
12    Workload : constant Workload_Type (Monday .. Friday) :=
13    --           ^ Specify the bounds
14    --           when declaring
15    (Friday => 7, others => 8);
16    --           ^ Default value
17    -- ^ Specify element by name of index
18  begin
19    for I in Workload'Range loop
20      Put_Line (Integer'Image (Workload (I)));
21    end loop;
22  end Unconstrained_Array_Example;

```

Build output

```

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

```

Runtime output

```

8
8
8
8
7

```

The fact that the bounds of the array are not known is indicated by the `Days range <>` syntax. Given a discrete type `Discrete_Type`, if we use `Discrete_Type` for the index in an array type then `Discrete_Type` serves as the type of the index and comprises the range of index values for each array instance.

If we define the index as `Discrete_Type range <>` then `Discrete_Type` serves as the type of the index, but different array instances may have different bounds from this type

An array type that is defined with the `Discrete_Type range <>` syntax for its index is referred to as an unconstrained array type, and, as illustrated above, the bounds need to be provided when an instance is created.

The above example also shows other forms of the aggregate syntax. You can specify associations by name, by giving the value of the index on the left side of an arrow association. `1 => 2` thus means "assign value 2 to the element at index 1 in my array". `others => 8` means "assign value 8 to every element that wasn't previously assigned in this aggregate".

Attention: The so-called "box" notation (`<>`) is commonly used as a wildcard or placeholder in Ada. You will often see it when the meaning is "what is expected here can be anything".

In other languages

While unconstrained arrays in Ada might seem similar to variable length arrays in C, they are in reality much more powerful, because they're truly first-class values in the language. You can pass them as parameters to subprograms or return them from functions, and they implicitly contain their bounds as part of their value. This means that it is useless to pass the bounds or length of an array explicitly along with the array, because they are accessible via the 'First, 'Last, 'Range and 'Length attributes explained earlier.

Although different instances of the same unconstrained array type can have different bounds, a specific instance has the same bounds throughout its lifetime. This allows Ada to implement unbounded arrays efficiently; instances can be stored on the stack and do not require heap allocation as in languages like Java.

7.6 Predefined array type: String

A recurring theme in our introduction to Ada types has been the way important built-in types like Boolean or Integer are defined through the same facilities that are available to the user. This is also true for strings: The String type in Ada is a simple array.

Here is how the string type is defined in Ada:

```
type String is array (Positive range <>) of Character;
```

The only built-in feature Ada adds to make strings more ergonomic is custom literals, as we can see in the example below.

Hint: String literals are a syntactic sugar for aggregates, so that in the following example, A and B have the same value.

Listing 10: string_literals.ads

```
1 package String_Literals is
2   -- Those two declarations are equivalent
3   A : String (1 .. 11) := "Hello World";
4   B : String (1 .. 11) := ('H', 'e', 'l', 'l', 'o', ' ',
5                           'W', 'o', 'r', 'l', 'd');
6 end String_Literals;
```

Listing 11: greet.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet is
4   Message : String (1 .. 11) := "dlroW olleH";
5   --     ^ Pre-defined array type.
6   --     Component type is Character
7 begin
8   for I in reverse Message'Range loop
9     --     ^ Iterate in reverse order
10    Put (Message (I));
11  end loop;
12  New_Line;
13 end Greet;
```

However, specifying the bounds of the object explicitly is a bit of a hassle; you have to manually count the number of characters in the literal. Fortunately, Ada gives you an easier way.

You can omit the bounds when creating an instance of an unconstrained array type if you supply an initialization, since the bounds can be deduced from the initialization expression.

Listing 12: greet.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Greet is
4   Message : constant String := "dlroW olleH";
5   --           ^ Bounds are automatically computed
6   --           from initialization value
7 begin
8   for I in reverse Message'Range loop
9     Put (Message (I));
10  end loop;
11  New_Line;
12 end Greet;
```

Build output

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

Runtime output

```
Hello World
```

Listing 13: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   type Integer_Array is array (Natural range <>) of Integer;
5
6   My_Array : constant Integer_Array := (1, 2, 3, 4);
7   --           ^ Bounds are automatically computed
8   --           from initialization value
9 begin
10  null;
11 end Main;
```

Attention: As you can see above, the standard String type in Ada is an array. As such, it shares the advantages and drawbacks of arrays: a String value is stack allocated, it is accessed efficiently, and its bounds are immutable.

If you want something akin to C++'s `std::string`, you can use *Unbounded Strings* (page 230) from Ada's standard library. This type is more like a mutable, automatically managed string buffer to which you can add content.

7.7 Restrictions

A very important point about arrays: bounds *have* to be known when instances are created. It is for example illegal to do the following.

```
declare
  A : String;
begin
  A := "World";
end;
```

Also, while you of course can change the values of elements in an array, you cannot change the array's bounds (and therefore its size) after it has been initialized. So this is also illegal:

```
declare
  A : String := "Hello";
begin
  A := "World";           -- OK: Same size
  A := "Hello World";    -- Not OK: Different size
end;
```

Also, while you can expect a warning for this kind of error in very simple cases like this one, it is impossible for a compiler to know in the general case if you are assigning a value of the correct length, so this violation will generally result in a run-time error.

Attention

While we will learn more about this later, it is important to know that arrays are not the only types whose instances might be of unknown size at compile-time.

Such objects are said to be of an *indefinite subtype*, which means that the subtype size is not known at compile time, but is dynamically computed (at run time).

Listing 14: indefinite_subtypes.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Indefinite_Subtypes is
4   function Get_Number return Integer is
5     begin
6       return Integer'Value (Get_Line);
7     end Get_Number;
8
9   A : String := "Hello";
10  -- Indefinite subtype
11
12  B : String (1 .. 5) := "Hello";
13  -- Definite subtype
14
15  C : String (1 .. Get_Number);
16  -- Indefinite subtype (Get_Number's value is computed at run-time)
17 begin
18   null;
19 end Indefinite_Subtypes;
```

7.8 Returning unconstrained arrays

The return type of a function can be any type; a function can return a value whose size is unknown at compile time. Likewise, the parameters can be of any type.

For example, this is a function that returns an unconstrained String:

Listing 15: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4
5     type Days is (Monday, Tuesday, Wednesday,
6                  Thursday, Friday, Saturday, Sunday);
7
8     function Get_Day_Name (Day : Days := Monday) return String is
9     begin
10        return
11        (case Day is
12         when Monday => "Monday",
13         when Tuesday => "Tuesday",
14         when Wednesday => "Wednesday",
15         when Thursday => "Thursday",
16         when Friday => "Friday",
17         when Saturday => "Saturday",
18         when Sunday => "Sunday");
19     end Get_Day_Name;
20
21 begin
22     Put_Line ("First day is " & Get_Day_Name (Days'First));
23 end Main;
```

Build output

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

Runtime output

```
First day is Monday
```

(This example is for illustrative purposes only. There is a built-in mechanism, the 'Image attribute for scalar types, that returns the name (as a String) of any element of an enumeration type. For example Days'Image(Monday) is "MONDAY".)

In other languages

Returning variable size objects in languages lacking a garbage collector is a bit complicated implementation-wise, which is why C and C++ don't allow it, preferring to depend on explicit dynamic allocation / free from the user.

The problem is that explicit storage management is unsafe as soon as you want to collect unused memory. Ada's ability to return variable size objects will remove one use case for dynamic allocation, and hence, remove one potential source of bugs from your programs.

Rust follows the C/C++ model, but with safe pointer semantics. However, dynamic allocation is still used. Ada can benefit from an eventual performance edge because it can use any model.

7.9 Declaring arrays (2)

While we can have array types whose size and bounds are determined at run time, the array's component type needs to be of a definite and constrained type.

Thus, if you need to declare, for example, an array of Strings, the String subtype used as component will need to have a fixed size.

Listing 16: show_days.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Days is
4   type Days is (Monday, Tuesday, Wednesday,
5                Thursday, Friday, Saturday, Sunday);
6
7   subtype Day_Name is String (1 .. 2);
8   -- Subtype of string with known size
9
10  type Days_Name_Type
11  is array (Days) of Day_Name;
12  --      ^ Type of the index
13  --      ^ Type of the element. Must be
14  --      definite
15
16  Names : constant Days_Name_Type :=
17    ("Mo", "Tu", "We", "Th", "Fr", "Sa", "Su");
18  -- Initial value given by aggregate
19 begin
20   for I in Names'Range loop
21     Put_Line (Names (I));
22   end loop;
23 end Show_Days;
```

Build output

```

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

Runtime output

```

Mo
Tu
We
Th
Fr
Sa
Su
```

7.10 Array slices

One last feature of Ada arrays that we're going to cover is array slices. It is possible to take and use a slice of an array (a contiguous sequence of elements) as a name or a value.

Listing 17: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Main is
4   Buf : String := "Hello ...";
5
6   Full_Name : String := "John Smith";
7 begin
8   Buf (7 .. 9) := "Bob";
9   -- Careful! This works because the string on the right side is the
10  -- same length as the replaced slice!
11
12  Put_Line (Buf); -- Prints "Hello Bob"
13
14  Put_Line ("Hi " & Full_Name (1 .. 4)); -- Prints "Hi John"
15 end Main;

```

Build output

```

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

```

Runtime output

```

Hello Bob
Hi John

```

As we can see above, you can use a slice on the left side of an assignment, to replace only part of an array.

A slice of an array is of the same type as the array, but has a different subtype, constrained by the bounds of the slice.

Attention: Ada has [multidimensional arrays](http://www.adaic.org/resources/add_content/standards/12rm/html/RM-3-6.html)¹⁰, which are not covered in this course. Slices will only work on one dimensional arrays.

¹⁰ http://www.adaic.org/resources/add_content/standards/12rm/html/RM-3-6.html

7.11 Renaming

So far, we've seen that the following elements can be renamed: *subprograms* (page 26), *packages* (page 39), and *record components* (page 61). We can also rename objects by using the `renames` keyword. This allows for creating alternative names for these objects. Let's look at an example:

Listing 18: measurements.ads

```

1 package Measurements is
2     subtype Degree_Celsius is Float;
3
4     Current_Temperature : Degree_Celsius;
5
6
7 end Measurements;
```

Listing 19: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Measurements;
3
4 procedure Main is
5     subtype Degrees is Measurements.Degree_Celsius;
6
7     T : Degrees renames Measurements.Current_Temperature;
8 begin
9     T := 5.0;
10
11     Put_Line (Degrees'Image (T));
12     Put_Line (Degrees'Image (Measurements.Current_Temperature));
13
14     T := T + 2.5;
15
16     Put_Line (Degrees'Image (T));
17     Put_Line (Degrees'Image (Measurements.Current_Temperature));
18 end Main;
```

Build output

```

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

Runtime output

```

5.00000E+00
5.00000E+00
7.50000E+00
7.50000E+00
```

In the example above, we declare a variable `T` by renaming the `Current_Temperature` object from the `Measurements` package. As you can see by running this example, both `Current_Temperature` and its alternative name `T` have the same values:

- first, they show the value 5.0
- after the addition, they show the value 7.5.

This is because they are essentially referring to the same object, but with two different names.

Note that, in the example above, we're using `Degrees` as an alias of `Degree_Celsius`. We discussed this method *earlier in the course* (page 56).

Renaming can be useful for improving the readability of more complicated array indexing. Instead of explicitly using indices every time we're accessing certain positions of the array, we can create shorter names for these positions by renaming them. Let's look at the following example:

Listing 20: colors.ads

```
1 package Colors is
2
3     type Color is (Black, Red, Green, Blue, White);
4
5     type Color_Array is array (Positive range <>) of Color;
6
7     procedure Reverse_It (X : in out Color_Array);
8
9 end Colors;
```

Listing 21: colors.adb

```
1 package body Colors is
2
3     procedure Reverse_It (X : in out Color_Array) is
4     begin
5         for I in X'First .. (X'Last + X'First) / 2 loop
6             declare
7                 Tmp      : Color;
8                 X_Left   : Color renames X (I);
9                 X_Right  : Color renames X (X'Last + X'First - I);
10            begin
11                Tmp      := X_Left;
12                X_Left   := X_Right;
13                X_Right  := Tmp;
14            end;
15        end loop;
16    end Reverse_It;
17
18 end Colors;
```

Listing 22: test_reverse_colors.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 with Colors; use Colors;
4
5 procedure Test_Reverse_Colors is
6
7     My_Colors : Color_Array (1 .. 5) := (Black, Red, Green, Blue, White);
8
9 begin
10     for C of My_Colors loop
11         Put_Line ("My_Color: " & Color'Image (C));
12     end loop;
13
14     New_Line;
15     Put_Line ("Reversing My_Color...");
16     New_Line;
17     Reverse_It (My_Colors);
18
```

(continues on next page)

(continued from previous page)

```

19   for C of My_Colors loop
20       Put_Line ("My_Color: " & Color'Image (C));
21   end loop;
22
23 end Test_Reverse_Colors;

```

Build output

```

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

```

Runtime output

```

My_Color: BLACK
My_Color: RED
My_Color: GREEN
My_Color: BLUE
My_Color: WHITE

Reversing My_Color...

My_Color: WHITE
My_Color: BLUE
My_Color: GREEN
My_Color: RED
My_Color: BLACK

```

In the example above, package `Colors` implements the procedure `Reverse_It` by declaring new names for two positions of the array. The actual implementation becomes easy to read:

```

begin
  Tmp      := X_Left;
  X_Left   := X_Right;
  X_Right  := Tmp;
end;

```

Compare this to the alternative version without renaming:

```

begin
  Tmp      := X (I);
  X (I)    := X (X'Last + X'First - I);
  X (X'Last + X'First - I) := Tmp;
end;

```


MORE ABOUT TYPES

8.1 Aggregates: A primer

So far, we have talked about aggregates quite a bit and have seen a number of examples. Now we will revisit this feature in some more detail.

An Ada aggregate is, in effect, a literal value for a composite type. It's a very powerful notation that helps you to avoid writing procedural code for the initialization of your data structures in many cases.

A basic rule when writing aggregates is that *every component* of the array or record has to be specified, even components that have a default value.

This means that the following code is incorrect:

Listing 1: incorrect.ads

```
1 package Incorrect is
2   type Point is record
3     X, Y : Integer := 0;
4   end record;
5
6   Origin : Point := (X => 0);
7 end Incorrect;
```

Compilation output

```
incorrect.ads:6:22: no value supplied for component "Y"
```

There are a few shortcuts that you can use to make the notation more convenient:

- To specify the default value for a component, you can use the `<>` notation.
- You can use the `|` symbol to give several components the same value.
- You can use the `others` choice to refer to every component that has not yet been specified, provided all those fields have the same type.
- You can use the range notation `..` to refer to specify a contiguous sequence of indices in an array.

However, note that as soon as you used a named association, all subsequent components likewise need to be specified with names associations.

Listing 2: points.ads

```
1 package Points is
2   type Point is record
3     X, Y : Integer := 0;
4   end record;
```

(continues on next page)

(continued from previous page)

```

5
6  type Point_Array is array (Positive range <>) of Point;
7
8  Origin   : Point := (X | Y => <>);  -- use the default values
9  Origin_2 : Point := (others => <>);  -- likewise use the defaults
10
11  Points_1 : Point_Array := ((1, 2), (3, 4));
12  Points_2 : Point_Array := (1 => (1, 2), 2 => (3, 4), 3 .. 20 => <>);
13 end Points;

```

8.2 Overloading and qualified expressions

Ada has a general concept of name overloading, which we saw earlier in the section on *enumeration types* (page 45).

Let's take a simple example: it is possible in Ada to have functions that have the same name, but different types for their parameters.

Listing 3: pkg.ads

```

1 package Pkg is
2   function F (A : Integer) return Integer;
3   function F (A : Character) return Integer;
4 end Pkg;

```

This is a common concept in programming languages, called *overloading*¹¹, or name overloading.

One of the novel aspects of Ada's overloading facility is the ability to resolve overloading based on the return type of a function.

Listing 4: pkg.ads

```

1 package Pkg is
2   type SSID is new Integer;
3
4   function Convert (Self : SSID) return Integer;
5   function Convert (Self : SSID) return String;
6 end Pkg;

```

Listing 5: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Pkg;         use Pkg;
3
4 procedure Main is
5   S : String := Convert (123_145_299);
6   --           ^ Valid, will choose the proper Convert
7 begin
8   Put_Line (S);
9 end Main;

```

Attention: Note that overload resolution based on the type is allowed for both functions and enumeration literals in Ada - which is why you can have multiple enumeration literals with the same name. Semantically, an enumeration literal is treated like a function that has no parameters.

¹¹ https://en.m.wikipedia.org/wiki/Function_overloading

However, sometimes an ambiguity makes it impossible to resolve which declaration of an overloaded name a given occurrence of the name refers to. This is where a qualified expression becomes useful.

Listing 6: pkg.ads

```

1 package Pkg is
2   type SSID is new Integer;
3
4   function Convert (Self : SSID) return Integer;
5   function Convert (Self : SSID) return String;
6   function Convert (Self : Integer) return String;
7 end Pkg;
```

Listing 7: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Pkg;         use Pkg;
3
4 procedure Main is
5   S : String := Convert (123_145_299);
6   --           ^ Invalid, which convert should we call?
7
8   S2 : String := Convert (SSID'(123_145_299));
9   --           ^ We specify that the type of the expression is
10  --             SSID.
11
12  -- We could also have declared a temporary
13
14  I : SSID := 123_145_299;
15
16  S3 : String := Convert (I);
17 begin
18   Put_Line (S);
19 end Main;
```

Compilation output

```

main.adb:5:18: ambiguous expression (cannot resolve "Convert")
main.adb:5:18: possible interpretation at pkg.ads:6
main.adb:5:18: possible interpretation at pkg.ads:5
```

Syntactically the target of a qualified expression can be either any expression in parentheses, or an aggregate:

Listing 8: qual_expr.ads

```

1 package Qual_Expr is
2   type Point is record
3     A, B : Integer;
4   end record;
5
6   P : Point := Point'(12, 15);
7
8   A : Integer := Integer'(12);
9 end Qual_Expr;
```

This illustrates that qualified expressions are a convenient (and sometimes necessary) way for the programmer to make the type of an expression explicit, for the compiler of course, but also for other programmers.

Attention: While they look and feel similar, type conversions and qualified expressions are *not* the same.

A qualified expression specifies the exact type that the target expression will be resolved to, whereas a type conversion will try to convert the target and issue a run-time error if the target value cannot be so converted.

Note that you can use a qualified expression to convert from one subtype to another, with an exception raised if a constraint is violated.

```
X : Integer := Natural'(1);
```

8.3 Access types (pointers)

Pointers are a potentially dangerous construct, which conflicts with Ada's underlying philosophy.

There are two ways in which Ada helps shield programmers from the dangers of pointers:

1. One approach, which we have already seen, is to provide alternative features so that the programmer does not need to use pointers. Parameter modes, arrays, and varying size types are all constructs that can replace typical pointer usages in C.
2. Second, Ada has made pointers as safe and restricted as possible, but allows "escape hatches" when the programmer explicitly requests them and presumably will be exercising such features with appropriate care.

Here is how you declare a simple pointer type, or access type, in Ada:

Listing 9: dates.ads

```
1 package Dates is
2   type Months is (January, February, March, April, May, June, July,
3     August, September, October, November, December);
4
5   type Date is record
6     Day   : Integer range 1 .. 31;
7     Month : Months;
8     Year  : Integer;
9   end record;
10 end Dates;
```

Listing 10: access_types.ads

```
1 with Dates; use Dates;
2
3 package Access_Types is
4   -- Declare an access type
5   type Date_Acc is access Date;
6   --           ^ "Designated type"
7   --           ^ Date_Acc values point to Date objects
8
9   D : Date_Acc := null;
10  --           ^ Literal for "access to nothing"
11  --           ^ Access to date
12 end Access_Types;
```

This illustrates how to:

- Declare an access type whose values point to ("designate") objects from a specific type
- Declare a variable (access value) from this access type

- Give it a value of null

In line with Ada's strong typing philosophy, if you declare a second access type whose designated type is Date, the two access types will be incompatible with each other, and you will need an explicit type conversion to convert from one to the other:

Listing 11: access_types.ads

```

1 with Dates; use Dates;
2
3 package Access_Types is
4   -- Declare an access type
5   type Date_Acc is access Date;
6   type Date_Acc_2 is access Date;
7
8   D : Date_Acc := null;
9   D2 : Date_Acc_2 := D;
10  --           ^ Invalid! Different types
11
12  D3 : Date_Acc_2 := Date_Acc_2 (D);
13  --           ^ Valid with type conversion
14 end Access_Types;
```

Compilation output

```

access_types.ads:9:24: expected type "Date_Acc_2" defined at line 6
access_types.ads:9:24: found type "Date_Acc" defined at line 5
access_types.ads:12:24: target type must be general access type
access_types.ads:12:24: add "all" to type "Date_Acc_2" defined at line 6
```

In other languages

In most other languages, pointer types are structurally, not nominally typed, like they are in Ada, which means that two pointer types will be the same as long as they share the same target type and accessibility rules.

Not so in Ada, which takes some time getting used to. A seemingly simple problem is, if you want to have a canonical access to a type, where should it be declared? A commonly used pattern is that if you need an access type to a specific type you "own", you will declare it along with the type:

```

package Access_Types is
  type Point is record
    X, Y : Natural;
  end record;

  type Point_Access is access Point;
end Access_Types;
```

8.3.1 Allocation (by type)

Once we have declared an access type, we need a way to give variables of the types a meaningful value! You can allocate a value of an access type with the new keyword in Ada.

Listing 12: access_types.ads

```

1 with Dates; use Dates;
2
3 package Access_Types is
4   type Date_Acc is access Date;
```

(continues on next page)

(continued from previous page)

```

5   D : Date_Acc := new Date;
6   --           ^ Allocate a new Date record
7   end Access_Types;
8

```

If the type you want to allocate needs constraints, you can put them in the subtype indication, just as you would do in a variable declaration:

Listing 13: access_types.ads

```

1  with Dates; use Dates;
2
3  package Access_Types is
4    type String_Acc is access String;
5    --           ^ Access to unconstrained array type
6    Msg : String_Acc;
7    --           ^ Default value is null
8
9    Buffer : String_Acc := new String (1 .. 10);
10   --           ^ Constraint required
11 end Access_Types;

```

In some cases, though, allocating just by specifying the type is not ideal, so Ada also allows you to initialize along with the allocation. This is done via the qualified expression syntax:

Listing 14: access_types.ads

```

1  with Dates; use Dates;
2
3  package Access_Types is
4    type Date_Acc is access Date;
5    type String_Acc is access String;
6
7    D : Date_Acc := new Date'(30, November, 2011);
8    Msg : String_Acc := new String'("Hello");
9  end Access_Types;

```

8.3.2 Dereferencing

The last important piece of Ada's access type facility is how to get from an access value to the object that is pointed to, that is, how to dereference the pointer. Dereferencing a pointer uses the `.all` syntax in Ada, but is often not needed - in many cases, the access value will be implicitly dereferenced for you:

Listing 15: access_types.ads

```

1  with Dates; use Dates;
2
3  package Access_Types is
4    type Date_Acc is access Date;
5
6    D : Date_Acc := new Date'(30, November, 2011);
7
8    Today : Date := D.all;
9    --           ^ Access value dereference
10   J : Integer := D.Day;
11   --           ^ Implicit dereference for record and array components
12   --           Equivalent to D.all.day
13 end Access_Types;

```

8.3.3 Other features

As you might know if you have used pointers in C or C++, we are still missing features that are considered fundamental to the use of pointers, such as:

- Pointer arithmetic (being able to increment or decrement a pointer in order to point to the next or previous object)
- Manual deallocation - what is called `free` or `delete` in C. This is a potentially unsafe operation. To keep within the realm of safe Ada, you need to never deallocate manually.

Those features exist in Ada, but are only available through specific standard library APIs.

Attention: The guideline in Ada is that most of the time you can avoid manual allocation, and you should.

There are many ways to avoid manual allocation, some of which have been covered (such as parameter modes). The language also provides library abstractions to avoid pointers:

1. One is the use of *containers* (page 185). Containers help users avoid pointers, because container memory is automatically managed.
2. A container to note in this context is the *Indefinite holder*¹². This container allows you to store a value of an indefinite type such as `String`.
3. GNATCOLL has a library for smart pointers, called *RefCount*¹³. Those pointers' memory is automatically managed, so that when an allocated object has no more references to it, the memory is automatically deallocated.

8.4 Mutually recursive types

The linked list is a common idiom in data structures; in Ada this would be most naturally defined through two types, a record type and an access type, that are mutually dependent. To declare mutually dependent types, you can use an incomplete type declaration:

¹² <http://www.ada-auth.org/standards/12rat/html/Rat12-8-5.html>

¹³ <https://github.com/AdaCore/gnatcoll-core/blob/master/src/gnatcoll-refcount.ads>

Listing 16: simple_list.ads

```
1 package Simple_List is
2   type Node;
3   -- This is an incomplete type declaration, which is
4   -- completed in the same declarative region.
5
6   type Node_Acc is access Node;
7
8   type Node is record
9     Content    : Natural;
10    Prev, Next : Node_Acc;
11  end record;
12 end Simple_List;
```

8.5 More about records

8.5.1 Dynamically sized record types

We have previously seen some simple examples of record types. Let's now look at some of the more advanced properties of this fundamental language feature.

One point to note is that object size for a record type does not need to be known at compile time. This is illustrated in the example below:

Listing 17: runtime_length.ads

```
1 package Runtime_Length is
2   function Compute_Max_Len return Natural;
3 end Runtime_Length;
```

Listing 18: var_size_record.ads

```
1 with Runtime_Length; use Runtime_Length;
2
3 package Var_Size_Record is
4   Max_Len : constant Natural := Compute_Max_Len;
5   -- ^ Not known at compile time
6
7   type Items_Array is array (Positive range <>) of Integer;
8
9   type Growable_Stack is record
10     Items : Items_Array (1 .. Max_Len);
11     Len   : Natural;
12  end record;
13  -- Growable_Stack is a definite type, but size is not known at compile
14  -- time.
15
16  G : Growable_Stack;
17 end Var_Size_Record;
```

It is completely fine to determine the size of your records at run time, but note that all objects of this type will have the same size.

8.5.2 Records with discriminant

In the example above, the size of the `Items` field is determined once, at run-time, but every `Growable_Stack` instance will be exactly the same size. But maybe that's not what you want to do. We saw that arrays in general offer this flexibility: for an unconstrained array type, different objects can have different sizes.

You can get analogous functionality for records, too, using a special kind of field that is called a discriminant:

Listing 19: `var_size_record_2.ads`

```

1 package Var_Size_Record_2 is
2   type Items_Array is array (Positive range <>) of Integer;
3
4   type Growable_Stack (Max_Len : Natural) is record
5     --           ^ Discriminant. Cannot be modified once initialized.
6     Items : Items_Array (1 .. Max_Len);
7     Len   : Natural := 0;
8   end record;
9   -- Growable_Stack is an indefinite type (like an array)
10 end Var_Size_Record_2;
```

Discriminants, in their simple forms, are constant: You cannot modify them once you have initialized the object. This intuitively makes sense since they determine the size of the object.

Also, they make a type indefinite: Whether or not the discriminant is used to specify the size of an object, a type with a discriminant will be indefinite if the discriminant is not declared with an initialization:

Listing 20: `test_discriminants.ads`

```

1 package Test_Discriminants is
2   type Point (X, Y : Natural) is record
3     null;
4   end record;
5
6   P : Point;
7   -- ERROR: Point is indefinite, so you need to specify the discriminants
8   -- or give a default value
9
10  P2 : Point (1, 2);
11  P3 : Point := (1, 2);
12  -- Those two declarations are equivalent.
13
14 end Test_Discriminants;
```

Compilation output

```

test_discriminants.ads:6:08: unconstrained subtype not allowed (need_
↳ initialization)
test_discriminants.ads:6:08: provide initial value or explicit discriminant values
test_discriminants.ads:6:08: or give default discriminant values for type "Point"
```

This also means that, in the example above, you cannot declare an array of `Point` values, because the size of a `Point` is not known.

In most other respects discriminants behave like regular fields: You have to specify their values in aggregates, as seen above, and you can access their values via the dot notation.

Listing 21: main.adb

```

1 with Var_Size_Record_2; use Var_Size_Record_2;
2 with Ada.Text_IO; use Ada.Text_IO;
3
4 procedure Main is
5   procedure Print_Stack (G : Growable_Stack) is
6     begin
7       Put ("<Stack, items: [");
8       for I in G.Items'Range loop
9         exit when I > G.Len;
10        Put (" " & Integer'Image (G.Items (I)));
11      end loop;
12      Put_Line ("]>");
13    end Print_Stack;
14
15    S : Growable_Stack :=
16      (Max_Len => 128, Items => (1, 2, 3, 4, others => <>), Len => 4);
17  begin
18    Print_Stack (S);
19  end Main;

```

Build output

```

Compile
  [Ada]      main.adb
  [Ada]      var_size_record_2.ads
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb

```

Runtime output

```
<Stack, items: [ 1 2 3 4]>
```

8.5.3 Variant records

The examples of discriminants thus far have illustrated the declaration of records of varying size, by having components whose size depends on the discriminant.

However, discriminants can also be used to obtain the functionality of what are sometimes called "variant records": records that can contain different sets of fields.

Listing 22: variant_record.ads

```

1 package Variant_Record is
2   type Expr; -- Forward declaration of Expr
3   type Expr_Access is access Expr; -- Access to a Expr
4
5   type Expr_Kind_Type is (Bin_Op_Plus, Bin_Op_Minus, Num);
6   -- A regular enumeration type
7
8   type Expr (Kind : Expr_Kind_Type) is record
9     -- ^ The discriminant is an enumeration value
10    case Kind is
11      when Bin_Op_Plus | Bin_Op_Minus =>
12        Left, Right : Expr_Access;
13      when Num =>

```

(continues on next page)

(continued from previous page)

```

14     Val : Integer;
15     end case;
16     -- Variant part. Only one, at the end of the record
17     -- definition, but can be nested
18     end record;
19 end Variant_Record;

```

The fields that are in a when branch will be only available when the value of the discriminant is covered by the branch. In the example above, you will only be able to access the fields `Left` and `Right` when the `Kind` is `Bin_Op_Plus` or `Bin_Op_Minus`.

If you try to access a field that is not valid for your record, a `Constraint_Error` will be raised.

Listing 23: main.adb

```

1 with Variant_Record; use Variant_Record;
2
3 procedure Main is
4   E : Expr := (Num, 12);
5 begin
6   E.Left := new Expr'(Num, 15);
7   -- Will compile but fail at runtime
8 end Main;

```

Build output

```

Compile
  [Ada]          main.adb
main.adb:6:05: warning: component not present in subtype of "Expr" defined at line_
↪4
main.adb:6:05: warning: "Constraint_Error" will be raised at run time
  [Ada]          variant_record.ads
Bind
  [gprbind]     main.bexch
  [Ada]         main.ali
Link
  [link]        main.adb

```

Runtime output

```

raised CONSTRAINT_ERROR : main.adb:6 discriminant check failed

```

Here is how you could write an evaluator for expressions:

Listing 24: main.adb

```

1 with Variant_Record; use Variant_Record;
2 with Ada.Text_IO; use Ada.Text_IO;
3
4 procedure Main is
5   function Eval_Expr (E : Expr) return Integer is
6     (case E.Kind is
7     when Bin_Op_Plus => Eval_Expr (E.Left.all) + Eval_Expr (E.Right.all),
8     when Bin_Op_Minus => Eval_Expr (E.Left.all) - Eval_Expr (E.Right.all),
9     when Num => E.Val);
10
11   E : Expr := (Bin_Op_Plus,
12               new Expr'(Bin_Op_Minus,
13                           new Expr'(Num, 12), new Expr'(Num, 15)),
14               new Expr'(Num, 3));

```

(continues on next page)

(continued from previous page)

```

15 begin
16   Put_Line (Integer'Image (Eval_Expr (E)));
17 end Main;
```

Build output

```

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

Runtime output

```
0
```

In other languages

Ada's variant records are very similar to Sum types in functional languages such as OCaml or Haskell. A major difference is that the discriminant is a separate field in Ada, whereas the 'tag' of a Sum type is kind of built in, and only accessible with pattern matching.

There are other differences (you can have several discriminants in a variant record in Ada). Nevertheless, they allow the same kind of type modeling as sum types in functional languages.

Compared to C/C++ unions, Ada variant records are more powerful in what they allow, and are also checked at run time, which makes them safer.

8.6 Fixed-point types

8.6.1 Decimal fixed-point types

We have already seen how to specify floating-point types. However, in some applications floating-point is not appropriate since, for example, the roundoff error from binary arithmetic may be unacceptable or perhaps the hardware does not support floating-point instructions. Ada provides a category of types, the decimal fixed-point types, that allows the programmer to specify the required decimal precision (number of digits) as well as the scaling factor (a power of ten) and, optionally, a range. In effect the values will be represented as integers implicitly scaled by the specified power of 10. This is useful, for example, for financial applications.

The syntax for a simple decimal fixed-point type is

```
type <type-name> is delta <delta-value> digits <digits-value>;
```

In this case, the `delta` and the `digits` will be used by the compiler to derive a range.

Several attributes are useful for dealing with decimal types:

Attribute Name	Meaning
First	The first value of the type
Last	The last value of the type
Delta	The delta value of the type

In the example below, we declare two data types: T3_D3 and T6_D3. For both types, the delta value is the same: 0.001.

Listing 25: decimal_fixed_point_types.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Decimal_Fixed_Point_Types is
4   type T3_D3 is delta 10.0 ** (-3) digits 3;
5   type T6_D3 is delta 10.0 ** (-3) digits 6;
6 begin
7   Put_Line ("The delta value of T3_D3 is " & T3_D3'Image (T3_D3'Delta));
8   Put_Line ("The minimum value of T3_D3 is " & T3_D3'Image (T3_D3'First));
9   Put_Line ("The maximum value of T3_D3 is " & T3_D3'Image (T3_D3'Last));
10  New_Line;
11  Put_Line ("The delta value of T6_D3 is " & T6_D3'Image (T6_D3'Delta));
12  Put_Line ("The minimum value of T6_D3 is " & T6_D3'Image (T6_D3'First));
13  Put_Line ("The maximum value of T6_D3 is " & T6_D3'Image (T6_D3'Last));
14 end Decimal_Fixed_Point_Types;
```

Build output

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

Runtime output

```

The delta value of T3_D3 is 0.001
The minimum value of T3_D3 is -0.999
The maximum value of T3_D3 is 0.999

The delta value of T6_D3 is 0.001
The minimum value of T6_D3 is -999.999
The maximum value of T6_D3 is 999.999
```

When running the application, we see that the delta value of both types is indeed the same: 0.001. However, because T3_D3 is restricted to 3 digits, its range is -0.999 to 0.999. For the T6_D3, we have defined a precision of 6 digits, so the range is -999.999 to 999.999.

Similar to the type definition using the range syntax, because we have an implicit range, the compiled code will check that the variables contain values that are not out-of-range. Also, if the result of a multiplication or division on decimal fixed-point types is smaller than the delta value required for the context, the actual result will be zero. For example:

Listing 26: decimal_fixed_point_smaller.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Decimal_Fixed_Point_Smaller is
4   type T3_D3 is delta 10.0 ** (-3) digits 3;
5   type T6_D6 is delta 10.0 ** (-6) digits 6;
6   A : T3_D3 := T3_D3'Delta;
7   B : T3_D3 := 0.5;
8   C : T6_D6;
9 begin
10  Put_Line ("The value of A is " & T3_D3'Image (A));
11  A := A * B;
```

(continues on next page)

(continued from previous page)

```

12   Put_Line ("The value of A * B is " & T3_D3'Image (A));
13   A := T3_D3'Delta;
14   C := A * B;
15   Put_Line ("The value of A * B is " & T6_D6'Image (C));
16   end Decimal_Fixed_Point_Smaller;

```

Build output

```

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

```

Runtime output

```

The value of A      is 0.001
The value of A * B is 0.000
The value of A * B is 0.000500

```

In this example, the result of the operation $0.001 * 0.5$ is 0.0005 . Since this value is not representable for the `T3_D3` type because the `delta` value is 0.001 , the actual value stored in variable `A` is zero. However, accuracy is preserved during the arithmetic operations if the target has sufficient precision, and the value displayed for `C` is 0.000500 .

8.6.2 Fixed-point types

Ordinary fixed-point types are similar to decimal fixed-point types in that the values are, in effect, scaled integers. The difference between them is in the scale factor: for a decimal fixed-point type, the scaling, given explicitly by the type's `delta`, is always a power of ten.

In contrast, for an ordinary fixed-point type, the scaling 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.

Note: Ordinary fixed-point types can be thought of being closer to the actual representation on the machine, since hardware support for decimal fixed-point arithmetic is not widespread (rescalings by a power of ten), while ordinary fixed-point types make use of the available integer shift instructions.

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:

Listing 27: `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);

```

(continues on next page)

(continued from previous page)

```

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

```

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 exact one. This is a typical effect of fixed-point data types — you can find more details in this discussion about the *Q format*¹⁴. We may also rewrite this code with an exact type definition:

Listing 28: normalized_adapted_fixed_point_type.adb

```

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

```

We may also use any other range. For example:

Listing 29: 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 T_Inv_Trig is delta 2.0 ** (-15) * Pi range -Pi / 2.0 .. Pi / 2.0;
6  begin
7    Put_Line ("T_Inv_Trig requires " & Integer'Image (T_Inv_Trig'Size)
8              & " bits");
9    Put_Line ("The delta value of T_Inv_Trig is "
10             & T_Inv_Trig'Image (T_Inv_Trig'Delta));
11   Put_Line ("The minimum value of T_Inv_Trig is "
12             & T_Inv_Trig'Image (T_Inv_Trig'First));
13   Put_Line ("The maximum value of T_Inv_Trig is "
14             & T_Inv_Trig'Image (T_Inv_Trig'Last));
15  end Custom_Fixed_Point_Range;

```

Build output

¹⁴ [https://en.wikipedia.org/wiki/Q_\(number_format\)](https://en.wikipedia.org/wiki/Q_(number_format))

```
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

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

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

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

Listing 30: `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;
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 fact the language is more general than these examples imply, since in practice it is typical to need to multiply or divide values from different fixed-point types, and obtain a result that may be of a third fixed-point type. The details are outside the scope of this introductory course.

It is also worth noting, although again the details are outside the scope of this course, that you can explicitly specify a value for an ordinary fixed-point type's `small`. This allows non-binary scaling, for example:

```
type Angle is delta 1.0/3600.0 range 0.0 .. 360.0 - 1.0/3600.0;
for Angle'Small use Angle'Delta;
```


8.7 Character types

As noted earlier, each enumeration type is distinct and incompatible with every other enumeration type. However, what we did not mention previously is that character literals are permitted as enumeration literals. This means that in addition to the language's strongly typed character types, user-defined character types are also permitted:

Listing 31: character_example.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Character_Example is
4   type My_Char is ('a', 'b', 'c');
5   -- Our custom character type, an enumeration type with 3 valid values.
6
7   C : Character;
8   -- ^ Built-in character type (it's an enumeration type)
9
10  M : My_Char;
11 begin
12   C := '?';
13   -- ^ Character literal (enumeration literal)
14
15   M := 'a';
16
17   C := 65;
18   -- ^ Invalid: 65 is not a Character value
19
20   C := Character'Val (65);
21   -- Assign the character at position 65 in the enumeration (which is 'A')
22
23   M := C;
24   -- ^ Invalid: C is of type Character, and M is a My_Char
25
26   M := 'd';
27   -- ^ Invalid: 'd' is not a valid literal for type My_Char
28 end Character_Example;

```


PRIVACY

One of the main principles of modular programming, as well as object oriented programming, is [encapsulation](https://en.wikipedia.org/wiki/Encapsulation_(computer_programming))¹⁵.

Encapsulation, briefly, is the concept that the implementer of a piece of software will distinguish between the code's public interface and its private implementation.

This is not only applicable to software libraries but wherever abstraction is used.

In Ada, the granularity of encapsulation is a bit different from most object-oriented languages, because privacy is generally specified at the package level.

9.1 Basic encapsulation

Listing 1: encapsulate.ads

```
1 package Encapsulate is
2   procedure Hello;
3
4 private
5
6   procedure Hello2;
7   -- Not visible from external units
8 end Encapsulate;
```

Listing 2: main.adb

```
1 with Encapsulate;
2
3 procedure Main is
4 begin
5   Encapsulate.Hello;
6   Encapsulate.Hello2;
7   -- Invalid: Hello2 is not visible
8 end Main;
```

Compilation output

```
main.adb:6:15: "Hello2" is not a visible entity of "Encapsulate"
```

¹⁵ [https://en.wikipedia.org/wiki/Encapsulation_\(computer_programming\)](https://en.wikipedia.org/wiki/Encapsulation_(computer_programming))

9.2 Abstract data types

With this high-level granularity, it might not seem obvious how to hide the implementation details of a type. Here is how it can be done in Ada:

Listing 3: stacks.ads

```

1 package Stacks is
2   type Stack is private;
3   -- Declare a private type: You cannot depend on its
4   -- implementation. You can only assign and test for
5   -- equality.
6
7   procedure Push (S : in out Stack; Val : Integer);
8   procedure Pop (S : in out Stack; Val : out Integer);
9 private
10
11   subtype Stack_Index is Natural range 1 .. 10;
12   type Content_Type is array (Stack_Index) of Natural;
13
14   type Stack is record
15     Top      : Stack_Index;
16     Content : Content_Type;
17   end record;
18 end Stacks;
```

In the above example, we define a stack type in the public part (known as the *visible part* of the package spec in Ada), but the exact representation of that type is private.

Then, in the private part, we define the representation of that type. We can also declare other types that will be used as *helpers* for our main public type. This is useful since declaring helper types is common in Ada.

A few words about terminology:

- The Stack type as viewed from the public part is called the partial view of the type. This is what clients have access to.
- The Stack type as viewed from the private part or the body of the package is called the full view of the type. This is what implementers have access to.

From the point of view of the client (the *with*'ing unit), only the public (visible) part is important, and the private part could as well not exist. It makes it very easy to read linearly the part of the package that is important for you.

```

-- No need to read the private part to use the package
package Stacks is
  type Stack is private;

  procedure Push (S : in out Stack; Val : Integer);
  procedure Pop (S : in out Stack; Val : out Integer);
private
  ...
end Stacks;
```

Here is how the Stacks package would be used:

```

-- Example of use
with Stacks; use Stacks;

procedure Test_Stack is
  S : Stack;
```

(continues on next page)

(continued from previous page)

```

    Res : Integer;
begin
    Push (S, 5);
    Push (S, 7);
    Pop (S, Res);
end Test_Stack;

```

9.3 Limited types

Ada's *limited type* facility allows you to declare a type for which assignment and comparison operations are not automatically provided.

Listing 4: stacks.ads

```

1 package Stacks is
2   type Stack is limited private;
3   -- Limited type. Cannot assign nor compare.
4
5   procedure Push (S : in out Stack; Val : Integer);
6   procedure Pop (S : in out Stack; Val : out Integer);
7 private
8   subtype Stack_Index is Natural range 1 .. 10;
9   type Content_Type is array (Stack_Index) of Natural;
10
11  type Stack is limited record
12    Top      : Stack_Index;
13    Content : Content_Type;
14  end record;
15 end Stacks;

```

Listing 5: main.adb

```

1 with Stacks; use Stacks;
2
3 procedure Main is
4   S, S2 : Stack;
5 begin
6   S := S2;
7   -- Illegal: S is limited.
8 end Main;

```

Compilation output

```
main.adb:6:04: left hand of assignment must not be limited type
```

This is useful because, for example, for some data types the built-in assignment operation might be incorrect (for example when a deep copy is required).

Ada does allow you to overload the comparison operators = and /= for limited types (and to override the built-in declarations for non-limited types).

Ada also allows you to implement special semantics for assignment via [controlled types](#)¹⁶. However, in some cases assignment is simply inappropriate; one example is the `File_Type` from the `Ada.Text_IO` package, which is declared as a limited type and thus attempts to assign one file to another would be detected as illegal.

¹⁶ https://www.adaic.org/resources/add_content/standards/12rm/html/RM-7-6.html

9.4 Child packages & privacy

We've seen previously (in the *child packages section* (page 33)) that packages can have child packages. Privacy plays an important role in child packages. This section discusses some of the privacy rules that apply to child packages.

Although the private part of a package P is meant to encapsulate information, certain parts of a child package $P.C$ can have access to this private part of P . In those cases, information from the private part of P can then be used as if it were declared in the public part of its specification. To be more specific, the body of $P.C$ and the private part of the specification of $P.C$ have access to the private part of P . However, the public part of the specification of $P.C$ only has access to the public part of P 's specification. The following table summarizes this:

Part of a child package	Access to the private part of its parent's specification
Specification: public part	
Specification: private part	X
Body	X

The rest of this section shows examples of how this access to private information actually works for child packages

Let's first look at an example where the body of a child package $P.C$ has access to the private part of the specification of its parent P . We've seen, in a previous source-code example, that the `Hello2` procedure declared in the private part of the `Encapsulate` package cannot be used in the `Main` procedure, since it's not visible there. This limitation doesn't apply, however, for parts of the child packages of the `Encapsulate` package. In fact, the body of its child package `Encapsulate.Child` has access to the `Hello2` procedure and can call it there, as you can see in the implementation of the `Hello3` procedure of the `Child` package:

Listing 6: encapsulate.ads

```

1 package Encapsulate is
2   procedure Hello;
3
4 private
5
6   procedure Hello2;
7   -- Not visible from external units
8   -- But visible in child packages
9 end Encapsulate;
```

Listing 7: encapsulate.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body Encapsulate is
4
5   procedure Hello is
6   begin
7     Put_Line ("Hello");
8   end Hello;
9
10  procedure Hello2 is
11  begin
12    Put_Line ("Hello #2");
13  end Hello2;
14
15 end Encapsulate;
```

Listing 8: encapsulate-child.ads

```

1 package Encapsulate.Child is
2
3   procedure Hello3;
4
5 end Encapsulate.Child;
```

Listing 9: encapsulate-child.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body Encapsulate.Child is
4
5   procedure Hello3 is
6   begin
7     -- Using private procedure Hello2 from the parent package
8     Hello2;
9     Put_Line ("Hello #3");
10  end Hello3;
11
12 end Encapsulate.Child;
```

Listing 10: main.adb

```

1 with Encapsulate.Child;
2
3 procedure Main is
4 begin
5   Encapsulate.Child.Hello3;
6 end Main;
```

Build output

```

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

Runtime output

```

Hello #2
Hello #3
```

The same mechanism applies to types declared in the private part of a parent package. For instance, the body of a child package can access components of a record declared in the private part of its parent package. Let's look at an example:

Listing 11: my_types.ads

```

1 package My_Types is
2
3   type Priv_Rec is private;
4
5 private
```

(continues on next page)

(continued from previous page)

```
6
7  type Priv_Rec is record
8     Number : Integer := 42;
9  end record;
10
11 end My_Types;
```

Listing 12: my_types-ops.ads

```
1 package My_Types.Ops is
2
3     procedure Display (E : Priv_Rec);
4
5 end My_Types.Ops;
```

Listing 13: my_types-ops.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body My_Types.Ops is
4
5     procedure Display (E : Priv_Rec) is
6     begin
7         Put_Line ("Priv_Rec.Number: " & Integer'Image (E.Number));
8     end Display;
9
10 end My_Types.Ops;
```

Listing 14: main.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 with My_Types; use My_Types;
4 with My_Types.Ops; use My_Types.Ops;
5
6 procedure Main is
7     E : Priv_Rec;
8 begin
9     Put_Line ("Presenting information:");
10
11     -- The following line would trigger a compilation error here:
12     -- Put_Line ("Priv_Rec.Number: " & Integer'Image (E.Number));
13
14     Display (E);
15 end Main;
```

Build output

```
Compile
  [Ada]      main.adb
  [Ada]      my_types.ads
  [Ada]      my_types-ops.adb
Bind
  [gprbind]  main.bexch
  [Ada]      main.ali
Link
  [link]     main.adb
```

Runtime output


```
Presenting information:
Priv_Rec.Number: 42
```

In this example, we don't have access to the `Number` component of the record type `Priv_Rec` in the `Main` procedure. You can see this in the call to `Put_Line` that has been commented-out in the implementation of `Main`. Trying to access the `Number` component there would trigger a compilation error. But we do have access to this component in the body of the `My_Types.Ops` package, since it's a child package of the `My_Types` package. Therefore, `Ops`'s body has access to the declaration of the `Priv_Rec` type — which is in the private part of its parent, the `My_Types` package. For this reason, the same call to `Put_Line` that would trigger a compilation error in the `Main` procedure works fine in the `Display` procedure of the `My_Types.Ops` package.

This kind of privacy rules for child packages allows for extending the functionality of a parent package and, at the same time, retain its encapsulation.

As we mentioned previously, in addition to the package body, the private part of the specification of a child package `P.C` also has access to the private part of the specification of its parent `P`. Let's look at an example where we declare an object of private type `Priv_Rec` in the private part of the child package `My_Types.Child` and initialize the `Number` component of the `Priv_Rec` record directly:

```
package My_Types.Child is
private
  E : Priv_Rec := (Number => 99);
end My_Types.Ops;
```

As expected, we wouldn't be able to initialize this component if we moved this declaration to the public (visible) part of the same child package:

```
package My_Types.Child is
  E : Priv_Rec := (Number => 99);
end My_Types.Ops;
```

The declaration above triggers a compilation error, since type `Priv_Rec` is private. Because the public part of `My_Types.Child` is also visible outside the child package, Ada cannot allow accessing private information in this part of the specification.

GENERICICS

10.1 Introduction

Generics are used for metaprogramming in Ada. They are useful for abstract algorithms that share common properties with each other.

Either a subprogram or a package can be generic. A generic is declared by using the keyword `generic`. For example:

Listing 1: operator.ads

```
1 generic
2   type T is private;
3   -- Declaration of formal types and objects
4   -- Below, we could use one of the following:
5   -- <procedure | function | package>
6   procedure Operator (Dummy : in out T);
```

Listing 2: operator.adb

```
1 procedure Operator (Dummy : in out T) is
2   begin
3     null;
4   end Operator;
```

10.2 Formal type declaration

Formal types are abstractions of a specific type. For example, we may want to create an algorithm that works on any integer type, or even on any type at all, whether a numeric type or not. The following example declares a formal type `T` for the `Set` procedure.

Listing 3: set.ads

```
1 generic
2   type T is private;
3   -- T is a formal type that indicates that any type can be used,
4   -- possibly a numeric type or possibly even a record type.
5   procedure Set (Dummy : T);
```

Listing 4: set.adb

```
1 procedure Set (Dummy : T) is
2   begin
```

(continues on next page)

(continued from previous page)

```

3   null;
4 end Set;
```

The declaration of T as private indicates that you can map any definite type to it. But you can also restrict the declaration to allow only some types to be mapped to that formal type. Here are some examples:

Formal Type	Format
Any type	type T is private;
Any discrete type	type T is (<>);
Any floating-point type	type T is digits <>;

10.3 Formal object declaration

Formal objects are similar to subprogram parameters. They can reference formal types declared in the formal specification. For example:

Listing 5: set.ads

```

1 generic
2   type T is private;
3   X : in out T;
4   -- X can be used in the Set procedure
5 procedure Set (E : T);
```

Listing 6: set.adb

```

1 procedure Set (E : T) is
2   pragma Unreferenced (E, X);
3 begin
4   null;
5 end Set;
```

Formal objects can be either input parameters or specified using the in out mode.

10.4 Generic body definition

We don't repeat the generic keyword for the body declaration of a generic subprogram or package. Instead, we start with the actual declaration and use the generic types and objects we declared. For example:

Listing 7: set.ads

```

1 generic
2   type T is private;
3   X : in out T;
4 procedure Set (E : T);
```

Listing 8: set.adb

```

1 procedure Set (E : T) is
2 -- Body definition: "generic" keyword is not used
3 begin
```

(continues on next page)

(continued from previous page)

```

4   X := E;
5  end Set;

```

10.5 Generic instantiation

Generic subprograms or packages can't be used directly. Instead, they need to be instantiated, which we do using the new keyword, as shown in the following example:

Listing 9: set.ads

```

1  generic
2    type T is private;
3    X : in out T;
4    -- X can be used in the Set procedure
5  procedure Set (E : T);

```

Listing 10: set.adb

```

1  procedure Set (E : T) is
2  begin
3    X := E;
4  end Set;

```

Listing 11: show_generic_instantiation.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2  with Set;
3
4  procedure Show_Generic_Instantiation is
5
6    Main      : Integer := 0;
7    Current  : Integer;
8
9    procedure Set_Main is new Set (T => Integer,
10                                   X => Main);
11    -- Here, we map the formal parameters to actual types and objects.
12    --
13    -- The same approach can be used to instantiate functions or
14    -- packages, e.g.:
15    -- function Get_Main is new ...
16    -- package Integer_Queue is new ...
17
18  begin
19    Current := 10;
20
21    Set_Main (Current);
22    Put_Line ("Value of Main is " & Integer'Image (Main));
23  end Show_Generic_Instantiation;

```

Build output

```

Compile
  [Ada]      show_generic_instantiation.adb
  [Ada]      set.adb
Bind
  [gprbind]  show_generic_instantiation.bexch
  [Ada]      show_generic_instantiation.ali

```

(continues on next page)

(continued from previous page)

```
Link
  [link]          show_generic_instantiation.adb
```

Runtime output

```
Value of Main is 10
```

In the example above, we instantiate the procedure `Set` by mapping the formal parameters `T` and `X` to actual existing elements, in this case the `Integer` type and the `Main` variable.

10.6 Generic packages

The previous examples focused on generic subprograms. In this section, we look at generic packages. The syntax is similar to that used for generic subprograms: we start with the `generic` keyword and continue with formal declarations. The only difference is that `package` is specified instead of a subprogram keyword.

Here's an example:

Listing 12: element.ads

```

1  generic
2    type T is private;
3  package Element is
4
5    procedure Set (E : T);
6    procedure Reset;
7    function Get return T;
8    function Is_Valid return Boolean;
9
10   Invalid_Element : exception;
11
12  private
13   Value : T;
14   Valid : Boolean := False;
15  end Element;
```

Listing 13: element.adb

```

1  package body Element is
2
3    procedure Set (E : T) is
4    begin
5      Value := E;
6      Valid := True;
7    end Set;
8
9    procedure Reset is
10   begin
11     Valid := False;
12   end Reset;
13
14   function Get return T is
15   begin
16     if not Valid then
17       raise Invalid_Element;
18     end if;
19     return Value;
```

(continues on next page)

(continued from previous page)

```

20   end Get;
21
22   function Is_Valid return Boolean is (Valid);
23 end Element;
```

Listing 14: show_generic_package.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2  with Element;
3
4  procedure Show_Generic_Package is
5
6     package I is new Element (T => Integer);
7
8     procedure Display_Initialized is
9     begin
10        if I.Is_Valid then
11            Put_Line ("Value is initialized");
12        else
13            Put_Line ("Value is not initialized");
14        end if;
15    end Display_Initialized;
16
17    begin
18        Display_Initialized;
19
20        Put_Line ("Initializing...");
21        I.Set (5);
22        Display_Initialized;
23        Put_Line ("Value is now set to " & Integer'Image (I.Get));
24
25        Put_Line ("Reseting...");
26        I.Reset;
27        Display_Initialized;
28
29    end Show_Generic_Package;
```

Build output

```

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

Runtime output

```

Value is not initialized
Initializing...
Value is initialized
Value is now set to 5
Reseting...
Value is not initialized
```

In the example above, we created a simple container named `Element`, with just one single element. This container tracks whether the element has been initialized or not.

After writing package definition, we create the instance `I` of the `Element`. We use the instance by calling the package subprograms (`Set`, `Reset`, and `Get`).

10.7 Formal subprograms

In addition to formal types and objects, we can also declare formal subprograms or packages. This course only describes formal subprograms; formal packages are discussed in the advanced course.

We use the `with` keyword to declare a formal subprogram. In the example below, we declare a formal function (`Comparison`) to be used by the generic procedure `Check`.

Listing 15: check.ads

```

1 generic
2   Description : String;
3   type T is private;
4   with function Comparison (X, Y : T) return Boolean;
5 procedure Check (X, Y : T);

```

Listing 16: check.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Check (X, Y : T) is
4   Result : Boolean;
5 begin
6   Result := Comparison (X, Y);
7   if Result then
8     Put_Line ("Comparison (" & Description &
9              ") between arguments is OK!");
10
11  else
12    Put_Line ("Comparison (" & Description &
13             ") between arguments is not OK!");
14  end if;
15 end Check;

```

Listing 17: show_formal_subprogram.adb

```

1 with Check;
2
3 procedure Show_Forma1_Subprogram is
4
5   A, B : Integer;
6
7   procedure Check_Is_Equal is new Check (Description => "equality",
8                                          T           => Integer,
9                                          Comparison => Standard."=");
10  -- Here, we are mapping the standard equality operator for Integer
11  -- types to the Comparison formal function
12 begin
13   A := 0;
14   B := 1;
15   Check_Is_Equal (A, B);
16 end Show_Forma1_Subprogram;

```

Build output

```

Compile
  [Ada]      show_formal_subprogram.adb
  [Ada]      check.adb
Bind
  [gprbind]  show_formal_subprogram.bexch
  [Ada]      show_formal_subprogram.ali
Link

```

(continues on next page)

(continued from previous page)

```
[link]      show_formal_subprogram.adb
```

Runtime output

```
Comparison (equality) between arguments is not OK!
```

10.8 Example: I/O instances

Ada offers generic I/O packages that can be instantiated for standard and derived types. One example is the generic `Float_IO` package, which provides procedures such as `Put` and `Get`. In fact, `Float_Text_IO` — available from the standard library — is an instance of the `Float_IO` package, and it's defined as:

```
with Ada.Text_IO;

package Ada.Float_Text_IO is new Ada.Text_IO.Float_IO (Float);
```

You can use it directly with any object of floating-point type. For example:

Listing 18: show_float_text_io.adb

```
1 with Ada.Float_Text_IO;
2
3 procedure Show_Float_Text_IO is
4   X : constant Float := 2.5;
5
6   use Ada.Float_Text_IO;
7 begin
8   Put (X);
9 end Show_Float_Text_IO;
```

Build output

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

Runtime output

```
2.50000E+00
```

Instantiating generic I/O packages can be useful for derived types. For example, let's create a new type `Price` that must be displayed with two decimal digits after the point, and no exponent.

Listing 19: show_float_io_inst.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Float_IO_Inst is
4
5   type Price is digits 3;
6
7   package Price_IO is new Ada.Text_IO.Float_IO (Price);
```

(continues on next page)

(continued from previous page)

```

8
9   P : Price;
10  begin
11    -- Set to zero => don't display exponent
12    Price_IO.Default_Exp := 0;
13
14    P := 2.5;
15    Price_IO.Put (P);
16    New_Line;
17
18    P := 5.75;
19    Price_IO.Put (P);
20    New_Line;
21  end Show_Float_IO_Inst;
```

Build output

```

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

Runtime output

```

2.50
5.75
```

By adjusting `Default_Exp` from the `Price_IO` instance to *remove* the exponent, we can control how variables of `Price` type are displayed. Just as a side note, we could also have written:

```

-- [...]

type Price is new Float;

package Price_IO is new Ada.Text_IO.Float_IO (Price);

begin
  Price_IO.Default_Aft := 2;
  Price_IO.Default_Exp := 0;
```

In this case, we're adjusting `Default_Aft`, too, to get two decimal digits after the point when calling `Put`.

In addition to the generic `Float_IO` package, the following generic packages are available from `Ada.Text_IO`:

- `Enumeration_IO` for enumeration types;
- `Integer_IO` for integer types;
- `Modular_IO` for modular types;
- `Fixed_IO` for fixed-point types;
- `Decimal_IO` for decimal types.

In fact, we could rewrite the example above using decimal types:

Listing 20: show_decimal_io_inst.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Decimal_IO_Inst is
4
5     type Price is delta 10.0 ** (-2) digits 12;
6
7     package Price_IO is new Ada.Text_IO.Decimal_IO (Price);
8
9     P : Price;
10 begin
11     Price_IO.Default_Exp := 0;
12
13     P := 2.5;
14     Price_IO.Put (P);
15     New_Line;
16
17     P := 5.75;
18     Price_IO.Put (P);
19     New_Line;
20 end Show_Decimal_IO_Inst;

```

Build output

```

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

```

Runtime output

```

2.50
5.75

```

10.9 Example: ADTs

An important application of generics is to model abstract data types (ADTs). In fact, Ada includes a library with numerous ADTs using generics: `Ada.Containers` (described in the [containers section](#) (page 185)).

A typical example of an ADT is a stack:

Listing 21: stacks.ads

```

1 generic
2     Max : Positive;
3     type T is private;
4 package Stacks is
5
6     type Stack is limited private;
7
8     Stack_Underflow, Stack_Overflow : exception;
9
10    function Is_Empty (S : Stack) return Boolean;
11

```

(continues on next page)

(continued from previous page)

```

12  function Pop (S : in out Stack) return T;
13
14  procedure Push (S : in out Stack; V : T);
15
16  private
17
18  type Stack_Array is array (Natural range <>) of T;
19
20  Min : constant := 1;
21
22  type Stack is record
23      Container : Stack_Array (Min .. Max);
24      Top       : Natural := Min - 1;
25  end record;
26
27  end Stacks;

```

Listing 22: stacks.adb

```

1  package body Stacks is
2
3      function Is_Empty (S : Stack) return Boolean is
4          (S.Top < S.Container'First);
5
6      function Is_Full (S : Stack) return Boolean is
7          (S.Top >= S.Container'Last);
8
9      function Pop (S : in out Stack) return T is
10         begin
11             if Is_Empty (S) then
12                 raise Stack_Underflow;
13             else
14                 return X : T do
15                     X := S.Container (S.Top);
16                     S.Top := S.Top - 1;
17                 end return;
18             end if;
19         end Pop;
20
21         procedure Push (S : in out Stack; V : T) is
22             begin
23                 if Is_Full (S) then
24                     raise Stack_Overflow;
25                 else
26                     S.Top := S.Top + 1;
27                     S.Container (S.Top) := V;
28                 end if;
29             end Push;
30
31         end Stacks;

```

Listing 23: show_stack.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2  with Stacks;
3
4  procedure Show_Stack is
5
6      package Integer_Stacks is new Stacks (Max => 10,
7                                              T => Integer);

```

(continues on next page)

(continued from previous page)

```

8   use Integer_Stacks;
9
10  Values : Integer_Stacks.Stack;
11
12  begin
13    Push (Values, 10);
14    Push (Values, 20);
15
16    Put_Line ("Last value was " & Integer'Image (Pop (Values)));
17  end Show_Stack;

```

Build output

```

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

```

Runtime output

```
Last value was 20
```

In this example, we first create a generic stack package (Stacks) and then instantiate it to create a stack of up to 10 integer values.

10.10 Example: Swap

Let's look at a simple procedure that swaps variables of type Color:

Listing 24: colors.ads

```

1  package Colors is
2    type Color is (Black, Red, Green, Blue, White);
3
4    procedure Swap_Colors (X, Y : in out Color);
5  end Colors;

```

Listing 25: colors.adb

```

1  package body Colors is
2
3    procedure Swap_Colors (X, Y : in out Color) is
4      Tmp : constant Color := X;
5    begin
6      X := Y;
7      Y := Tmp;
8    end Swap_Colors;
9
10 end Colors;

```

Listing 26: test_non_generic_swap_colors.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Colors;      use Colors;
3
4 procedure Test_Non_Generic_Swap_Colors is
5   A, B, C : Color;
6 begin
7   A := Blue;
8   B := White;
9   C := Red;
10
11   Put_Line ("Value of A is " & Color'Image (A));
12   Put_Line ("Value of B is " & Color'Image (B));
13   Put_Line ("Value of C is " & Color'Image (C));
14
15   New_Line;
16   Put_Line ("Swapping A and C...");
17   New_Line;
18   Swap_Colors (A, C);
19
20   Put_Line ("Value of A is " & Color'Image (A));
21   Put_Line ("Value of B is " & Color'Image (B));
22   Put_Line ("Value of C is " & Color'Image (C));
23 end Test_Non_Generic_Swap_Colors;

```

Build output

```

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

```

Runtime output

```

Value of A is BLUE
Value of B is WHITE
Value of C is RED

Swapping A and C...

Value of A is RED
Value of B is WHITE
Value of C is BLUE

```

In this example, `Swap_Colors` can only be used for the `Color` type. However, this algorithm can theoretically be used for any type, whether an enumeration type or a complex record type with many elements. The algorithm itself is the same: it's only the type that differs. If, for example, we want to swap variables of `Integer` type, we don't want to duplicate the implementation. Therefore, such an algorithm is a perfect candidate for abstraction using generics.

In the example below, we create a generic version of `Swap_Colors` and name it `Generic_Swap`. This generic version can operate on any type due to the declaration of formal type `T`.

Listing 27: generic_swap.ads

```

1 generic
2   type T is private;
3   procedure Generic_Swap (X, Y : in out T);

```

Listing 28: generic_swap.adb

```

1 procedure Generic_Swap (X, Y : in out T) is
2   Tmp : constant T := X;
3 begin
4   X := Y;
5   Y := Tmp;
6 end Generic_Swap;

```

Listing 29: colors.ads

```

1 with Generic_Swap;
2
3 package Colors is
4
5   type Color is (Black, Red, Green, Blue, White);
6
7   procedure Swap_Colors is new Generic_Swap (T => Color);
8
9 end Colors;

```

Listing 30: test_swap_colors.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Colors; use Colors;
3
4 procedure Test_Swap_Colors is
5   A, B, C : Color;
6 begin
7   A := Blue;
8   B := White;
9   C := Red;
10
11   Put_Line ("Value of A is " & Color'Image (A));
12   Put_Line ("Value of B is " & Color'Image (B));
13   Put_Line ("Value of C is " & Color'Image (C));
14
15   New_Line;
16   Put_Line ("Swapping A and C...");
17   New_Line;
18   Swap_Colors (A, C);
19
20   Put_Line ("Value of A is " & Color'Image (A));
21   Put_Line ("Value of B is " & Color'Image (B));
22   Put_Line ("Value of C is " & Color'Image (C));
23 end Test_Swap_Colors;

```

Build output

```

Compile
  [Ada]      test_swap_colors.adb
  [Ada]      colors.ads
  [Ada]      generic_swap.adb
Bind
  [gprbind]  test_swap_colors.bexch
  [Ada]      test_swap_colors.ali
Link
  [link]     test_swap_colors.adb

```

Runtime output

```
Value of A is BLUE
Value of B is WHITE
Value of C is RED
```

```
Swapping A and C...
```

```
Value of A is RED
Value of B is WHITE
Value of C is BLUE
```

As we can see in the example, we can create the same `Swap_Colors` procedure as we had in the non-generic version of the algorithm by declaring it as an instance of the generic `Generic_Swap` procedure. We specify that the generic `T` type will be mapped to the `Color` type by passing it as an argument to the `Generic_Swap` instantiation,

10.11 Example: Reversing

The previous example, with an algorithm to swap two values, is one of the simplest examples of using generics. Next we study an algorithm for reversing elements of an array. First, let's start with a non-generic version of the algorithm, one that works specifically for the `Color` type:

Listing 31: colors.ads

```
1 package Colors is
2
3   type Color is (Black, Red, Green, Blue, White);
4
5   type Color_Array is array (Integer range <>) of Color;
6
7   procedure Reverse_Color_Array (X : in out Color_Array);
8
9 end Colors;
```

Listing 32: colors.adb

```
1 package body Colors is
2
3   procedure Reverse_Color_Array (X : in out Color_Array) is
4   begin
5     for I in X'First .. (X'Last + X'First) / 2 loop
6       declare
7         Tmp      : Color;
8         X_Left   : Color renames X (I);
9         X_Right  : Color renames X (X'Last + X'First - I);
10      begin
11        Tmp      := X_Left;
12        X_Left   := X_Right;
13        X_Right  := Tmp;
14      end;
15    end loop;
16  end Reverse_Color_Array;
17
18 end Colors;
```

Listing 33: test_non_generic_reverse_colors.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Colors;      use Colors;
```

(continues on next page)

(continued from previous page)

```

3
4 procedure Test_Non_Generic_Reverse_Colors is
5
6   My_Colors : Color_Array (1 .. 5) := (Black, Red, Green, Blue, White);
7
8 begin
9   for C of My_Colors loop
10    Put_Line ("My_Color: " & Color'Image (C));
11  end loop;
12
13  New_Line;
14  Put_Line ("Reversing My_Color...");
15  New_Line;
16  Reverse_Color_Array (My_Colors);
17
18  for C of My_Colors loop
19    Put_Line ("My_Color: " & Color'Image (C));
20  end loop;
21
22 end Test_Non_Generic_Reverse_Colors;

```

Build output

```

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

```

Runtime output

```

My_Color: BLACK
My_Color: RED
My_Color: GREEN
My_Color: BLUE
My_Color: WHITE

Reversing My_Color...

My_Color: WHITE
My_Color: BLUE
My_Color: GREEN
My_Color: RED
My_Color: BLACK

```

The procedure `Reverse_Color_Array` takes an array of colors, starts by swapping the first and last elements of the array, and continues doing that with successive elements until it reaches the middle of array. At that point, the entire array has been reversed, as we see from the output of the test program.

To abstract this procedure, we declare formal types for three components of the algorithm:

- the elements of the array (`Color` type in the example)
- the range used for the array (`Integer` range in the example)
- the actual array type (`Color_Array` type in the example)

This is a generic version of the algorithm:

Listing 34: generic_reverse_array.ads

```

1 generic
2   type T is private;
3   type Index is range <>;
4   type Array_T is array (Index range <>) of T;
5   procedure Generic_Reverse_Array (X : in out Array_T);

```

Listing 35: generic_reverse_array.adb

```

1 procedure Generic_Reverse_Array (X : in out Array_T) is
2 begin
3   for I in X'First .. (X'Last + X'First) / 2 loop
4     declare
5       Tmp      : T;
6       X_Left   : T renames X (I);
7       X_Right  : T renames X (X'Last + X'First - I);
8     begin
9       Tmp      := X_Left;
10      X_Left   := X_Right;
11      X_Right  := Tmp;
12    end;
13  end loop;
14 end Generic_Reverse_Array;

```

Listing 36: colors.ads

```

1 with Generic_Reverse_Array;
2
3 package Colors is
4
5   type Color is (Black, Red, Green, Blue, White);
6
7   type Color_Array is array (Integer range <>) of Color;
8
9   procedure Reverse_Color_Array is new Generic_Reverse_Array
10    (T => Color, Index => Integer, Array_T => Color_Array);
11
12 end Colors;

```

Listing 37: test_reverse_colors.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Colors;      use Colors;
3
4 procedure Test_Reverse_Colors is
5
6   My_Colors : Color_Array (1 .. 5) := (Black, Red, Green, Blue, White);
7
8 begin
9   for C of My_Colors loop
10    Put_Line ("My_Color: " & Color'Image (C));
11  end loop;
12
13  New_Line;
14  Put_Line ("Reversing My_Color...");
15  New_Line;
16  Reverse_Color_Array (My_Colors);
17
18  for C of My_Colors loop

```

(continues on next page)

(continued from previous page)

```

19     Put_Line ("My_Color: " & Color'Image (C));
20     end loop;
21
22 end Test_Reverse_Colors;

```

Build output

```

Compile
  [Ada]      test_reverse_colors.adb
  [Ada]      colors.ads
  [Ada]      generic_reverse_array.adb
Bind
  [gprbind]  test_reverse_colors.bexch
  [Ada]      test_reverse_colors.ali
Link
  [link]     test_reverse_colors.adb

```

Runtime output

```

My_Color: BLACK
My_Color: RED
My_Color: GREEN
My_Color: BLUE
My_Color: WHITE

Reversing My_Color...

My_Color: WHITE
My_Color: BLUE
My_Color: GREEN
My_Color: RED
My_Color: BLACK

```

As mentioned above, we're abstracting three components of the algorithm:

- the T type abstracts the elements of the array
- the Index type abstracts the range used for the array
- the Array_T type abstracts the array type and uses the formal declarations of the T and Index types.

10.12 Example: Test application

In the previous example we've focused only on abstracting the reversing algorithm itself. However, we could have decided to also abstract our small test application. This could be useful if we, for example, decide to test other procedures that change elements of an array.

In order to do this, we again have to choose the elements to abstract. We therefore declare the following formal parameters:

- S: the string containing the array name
- a function Image that converts an element of type T to a string
- a procedure Test that performs some operation on the array

Note that Image and Test are examples of formal subprograms and S is an example of a formal object.

Here is a version of the test application making use of the generic Perform_Test procedure:

Listing 38: generic_reverse_array.ads

```

1 generic
2   type T is private;
3   type Index is range <>;
4   type Array_T is array (Index range <>) of T;
5   procedure Generic_Reverse_Array (X : in out Array_T);

```

Listing 39: generic_reverse_array.adb

```

1 procedure Generic_Reverse_Array (X : in out Array_T) is
2 begin
3   for I in X'First .. (X'Last + X'First) / 2 loop
4     declare
5       Tmp      : T;
6       X_Left   : T renames X (I);
7       X_Right  : T renames X (X'Last + X'First - I);
8     begin
9       Tmp      := X_Left;
10      X_Left   := X_Right;
11      X_Right  := Tmp;
12    end;
13  end loop;
14 end Generic_Reverse_Array;

```

Listing 40: perform_test.ads

```

1 generic
2   type T is private;
3   type Index is range <>;
4   type Array_T is array (Index range <>) of T;
5   S : String;
6   with function Image (E : T) return String is <>;
7   with procedure Test (X : in out Array_T);
8   procedure Perform_Test (X : in out Array_T);

```

Listing 41: perform_test.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Perform_Test (X : in out Array_T) is
4 begin
5   for C of X loop
6     Put_Line (S & " : " & Image (C));
7   end loop;
8
9   New_Line;
10  Put_Line ("Testing " & S & "...");
11  New_Line;
12  Test (X);
13
14  for C of X loop
15    Put_Line (S & " : " & Image (C));
16  end loop;
17 end Perform_Test;

```

Listing 42: colors.ads

```

1 with Generic_Reverse_Array;
2

```

(continues on next page)

(continued from previous page)

```

3 package Colors is
4
5     type Color is (Black, Red, Green, Blue, White);
6
7     type Color_Array is array (Integer range <>) of Color;
8
9     procedure Reverse_Color_Array is new Generic_Reverse_Array
10        (T => Color, Index => Integer, Array_T => Color_Array);
11
12 end Colors;

```

Listing 43: test_reverse_colors.adb

```

1 with Colors;      use Colors;
2 with Perform_Test;
3
4 procedure Test_Reverse_Colors is
5
6     procedure Perform_Test_Reverse_Color_Array is new
7         Perform_Test (T      => Color,
8                       Index  => Integer,
9                       Array_T => Color_Array,
10                      S      => "My_Color",
11                      Image  => Color'Image,
12                      Test   => Reverse_Color_Array);
13
14     My_Colors : Color_Array (1 .. 5) := (Black, Red, Green, Blue, White);
15
16 begin
17     Perform_Test_Reverse_Color_Array (My_Colors);
18 end Test_Reverse_Colors;

```

Build output

```

Compile
  [Ada]      test_reverse_colors.adb
  [Ada]      colors.ads
  [Ada]      perform_test.adb
  [Ada]      generic_reverse_array.adb
Bind
  [gprbind]  test_reverse_colors.bexch
  [Ada]      test_reverse_colors.ali
Link
  [link]     test_reverse_colors.adb

```

Runtime output

```

My_Color: BLACK
My_Color: RED
My_Color: GREEN
My_Color: BLUE
My_Color: WHITE

Testing My_Color...

My_Color: WHITE
My_Color: BLUE
My_Color: GREEN
My_Color: RED
My_Color: BLACK

```

In this example, we create the procedure `Perform_Test_Reverse_Color_Array` as an instance of the generic procedure (`Perform_Test`). Note that:

- For the formal `Image` function, we use the `'Image` attribute of the `Color` type
- For the formal `Test` procedure, we reference the `Reverse_Array` procedure from the package.

EXCEPTIONS

Ada uses exceptions for error handling. Unlike many other languages, Ada speaks about *raising*, not *throwing*, an exception and *handling*, not *catching*, an exception.

11.1 Exception declaration

Ada exceptions are not types, but instead objects, which may be peculiar to you if you're used to the way Java or Python support exceptions. Here's how you declare an exception:

Listing 1: exceptions.ads

```
1 package Exceptions is
2   My_Except : exception;
3   -- Like an object. *NOT* a type !
4 end Exceptions;
```

Even though they're objects, you're going to use each declared exception object as a "kind" or "family" of exceptions. Ada does not require that a subprogram declare every exception it can potentially raise.

11.2 Raising an exception

To raise an exception of our newly declared exception kind, do the following:

Listing 2: main.adb

```
1 with Exceptions; use Exceptions;
2
3 procedure Main is
4 begin
5   raise My_Except;
6   -- Execution of current control flow abandoned; an exception of kind
7   -- "My_Except" will bubble up until it is caught.
8
9   raise My_Except with "My exception message";
10  -- Execution of current control flow abandoned; an exception of
11  -- kind "My_Except" with associated string will bubble up until
12  -- it is caught.
13 end Main;
```

Build output

```
Compile
  [Ada]          main.adb
main.adb:9:04: warning: unreachable code
  [Ada]          exceptions.ads
Bind
  [gprbind]     main.bexch
  [Ada]         main.ali
Link
  [link]        main.adb
```

Runtime output

```
raised EXCEPTIONS.MY_EXCEPT : main.adb:5
```

11.3 Handling an exception

Next, we address how to handle exceptions that were raised by us or libraries that we call. The neat thing in Ada is that you can add an exception handler to any statement block as follows:

Listing 3: open_file.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Exceptions; use Ada.Exceptions;
3
4 procedure Open_File is
5   File : File_Type;
6 begin
7   -- Block (sequence of statements)
8   begin
9     Open (File, In_File, "input.txt");
10    exception
11      when E : Name_Error =>
12        -- ^ Exception to be handled
13        Put ("Cannot open input file : ");
14        Put_Line (Exception_Message (E));
15        raise;
16        -- Reraise current occurrence
17    end;
18 end Open_File;
```

Build output

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

Runtime output

```
Cannot open input file : input.txt: No such file or directory

raised ADA.IO_EXCEPTIONS.NAME_ERROR : input.txt: No such file or directory
```

In the example above, we're using the `Exception_Message` function from the `Ada.Exceptions` package. This function returns the message associated with the exception as a string.

You don't need to introduce a block just to handle an exception: you can add it to the statements block of your current subprogram:

Listing 4: open_file.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Exceptions; use Ada.Exceptions;
3
4 procedure Open_File is
5   File : File_Type;
6 begin
7   Open (File, In_File, "input.txt");
8   -- Exception block can be added to any block
9 exception
10  when Name_Error =>
11    Put ("Cannot open input file");
12 end Open_File;
```

Build output

```

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

Runtime output

```
Cannot open input file
```

Attention

Exception handlers have an important restriction that you need to be careful about: Exceptions raised in the declarative section are not caught by the handlers of that block. So for example, in the following code, the exception will not be caught.

Listing 5: be_careful.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Exceptions; use Ada.Exceptions;
3
4 procedure Be_Careful is
5   function Dangerous return Integer is
6   begin
7     raise Constraint_Error;
8     return 42;
9   end Dangerous;
10
11 begin
12   declare
13     A : Integer := Dangerous;
14   begin
15     Put_Line (Integer'Image (A));
16   exception
17     when Constraint_Error => Put_Line ("error!");
18   end;
19 end Be_Careful;
```

Build output

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

Runtime output

```
raised CONSTRAINT_ERROR : be_careful.adb:7 explicit raise
```

This is also the case for the top-level exception block that is part of the current subprogram.

11.4 Predefined exceptions

Ada has a very small number of predefined exceptions:

- `Constraint_Error` is the main one you might see. It's raised:
 - When bounds don't match or, in general, any violation of constraints.
 - In case of overflow
 - In case of null dereferences
 - In case of division by 0
- `Program_Error` might appear, but probably less often. It's raised in more arcane situations, such as for order of elaboration issues and some cases of detectable erroneous execution.
- `Storage_Error` will happen because of memory issues, such as:
 - Not enough memory (allocator)
 - Not enough stack
- `Tasking_Error` will happen with task related errors, such as any error happening during task activation.

You should not reuse predefined exceptions. If you do then, it won't be obvious when one is raised that it is because something went wrong in a built-in language operation.

TASKING

Tasks and protected objects allow the implementation of concurrency in Ada. The following sections explain these concepts in more details.

12.1 Tasks

A task can be thought as an application that runs *concurrently* with the main application. In other programming languages, a task can be called a [thread](#)¹⁷, and tasking can be called [multithreading](#)¹⁸.

Tasks may synchronize with the main application but may also process information completely independent from the main application. Here we show how this is accomplished.

12.1.1 Simple task

Tasks are declared using the keyword `task`. The task implementation is specified in a `task body` block. For example:

Listing 1: show_simple_task.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Simple_Task is
4   task T;
5
6   task body T is
7   begin
8     Put_Line ("In task T");
9   end T;
10 begin
11   Put_Line ("In main");
12 end Show_Simple_Task;
```

Build output

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

Runtime output

¹⁷ [https://en.wikipedia.org/wiki/Thread_\(computing\)](https://en.wikipedia.org/wiki/Thread_(computing))

¹⁸ [https://en.wikipedia.org/wiki/Thread_\(computing\)#Multithreading](https://en.wikipedia.org/wiki/Thread_(computing)#Multithreading)

```
In task T
In main
```

Here, we're declaring and implementing the task T. As soon as the main application starts, task T starts automatically — it's not necessary to manually start this task. By running the application above, we can see that both calls to `Put_Line` are performed.

Note that:

- The main application is itself a task (the main task).
 - In this example, the subprogram `Show_Simple_Task` is the main task of the application.
- Task T is a subtask.
 - Each subtask has a master task.
 - Therefore the main task is also the master task of task T.
- The number of tasks is not limited to one: we could include a task T2 in the example above.
 - This task also starts automatically and runs *concurrently* with both task T and the main task. For example:

Listing 2: `show_simple_tasks.adb`

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Simple_Tasks is
4   task T;
5   task T2;
6
7   task body T is
8   begin
9     Put_Line ("In task T");
10  end T;
11
12  task body T2 is
13  begin
14    Put_Line ("In task T2");
15  end T2;
16
17  begin
18    Put_Line ("In main");
19  end Show_Simple_Tasks;
```

Build output

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

Runtime output

```
In task T
In task T2
In main
```

12.1.2 Simple synchronization

As we've just seen, as soon as the main task starts, its subtasks also start automatically. The main task continues its processing until it has nothing more to do. At that point, however, it will not terminate. Instead, the task waits until its subtasks have finished before it allows itself to terminate. In other words, this waiting process provides synchronization between the main task and its subtasks. After this synchronization, the main task will terminate. For example:

Listing 3: show_simple_sync.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Simple_Sync is
4   task T;
5   task body T is
6     begin
7       for I in 1 .. 10 loop
8         Put_Line ("hello");
9       end loop;
10    end T;
11  begin
12    null;
13    -- Will wait here until all tasks have terminated
14  end Show_Simple_Sync;
```

Build output

```

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

Runtime output

```

hello
hello
hello
hello
hello
hello
hello
hello
hello
hello
hello
```

The same mechanism is used for other subprograms that contain subtasks: the subprogram's master task will wait for its subtasks to finish. So this mechanism is not limited to the main application and also applies to any subprogram called by the main application or its subprograms.

Synchronization also occurs if we move the task to a separate package. In the example below, we declare a task T in the package `Simple_Sync_Pkg`.

Listing 4: simple_sync_pkg.ads

```

1 package Simple_Sync_Pkg is
2   task T;
3 end Simple_Sync_Pkg;
```

This is the corresponding package body:

Listing 5: simple_sync_pkg.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body Simple_Sync_Pkg is
4   task body T is
5     begin
6       for I in 1 .. 10 loop
7         Put_Line ("hello");
8       end loop;
9     end T;
10 end Simple_Sync_Pkg;
```

Because the package is with'ed by the main procedure, the task T defined in the package is part of the main task. For example:

Listing 6: test_simple_sync_pkg.adb

```
1 with Simple_Sync_Pkg;
2
3 procedure Test_Simple_Sync_Pkg is
4   begin
5     null;
6     -- Will wait here until all tasks have terminated
7   end Test_Simple_Sync_Pkg;
```

Build output

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

Runtime output

```
hello
hello
hello
hello
hello
hello
hello
hello
hello
hello
hello
```

Again, as soon as the main task reaches its end, it synchronizes with task T from Simple_Sync_Pkg before terminating.

12.1.3 Delay

We can introduce a delay by using the keyword `delay`. This puts the task to sleep for the length of time (in seconds) specified in the delay statement. For example:

Listing 7: show_delay.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Delay is
4
5     task T;
6
7     task body T is
8     begin
9         for I in 1 .. 5 loop
10            Put_Line ("hello from task T");
11            delay 1.0;
12            --      ^ Wait 1.0 seconds
13        end loop;
14    end T;
15 begin
16     delay 1.5;
17     Put_Line ("hello from main");
18 end Show_Delay;

```

Build output

```

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

```

Runtime output

```

hello from task T
hello from task T
hello from main
hello from task T
hello from task T
hello from task T

```

In this example, we're making the task T wait one second after each time it displays the "hello" message. In addition, the main task is waiting 1.5 seconds before displaying its own "hello" message

12.1.4 Synchronization: rendez-vous

The only type of synchronization we've seen so far is the one that happens automatically at the end of the main task. You can also define custom synchronization points using the keyword `entry`. An *entry* can be viewed as a special kind of subprogram, which is called by the master task using a similar syntax, as we will see later.

In the task definition, you define which part of the task will accept the entries by using the keyword `accept`. A task proceeds until it reaches an `accept` statement and then waits for the master task to synchronize with it. Specifically,

- The subtask waits at that point (in the `accept` statement), ready to accept a call to the corresponding entry from the master task.

- The master task calls the task entry, in a manner similar to a procedure call, to synchronize with the subtask.

This synchronization between tasks is called *rendez-vous*. Let's see an example:

Listing 8: show_rendezvous.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Rendezvous is
4
5     task T is
6         entry Start;
7     end T;
8
9     task body T is
10        begin
11            accept Start; -- Waiting for somebody to call the entry
12            Put_Line ("In T");
13        end T;
14
15    begin
16        Put_Line ("In Main");
17        T.Start; -- Calling T's entry
18    end Show_Rendezvous;
```

Build output

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

Runtime output

```
In Main
In T
```

In this example, we declare an entry `Start` for task `T`. In the task body, we implement this entry using `accept Start`. When task `T` reaches this point, it waits for the master task. This synchronization occurs in the `T.Start` statement. After the synchronization completes, the main task and task `T` again run concurrently until they synchronize one final time when the main task finishes.

An entry may be used to perform more than a simple task synchronization: it also may perform multiple statements during the time both tasks are synchronized. We do this with a `do ... end` block. For the previous example, we would simply write `accept Start do <statements>; end;`. We use this kind of block in the next example.

12.1.5 Select loop

There's no limit to the number of times an entry can be accepted. We could even create an infinite loop in the task and accept calls to the same entry over and over again. An infinite loop, however, prevents the subtask from finishing, so it blocks the master task when it reaches the end of its processing. Therefore, a loop containing accept statements in a task body is normally used in conjunction with a `select ... or terminate` statement. In simple terms, this statement allows the master task to automatically terminate the subtask when the master task finishes. For example:

Listing 9: show_rendezvous_loop.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Rendezvous_Loop is
4
5     task T is
6         entry Reset;
7         entry Increment;
8     end T;
9
10    task body T is
11        Cnt : Integer := 0;
12    begin
13        loop
14            select
15                accept Reset do
16                    Cnt := 0;
17                end Reset;
18                Put_Line ("Reset");
19            or
20                accept Increment do
21                    Cnt := Cnt + 1;
22                end Increment;
23                Put_Line ("In T's loop (" & Integer'Image (Cnt) & ")");
24            or
25                terminate;
26            end select;
27        end loop;
28    end T;
29
30 begin
31     Put_Line ("In Main");
32
33     for I in 1 .. 4 loop
34         T.Increment; -- Calling T's entry multiple times
35     end loop;
36
37     T.Reset;
38     for I in 1 .. 4 loop
39         T.Increment; -- Calling T's entry multiple times
40     end loop;
41
42 end Show_Rendezvous_Loop;
```

Build output

```

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

(continues on next page)

(continued from previous page)

[link]	show_rendezvous_loop.adb
--------	--------------------------

Runtime output

```
In Main
In T's loop ( 1)
In T's loop ( 2)
In T's loop ( 3)
In T's loop ( 4)
Reset
In T's loop ( 1)
In T's loop ( 2)
In T's loop ( 3)
In T's loop ( 4)
```

In this example, the task body implements an infinite loop that accepts calls to the `Reset` and `Increment` entry. We make the following observations:

- The `accept E do ... end` block is used to increment a counter.
 - As long as task T is performing the `do ... end` block, the main task waits for the block to complete.
- The main task is calling the `Increment` entry multiple times in the loop from 1 .. 4. It is also calling the `Reset` entry before and the loop.
 - Because task T contains an infinite loop, it always accepts calls to the `Reset` and `Increment` entries.
 - When the main task finishes, it checks the status of the T task. Even though task T could accept new calls to the `Reset` or `Increment` entries, the master task is allowed to terminate task T due to the `or terminate` part of the `select` statement.

12.1.6 Cycling tasks

In a previous example, we saw how to delay a task a specified time by using the `delay` keyword. However, using delay statements in a loop is not enough to guarantee regular intervals between those delay statements. For example, we may have a call to a computationally intensive procedure between executions of successive delay statements:

```
while True loop
  delay 1.0;
  -- ^ Wait 1.0 seconds
  Computational_Intensive_App;
end loop;
```

In this case, we can't guarantee that exactly 10 seconds have elapsed after 10 calls to the `delay` statement because a time drift may be introduced by the `Computational_Intensive_App` procedure. In many cases, this time drift is not relevant, so using the `delay` keyword is good enough.

However, there are situations where a time drift isn't acceptable. In those cases, we need to use the `delay until` statement, which accepts a precise time for the end of the delay, allowing us to define a regular interval. This is useful, for example, in real-time applications.

We will soon see an example of how this time drift may be introduced and how the `delay until` statement circumvents the problem. But before we do that, we look at a package containing a procedure allowing us to measure the elapsed time (`Show_Elapsed_Time`) and a dummy `Computational_Intensive_App` procedure which is simulated by using a simple delay. This is the package specification:

Listing 10: delay_aux_pkg.ads

```

1 with Ada.Real_Time; use Ada.Real_Time;
2
3 package Delay_Aux_Pkg is
4
5     function Get_Start_Time return Time
6         with Inline;
7
8     procedure Show_Elapsed_Time
9         with Inline;
10
11    procedure Computational_Intensive_App;
12 private
13     Start_Time    : Time := Clock;
14
15     function Get_Start_Time return Time is (Start_Time);
16
17 end Delay_Aux_Pkg;

```

And this is the package body:

Listing 11: delay_aux_pkg.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body Delay_Aux_Pkg is
4
5     procedure Show_Elapsed_Time is
6         Now_Time    : Time;
7         Elapsed_Time : Time_Span;
8     begin
9         Now_Time    := Clock;
10        Elapsed_Time := Now_Time - Start_Time;
11        Put_Line ("Elapsed time "
12                & Duration'Image (To_Duration (Elapsed_Time))
13                & " seconds");
14    end Show_Elapsed_Time;
15
16    procedure Computational_Intensive_App is
17    begin
18        delay 0.5;
19    end Computational_Intensive_App;
20
21 end Delay_Aux_Pkg;

```

Using this auxiliary package, we're now ready to write our time-drifting application:

Listing 12: show_time_task.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Real_Time; use Ada.Real_Time;
3
4 with Delay_Aux_Pkg;
5
6 procedure Show_Time_Task is
7     package Aux renames Delay_Aux_Pkg;
8
9     task T;
10
11    task body T is
12        Cnt : Integer := 1;

```

(continues on next page)

(continued from previous page)

```

13   begin
14     for I in 1 .. 5 loop
15       delay 1.0;
16
17       Aux.Show_Elapsed_Time;
18       Aux.Computational_Intensive_App;
19
20       Put_Line ("Cycle # " & Integer'Image (Cnt));
21       Cnt := Cnt + 1;
22     end loop;
23     Put_Line ("Finished time-drifting loop");
24   end T;
25
26   begin
27     null;
28   end Show_Time_Task;

```

Build output

```

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

```

Runtime output

```

Elapsed time 1.000264399 seconds
Cycle # 1
Elapsed time 2.500680928 seconds
Cycle # 2
Elapsed time 4.001080091 seconds
Cycle # 3
Elapsed time 5.501382926 seconds
Cycle # 4
Elapsed time 7.001825218 seconds
Cycle # 5
Finished time-drifting loop

```

We can see by running the application that we already have a time difference of about four seconds after three iterations of the loop due to the drift introduced by `Computational_Intensive_App`. Using the `delay until` statement, however, we're able to avoid this time drift and have a regular interval of exactly one second:

Listing 13: show_time_task.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2  with Ada.Real_Time; use Ada.Real_Time;
3
4  with Delay_Aux_Pkg;
5
6  procedure Show_Time_Task is
7    package Aux renames Delay_Aux_Pkg;
8
9    task T;
10
11   task body T is
12     Cycle : constant Time_Span := Milliseconds (1000);

```

(continues on next page)

(continued from previous page)

```

13     Next : Time := Aux.Get_Start_Time + Cycle;
14
15     Cnt  : Integer := 1;
16     begin
17         for I in 1 .. 5 loop
18             delay until Next;
19
20             Aux.Show_Elapsed_Time;
21             Aux.Computational_Intensive_App;
22
23             -- Calculate next execution time using a
24             -- cycle of one second
25             Next := Next + Cycle;
26
27             Put_Line ("Cycle # " & Integer'Image (Cnt));
28             Cnt := Cnt + 1;
29         end loop;
30         Put_Line ("Finished cycling");
31     end T;
32
33 begin
34     null;
35 end Show_Time_Task;

```

Build output

```

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

```

Runtime output

```

Elapsed time 1.000158707 seconds
Cycle # 1
Elapsed time 2.000149219 seconds
Cycle # 2
Elapsed time 3.000103642 seconds
Cycle # 3
Elapsed time 4.000103182 seconds
Cycle # 4
Elapsed time 5.000170341 seconds
Cycle # 5
Finished cycling

```

Now, as we can see by running the application, the `delay until` statement ensures that the `Computational_Intensive_App` doesn't disturb the regular interval of one second between iterations.

12.2 Protected objects

When multiple tasks are accessing shared data, corruption of that data may occur. For example, data may be inconsistent if one task overwrites parts of the information that's being read by another task at the same time. In order to avoid these kinds of problems and ensure information is accessed in a coordinated way, we use *protected objects*.

Protected objects encapsulate data and provide access to that data by means of *protected operations*, which may be subprograms or protected entries. Using protected objects ensures that data is not corrupted by race conditions or other simultaneous access.

Important

Protected objects can be implemented using Ada tasks. In fact, this was the *only* possible way of implementing them in Ada 83 (the first version of the Ada language). However, the use of protected objects is much simpler than using similar mechanisms implemented using only tasks. Therefore, you should use protected objects when your main goal is only to protect data.

12.2.1 Simple object

You declare a protected object with the `protected` keyword. The syntax is similar to that used for packages: you can declare operations (e.g., procedures and functions) in the public part and data in the private part. The corresponding implementation of the operations is included in the `protected body` of the object. For example:

Listing 14: show_protected_objects.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Protected_Objects is
4
5     protected Obj is
6         -- Operations go here (only subprograms)
7         procedure Set (V : Integer);
8         function Get return Integer;
9     private
10        -- Data goes here
11        Local : Integer := 0;
12    end Obj;
13
14    protected body Obj is
15        -- procedures can modify the data
16        procedure Set (V : Integer) is
17            begin
18                Local := V;
19            end Set;
20
21        -- functions cannot modify the data
22        function Get return Integer is
23            begin
24                return Local;
25            end Get;
26    end Obj;
27
28 begin
29     Obj.Set (5);
30     Put_Line ("Number is: " & Integer'Image (Obj.Get));
31 end Show_Protected_Objects;
```

Build output

```

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

```

Runtime output

```
Number is: 5
```

In this example, we define two operations for `Obj`: `Set` and `Get`. The implementation of these operations is in the `Obj` body. The syntax used for writing these operations is the same as that for normal procedures and functions. The implementation of protected objects is straightforward — we simply access and update `Local` in these subprograms. To call these operations in the main application, we use prefixed notation, e.g., `Obj.Get`.

12.2.2 Entries

In addition to protected procedures and functions, you can also define protected entry points. Do this using the `entry` keyword. Protected entry points allow you to define barriers using the `when` keyword. Barriers are conditions that must be fulfilled before the entry can start performing its actual processing — we speak of *releasing* the barrier when the condition is fulfilled.

The previous example used procedures and functions to define operations on the protected objects. However, doing so permits reading protected information (via `Obj.Get`) before it's set (via `Obj.Set`). To allow that to be a defined operation, we specified a default value (0). Instead, by rewriting `Obj.Get` using an *entry* instead of a function, we implement a barrier, ensuring no task can read the information before it's been set.

The following example implements the barrier for the `Obj.Get` operation. It also contains two concurrent subprograms (main task and task T) that try to access the protected object.

Listing 15: `show_protected_objects_entries.adb`

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Show_Protected_Objects_Entries is
4
5     protected Obj is
6       procedure Set (V : Integer);
7       entry Get (V : out Integer);
8     private
9       Local  : Integer;
10      Is_Set  : Boolean := False;
11    end Obj;
12
13    protected body Obj is
14      procedure Set (V : Integer) is
15        begin
16          Local := V;
17          Is_Set := True;
18        end Set;
19
20      entry Get (V : out Integer)
21        when Is_Set is
22        -- Entry is blocked until the condition is true.

```

(continues on next page)

(continued from previous page)

```

23     -- The barrier is evaluated at call of entries and at exits of
24     -- procedures and entries.
25     -- The calling task sleeps until the barrier is released.
26     begin
27         V := Local;
28         Is_Set := False;
29     end Get;
30 end Obj;
31
32 N : Integer := 0;
33
34 task T;
35
36 task body T is
37 begin
38     Put_Line ("Task T will delay for 4 seconds...");
39     delay 4.0;
40     Put_Line ("Task T will set Obj...");
41     Obj.Set (5);
42     Put_Line ("Task T has just set Obj...");
43 end T;
44 begin
45     Put_Line ("Main application will get Obj...");
46     Obj.Get (N);
47     Put_Line ("Main application has just retrieved Obj...");
48     Put_Line ("Number is: " & Integer'Image (N));
49
50 end Show_Protected_Objects_Entries;

```

Build output

```

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

```

Runtime output

```

Task T will delay for 4 seconds...
Main application will get Obj...
Task T will set Obj...
Task T has just set Obj...
Main application has just retrieved Obj...
Number is: 5

```

As we see by running it, the main application waits until the protected object is set (by the call to `Obj.Set` in task T) before it reads the information (via `Obj.Get`). Because a 4-second delay has been added in task T, the main application is also delayed by 4 seconds. Only after this delay does task T set the object and release the barrier in `Obj.Get` so that the main application can then resume processing (after the information is retrieved from the protected object).

12.3 Task and protected types

In the previous examples, we defined single tasks and protected objects. We can, however, generalize tasks and protected objects using type definitions. This allows us, for example, to create multiple tasks based on just a single task type.

12.3.1 Task types

A task type is a generalization of a task. The declaration is similar to simple tasks: you replace `task` with `task type`. The difference between simple tasks and task types is that task types don't create actual tasks that automatically start. Instead, a task declaration is needed. This is exactly the way normal variables and types work: objects are only created by variable definitions, not type definitions.

To illustrate this, we repeat our first example:

Listing 16: show_simple_task.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Simple_Task is
4   task T;
5
6   task body T is
7   begin
8     Put_Line ("In task T");
9   end T;
10 begin
11   Put_Line ("In main");
12 end Show_Simple_Task;
```

Build output

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

Runtime output

```
In task T
In main
```

We now rewrite it by replacing `task T` with `task type TT`. We declare a task (`A_Task`) based on the task type `TT` after its definition:

Listing 17: show_simple_task_type.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Simple_Task_Type is
4   task type TT;
5
6   task body TT is
7   begin
8     Put_Line ("In task type TT");
```

(continues on next page)

(continued from previous page)

```

9   end TT;
10
11  A_Task : TT;
12  begin
13    Put_Line ("In main");
14  end Show_Simple_Task_Type;

```

Build output

```

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

```

Runtime output

```

In task type TT
In main

```

We can extend this example and create an array of tasks. Since we're using the same syntax as for variable declarations, we use a similar syntax for task types: `array (<>) of Task_Type`. Also, we can pass information to the individual tasks by defining a `Start` entry. Here's the updated example:

Listing 18: show_task_type_array.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  procedure Show_Task_Type_Array is
4    task type TT is
5      entry Start (N : Integer);
6    end TT;
7
8    task body TT is
9      Task_N : Integer;
10   begin
11     accept Start (N : Integer) do
12       Task_N := N;
13     end Start;
14     Put_Line ("In task T: " & Integer'Image (Task_N));
15   end TT;
16
17   My_Tasks : array (1 .. 5) of TT;
18  begin
19    Put_Line ("In main");
20
21    for I in My_Tasks'Range loop
22      My_Tasks (I).Start (I);
23    end loop;
24  end Show_Task_Type_Array;

```

Build output

```

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

```

(continues on next page)

(continued from previous page)

```
Link
  [link]          show_task_type_array.adb
```

Runtime output

```
In main
In task T: 1
In task T: 2
In task T: 3
In task T: 4
In task T: 5
```

In this example, we're declaring five tasks in the array `My_Tasks`. We pass the array index to the individual tasks in the entry point (`Start`). After the synchronization between the individual subtasks and the main task, each subtask calls `Put_Line` concurrently.

12.3.2 Protected types

A protected type is a generalization of a protected object. The declaration is similar to that for protected objects: you replace `protected` with `protected type`. Like task types, protected types require an object declaration to create actual objects. Again, this is similar to variable declarations and allows for creating arrays (or other composite objects) of protected objects.

We can reuse a previous example and rewrite it to use a protected type:

Listing 19: `show_protected_object_type.adb`

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Protected_Object_Type is
4
5   protected type Obj_Type is
6     procedure Set (V : Integer);
7     function Get return Integer;
8   private
9     Local : Integer := 0;
10  end Obj_Type;
11
12  protected body Obj_Type is
13    procedure Set (V : Integer) is
14      begin
15        Local := V;
16      end Set;
17
18    function Get return Integer is
19      begin
20        return Local;
21      end Get;
22  end Obj_Type;
23
24  Obj : Obj_Type;
25  begin
26    Obj.Set (5);
27    Put_Line ("Number is: " & Integer'Image (Obj.Get));
28 end Show_Protected_Object_Type;
```

Build output

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

Runtime output

```
Number is: 5
```

In this example, instead of directly defining the protected object `Obj`, we first define a protected type `Obj_Type` and then declare `Obj` as an object of that protected type. Note that the main application hasn't changed: we still use `Obj.Set` and `Obj.Get` to access the protected object, just like in the original example.

DESIGN BY CONTRACTS

Contracts are used in programming to codify expectations. Parameter modes of a subprogram can be viewed as a simple form of contracts. When the specification of subprogram `Op` declares a parameter using `in` mode, the caller of `Op` knows that the `in` argument won't be changed by `Op`. In other words, the caller expects that `Op` doesn't modify the argument it's providing, but just reads the information stored in the argument. Constraints and subtypes are other examples of contracts. In general, these specifications improve the consistency of the application.

Design-by-contract programming refers to techniques that include pre- and postconditions, subtype predicates, and type invariants. We study those topics in this chapter.

13.1 Pre- and postconditions

Pre- and postconditions provide expectations regarding input and output parameters of subprograms and return value of functions. If we say that certain requirements must be met before calling a subprogram `Op`, those are preconditions. Similarly, if certain requirements must be met after a call to the subprogram `Op`, those are postconditions. We can think of preconditions and postconditions as promises between the subprogram caller and the callee: a precondition is a promise from the caller to the callee, and a postcondition is a promise in the other direction.

Pre- and postconditions are specified using an aspect clause in the subprogram declaration. A `with Pre => <condition>` clause specifies a precondition and a `with Post => <condition>` clause specifies a postcondition.

The following code shows an example of preconditions:

Listing 1: show_simple_precondition.adb

```
1 procedure Show_Simple_Precondition is
2
3     procedure DB_Entry (Name : String; Age : Natural)
4         with Pre => Name'Length > 0
5     is
6     begin
7         -- Missing implementation
8         null;
9     end DB_Entry;
10 begin
11     DB_Entry ("John", 30);
12     DB_Entry ("", 21); -- Precondition will fail!
13 end Show_Simple_Precondition;
```

Build output

```
Compile
  [Ada]          show_simple_precondition.adb
```

(continues on next page)

(continued from previous page)

```

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

```

Runtime output

```

raised SYSTEM.ASSERTIONS.ASSERT_FAILURE : failed precondition from show_simple_
↳precondition.adb:4

```

In this example, we want to prevent the name field in our database from containing an empty string. We implement this requirement by using a precondition requiring that the length of the string used for the Name parameter of the DB_Entry procedure is greater than zero. If the DB_Entry procedure is called with an empty string for the Name parameter, the call will fail because the precondition is not met.

In the GNAT toolchain

GNAT handles pre- and postconditions by generating runtime assertions for them. By default, however, assertions aren't enabled. Therefore, in order to check pre- and postconditions at runtime, you need to enable assertions by using the *-gnata* switch.

Before we get to our next example, let's briefly discuss quantified expressions, which are quite useful in concisely writing pre- and postconditions. Quantified expressions return a Boolean value indicating whether elements of an array or container match the expected condition. They have the form: (for all I in A'Range => <condition on A(I)>), where A is an array and I is an index. Quantified expressions using for all check whether the condition is true for every element. For example:

```
(for all I in A'Range => A (I) = 0)
```

This quantified expression is only true when all elements of the array A have a value of zero.

Another kind of quantified expressions uses for some. The form looks similar: (for some I in A'Range => <condition on A(I)>). However, in this case the qualified expression tests whether the condition is true only on *some* elements (hence the name) instead of all elements.

We illustrate postconditions using the following example:

Listing 2: show_simple_postcondition.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Simple_Postcondition is
4
5   type Int_8 is range -2 ** 7 .. 2 ** 7 - 1;
6
7   type Int_8_Array is array (Integer range <>) of Int_8;
8
9   function Square (A : Int_8) return Int_8 is
10    (A * A)
11    with Post => (if abs A in 0 | 1 then Square'Result = abs A
12                else Square'Result > A);
13
14   procedure Square (A : in out Int_8_Array)
15     with Post => (for all I in A'Range =>
16                  A (I) = A'Old (I) * A'Old (I))

```

(continues on next page)

(continued from previous page)

```

17  is
18  begin
19      for V of A loop
20          V := Square (V);
21      end loop;
22  end Square;
23
24  V : Int_8_Array := (-2, -1, 0, 1, 10, 11);
25  begin
26      for E of V loop
27          Put_Line ("Original: " & Int_8'Image (E));
28      end loop;
29      New_Line;
30
31      Square (V);
32      for E of V loop
33          Put_Line ("Square:  " & Int_8'Image (E));
34      end loop;
35  end Show_Simple_Postcondition;

```

Build output

```

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

```

Runtime output

```

Original: -2
Original: -1
Original: 0
Original: 1
Original: 10
Original: 11

Square: 4
Square: 1
Square: 0
Square: 1
Square: 100
Square: 121

```

We declare a signed 8-bit type `Int_8` and an array of that type (`Int_8_Array`). We want to ensure each element of the array is squared after calling the procedure `Square` for an object of the `Int_8_Array` type. We do this with a postcondition using a `for all` expression. This postcondition also uses the `'Old` attribute to refer to the original value of the parameter (before the call).

We also want to ensure that the result of calls to the `Square` function for the `Int_8` type are greater than the input to that call. To do that, we write a postcondition using the `'Result` attribute of the function and comparing it to the input value.

We can use both pre- and postconditions in the declaration of a single subprogram. For example:

Listing 3: `show_simple_contract.adb`

```

1  with Ada.Text_IO; use Ada.Text_IO;
2

```

(continues on next page)

(continued from previous page)

```

3 procedure Show_Simple_Contract is
4
5   type Int_8 is range -2 ** 7 .. 2 ** 7 - 1;
6
7   function Square (A : Int_8) return Int_8 is
8     (A * A)
9     with
10      Pre  => (Integer'Size >= Int_8'Size * 2 and
11              Integer (A) * Integer (A) < Integer (Int_8'Last)),
12      Post => (if abs A in 0 | 1 then Square'Result = abs A
13              else Square'Result > A);
14
15   V : Int_8;
16 begin
17   V := Square (11);
18   Put_Line ("Square of 11 is " & Int_8'Image (V));
19
20   V := Square (12); -- Precondition will fail...
21   Put_Line ("Square of 12 is " & Int_8'Image (V));
22 end Show_Simple_Contract;

```

Build output

```

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

```

Runtime output

```

Square of 11 is 121

raised SYSTEM.ASSERTIONS.ASSERT_FAILURE : failed precondition from show_simple_
↳ contract.adb:10

```

In this example, we want to ensure that the input value of calls to the Square function for the Int_8 type won't cause overflow in that function. We do this by converting the input value to the Integer type, which is used for the temporary calculation, and check if the result is in the appropriate range for the Int_8 type. We have the same postcondition in this example as in the previous one.

13.2 Predicates

Predicates specify expectations regarding types. They're similar to pre- and postconditions, but apply to types instead of subprograms. Their conditions are checked for each object of a given type, which allows verifying that an object of type T is conformant to the requirements of its type.

There are two kinds of predicates: static and dynamic. In simple terms, static predicates are used to check objects at compile-time, while dynamic predicates are used for checks at run time. Normally, static predicates are used for scalar types and dynamic predicates for the more complex types.

Static and dynamic predicates are specified using the following clauses, respectively:

- with Static_Predicate => <property>
- with Dynamic_Predicate => <property>

Let's use the following example to illustrate dynamic predicates:

Listing 4: show_dynamic_predicate_courses.adb

```

1 with Ada.Strings.Unbounded; use Ada.Strings.Unbounded;
2 with Ada.Calendar;         use Ada.Calendar;
3 with Ada.Containers.Vectors;
4
5 procedure Show_Dynamic_Predicate_Courses is
6
7   package Courses is
8     type Course_Container is private;
9
10    type Course is record
11      Name      : Unbounded_String;
12      Start_Date : Time;
13      End_Date  : Time;
14    end record
15    with Dynamic_Predicate => Course.Start_Date <= Course.End_Date;
16
17    procedure Add (CC : in out Course_Container; C : Course);
18  private
19    package Course_Vectors is new Ada.Containers.Vectors
20      (Index_Type  => Natural,
21       Element_Type => Course);
22
23    type Course_Container is record
24      V : Course_Vectors.Vector;
25    end record;
26  end Courses;
27
28  package body Courses is
29    procedure Add (CC : in out Course_Container; C : Course) is
30      begin
31        CC.V.Append (C);
32      end Add;
33  end Courses;
34
35  use Courses;
36
37  CC : Course_Container;
38  begin
39    Add (CC,
40        Course'(
41          Name      => To_Unbounded_String ("Intro to Photography"),
42          Start_Date => Time_Of (2018, 5, 1),
43          End_Date  => Time_Of (2018, 5, 10)));
44
45    -- This should trigger an error in the dynamic predicate check
46    Add (CC,
47        Course'(
48          Name      => To_Unbounded_String ("Intro to Video Recording"),
49          Start_Date => Time_Of (2019, 5, 1),
50          End_Date  => Time_Of (2018, 5, 10)));
51
52  end Show_Dynamic_Predicate_Courses;

```

Build output

```

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

```

(continues on next page)

(continued from previous page)

```
Link
  [link]          show_dynamic_predicate_courses.adb
```

Runtime output

```
raised SYSTEM.ASSERTIONS.ASSERT_FAILURE : Dynamic_Predicate failed at show_dynamic_
↳predicate_courses.adb:47
```

In this example, the package `Courses` defines a type `Course` and a type `Course_Container`, an object of which contains all courses. We want to ensure that the dates of each course are consistent, specifically that the start date is no later than the end date. To enforce this rule, we declare a dynamic predicate for the `Course` type that performs the check for each object. The predicate uses the type name where a variable of that type would normally be used: this is a reference to the instance of the object being tested.

Note that the example above makes use of unbounded strings and dates. Both types are available in Ada's standard library. Please refer to the following sections for more information about:

- the unbounded string type (`Unbounded_String`): *Unbounded Strings* (page 230) section;
- dates and times: *Dates & Times* (page 213) section.

Static predicates, as mentioned above, are mostly used for scalar types and checked during compilation. They're particularly useful for representing non-contiguous elements of an enumeration. A classic example is a list of week days:

```
type Week is (Mon, Tue, Wed, Thu, Fri, Sat, Sun);
```

We can easily create a sub-list of work days in the week by specifying a subtype with a range based on `Week`. For example:

```
subtype Work_Week is Week range Mon .. Fri;
```

Ranges in Ada can only be specified as contiguous lists: they don't allow us to pick specific days. However, we may want to create a list containing just the first, middle and last day of the work week. To do that, we use a static predicate:

```
subtype Check_Days is Work_Week
  with Static_Predicate => Check_Days in Mon | Wed | Fri;
```

Let's look at a complete example:

Listing 5: show_predicates.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Predicates is
4
5   type Week is (Mon, Tue, Wed, Thu, Fri, Sat, Sun);
6
7   subtype Work_Week is Week range Mon .. Fri;
8
9   subtype Test_Days is Work_Week
10    with Static_Predicate => Test_Days in Mon | Wed | Fri;
11
12   type Tests_Week is array (Week) of Natural
13    with Dynamic_Predicate =>
14     (for all I in Tests_Week'Range =>
15      (case I is
16       when Test_Days => Tests_Week (I) > 0,
```

(continues on next page)

(continued from previous page)

```

17         when others => Tests_Week (I) = 0));
18
19     Num_Tests : Tests_Week :=
20         (Mon => 3, Tue => 0,
21          Wed => 4, Thu => 0,
22          Fri => 2, Sat => 0, Sun => 0);
23
24     procedure Display_Tests (N : Tests_Week) is
25     begin
26         for I in Test_Days loop
27             Put_Line ("# tests on " & Test_Days'Image (I)
28                      & " => " & Integer'Image (N (I)));
29         end loop;
30     end Display_Tests;
31
32 begin
33     Display_Tests (Num_Tests);
34
35     -- Assigning non-conformant values to individual elements of
36     -- the Tests_Week type does not trigger a predicate check:
37     Num_Tests (Tue) := 2;
38
39     -- However, assignments with the "complete" Tests_Week type
40     -- trigger a predicate check. For example:
41     --
42     -- Num_Tests := (others => 0);
43
44     -- Also, calling any subprogram with parameters of Tests_Week
45     -- type triggers a predicate check.
46     -- Therefore, the following line will fail:
47     Display_Tests (Num_Tests);
48 end Show_Predicates;
```

Build output

```

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

Runtime output

```

# tests on MON => 3
# tests on WED => 4
# tests on FRI => 2

raised SYSTEM.ASSERTIONS.ASSERT_FAILURE : Dynamic_Predicate failed at show_
↳ predicates.adb:47
```

Here we have an application that wants to perform tests only on three days of the work week. These days are specified in the `Test_Days` subtype. We want to track the number of tests that occur each day. We declare the type `Tests_Week` as an array, an object of which will contain the number of tests done each day. According to our requirements, these tests should happen only in the aforementioned three days; on other days, no tests should be performed. This requirement is implemented with a dynamic predicate of the type `Tests_Week`. Finally, the actual information about these tests is stored in the array `Num_Tests`, which is an instance of the `Tests_Week` type.

The dynamic predicate of the `Tests_Week` type is verified during the initialization of `Num_Tests`. If we have a non-conformant value there, the check will fail. However, as we can see in our example,

individual assignments to elements of the array do not trigger a check. We can't check for consistency at this point because the initialization of the a complex data structure (such as arrays or records) may not be performed with a single assignment. However, as soon as the object is passed as an argument to a subprogram, the dynamic predicate is checked because the subprogram requires the object to be consistent. This happens in the last call to `Display_Tests` in our example. Here, the predicate check fails because the previous assignment has a non-conformant value.

13.3 Type invariants

Type invariants are another way of specifying expectations regarding types. While predicates are used for *non-private* types, type invariants are used exclusively to define expectations about private types. If a type `T` from a package `P` has a type invariant, the results of operations on objects of type `T` are always consistent with that invariant.

Type invariants are specified with a `with Type_Invariant => <property>` clause. Like predicates, the *property* defines a condition that allows us to check if an object of type `T` is conformant to its requirements. In this sense, type invariants can be viewed as a sort of predicate for private types. However, there are some differences in terms of checks. The following table summarizes the differences:

Element	Subprogram parameter checks	Assignment checks
Predicates	On all <code>in</code> and <code>out</code> parameters	On assignments and explicit initializations
Type invariants	On <code>out</code> parameters returned from subprograms declared in the same public scope	On all initializations

We could rewrite our previous example and replace dynamic predicates by type invariants. It would look like this:

Listing 6: `show_type_invariant.adb`

```

1 with Ada.Text_IO;           use Ada.Text_IO;
2 with Ada.Strings.Unbounded; use Ada.Strings.Unbounded;
3 with Ada.Calendar;         use Ada.Calendar;
4 with Ada.Containers.Vectors;
5
6 procedure Show_Type_Invariant is
7
8   package Courses is
9     type Course is private
10      with Type_Invariant => Check (Course);
11
12     type Course_Container is private;
13
14     procedure Add (CC : in out Course_Container; C : Course);
15
16     function Init
17      (Name : String; Start_Date, End_Date : Time) return Course;
18
19     function Check (C : Course) return Boolean;
20
21   private
22     type Course is record
23       Name      : Unbounded_String;
24       Start_Date : Time;
25       End_Date  : Time;
26     end record;
27

```

(continues on next page)

(continued from previous page)

```

28     function Check (C : Course) return Boolean is
29         (C.Start_Date <= C.End_Date);
30
31     package Course_Vectors is new Ada.Containers.Vectors
32         (Index_Type => Natural,
33          Element_Type => Course);
34
35     type Course_Container is record
36         V : Course_Vectors.Vector;
37     end record;
38 end Courses;
39
40 package body Courses is
41     procedure Add (CC : in out Course_Container; C : Course) is
42     begin
43         CC.V.Append (C);
44     end Add;
45
46     function Init
47         (Name : String; Start_Date, End_Date : Time) return Course is
48     begin
49         return Course'(Name => To_Unbounded_String (Name),
50                        Start_Date => Start_Date,
51                        End_Date => End_Date);
52     end Init;
53 end Courses;
54
55 use Courses;
56
57 CC : Course_Container;
58 begin
59     Add (CC,
60         Init (Name => "Intro to Photography",
61              Start_Date => Time_Of (2018, 5, 1),
62              End_Date => Time_Of (2018, 5, 10)));
63
64     -- This should trigger an error in the type-invariant check
65     Add (CC,
66         Init (Name => "Intro to Video Recording",
67              Start_Date => Time_Of (2019, 5, 1),
68              End_Date => Time_Of (2018, 5, 10)));
69 end Show_Type_Invariant;

```

Build output

```

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

```

Runtime output

```

raised SYSTEM.ASSERTIONS.ASSERT_FAILURE : failed invariant from show_type_
↳ invariant.adb:10

```

The major difference is that the `Course` type was a visible (public) type of the `Courses` package in the previous example, but in this example is a private type.

INTERFACING WITH C

Ada allows us to interface with code in many languages, including C and C++. This section discusses how to interface with C.

14.1 Multi-language project

By default, when using **gprbuild** we only compile Ada source files. To compile C files as well, we need to modify the project file used by **gprbuild**. We use the `Languages` entry, as in the following example:

```
project Multilang is
  for Languages use ("ada", "c");
  for Source_Dirs use ("src");
  for Main use ("main.adb");
  for Object_Dir use "obj";
end Multilang;
```

14.2 Type convention

To interface with data types declared in a C application, you specify the `Convention` aspect on the corresponding Ada type declaration. In the following example, we interface with the `C_Enum` enumeration declared in a C source file:

```
procedure Show_C_Enum is
  type C_Enum is (A, B, C) with Convention => C;
  -- Use C convention for C_Enum
begin
  null;
end Show_C_Enum;
```

To interface with C's built-in types, we use the `Interfaces.C` package, which contains most of the type definitions we need. For example:

```
with Interfaces.C; use Interfaces.C;

procedure Show_C_Struct is
  type c_struct is record
    a : int;
```

(continues on next page)

(continued from previous page)

```

    b : long;
    c : unsigned;
    d : double;
end record
with Convention => C;

begin
    null;
end Show_C_Struct;

```

Here, we're interfacing with a C struct (C_Struct) and using the corresponding data types in C (int, long, unsigned and double). This is the declaration in C:

```

struct c_struct
{
    int      a;
    long     b;
    unsigned c;
    double   d;
};

```

14.3 Foreign subprograms

14.3.1 Calling C subprograms in Ada

We use a similar approach when interfacing with subprograms written in C. In this case, an additional aspect is required: Import. For example:

```

with Interfaces.C; use Interfaces.C;
with Ada.Text_IO; use Ada.Text_IO;

procedure Show_C_Func is

    function my_func (a : int) return int
    with
        Import      => True,
        Convention  => C;
    -- Imports function 'my_func' from C.
    -- You can now call it from Ada.

    V : int;
begin
    V := my_func (2);
    Put_Line ("Result is " & int'Image (V));
end Show_C_Func;

```

This code interfaces with the following declaration in the C header file:

```
int my_func (int a);
```

Here's the corresponding C definition:

```

#include "my_func.h"

int my_func (int a)
{
    return a * 2;
}

```


If you want, you can use a different subprogram name in the Ada code. For example, we could call the C function `Get_Value`:

```
with Interfaces.C; use Interfaces.C;
with Ada.Text_IO; use Ada.Text_IO;

procedure Show_C_Func is

  function Get_Value (a : int) return int
  with
    Import      => True,
    Convention  => C,
    External_Name => "my_func";

  -- Imports function 'my_func' from C and
  -- rename it to 'Get_Value'

  V : int;
begin
  V := Get_Value (2);
  Put_Line ("Result is " & int'Image (V));
end Show_C_Func;
```

14.3.2 Calling Ada subprograms in C

You can also call Ada subprograms from C applications. You do this with the `Export` aspect. For example:

```
with Interfaces.C; use Interfaces.C;

package C_API is

  function My_Func (a : int) return int
  with
    Export      => True,
    Convention  => C,
    External_Name => "my_func";

end C_API;
```

This is the corresponding body that implements that function:

```
package body C_API is

  function My_Func (a : int) return int is
  begin
    return a * 2;
  end My_Func;

end C_API;
```

On the C side, we do the same as we would if the function were written in C: simply declare it using the `extern` keyword. For example:

```
#include <stdio.h>

extern int my_func (int a);

int main (int argc, char **argv) {
```

(continues on next page)

(continued from previous page)

```

int v = my_func(2);

printf("Result is %d\n", v);

return 0;
}

```

14.4 Foreign variables

14.4.1 Using C global variables in Ada

To use global variables from C code, we use the same method as subprograms: we specify the `Import` and `Convention` aspects for each variable we want to import.

Let's reuse an example from the previous section. We'll add a global variable (`func_cnt`) to count the number of times the function (`my_func`) is called:

```

/*% filename: test.h */

extern int func_cnt;

int my_func (int a);

```

The variable is declared in the C file and incremented in `my_func`:

```

#include "test.h"

int func_cnt = 0;

int my_func (int a)
{
    func_cnt++;

    return a * 2;
}

```

In the Ada application, we just reference the foreign variable:

```

with Interfaces.C; use Interfaces.C;
with Ada.Text_IO; use Ada.Text_IO;

procedure Show_C_Func is

    function my_func (a : int) return int
        with
            Import          => True,
            Convention     => C;

    V : int;

    func_cnt : int
        with
            Import          => True,
            Convention     => C;
    -- We can access the func_cnt variable from test.c

begin

```

(continues on next page)

(continued from previous page)

```

V := my_func (1);
V := my_func (2);
V := my_func (3);
Put_Line ("Result is " & int'Image (V));

Put_Line ("Function was called " & int'Image (func_cnt) & " times");
end Show_C_Func;

```

As we see by running the application, the value of the counter is the number of times `my_func` was called.

We can use the `External_Name` aspect to give a different name for the variable in the Ada application in the same we do for subprograms.

14.4.2 Using Ada variables in C

You can also use variables declared in Ada files in C applications. In the same way as we did for subprograms, you do this with the `Export` aspect.

Let's reuse a past example and add a counter, as in the previous example, but this time have the counter incremented in Ada code:

```

with Interfaces.C; use Interfaces.C;

package C_API is

  func_cnt : int := 0
  with
    Export      => True,
    Convention => C;

  function My_Func (a : int) return int
  with
    Export      => True,
    Convention  => C,
    External_Name => "my_func";

end C_API;

```

The variable is then increment in `My_Func`:

```

--% filename: c_api.adb
package body C_API is

  function My_Func (a : int) return int is
  begin
    func_cnt := func_cnt + 1;
    return a * 2;
  end My_Func;

end C_API;

```

In the C application, we just need to declare the variable and use it:

```

#include <stdio.h>

extern int my_func (int a);

extern int func_cnt;

```

(continues on next page)

(continued from previous page)

```

int main (int argc, char **argv) {

    int v;

    v = my_func(1);
    v = my_func(2);
    v = my_func(3);

    printf("Result is %d\n", v);

    printf("Function was called %d times\n", func_cnt);

    return 0;
}

```

Again, by running the application, we see that the value from the counter is the number of times that `my_func` was called.

14.5 Generating bindings

In the examples above, we manually added aspects to our Ada code to correspond to the C source-code we're interfacing with. This is called creating a *binding*. We can automate this process by using the *Ada spec dump* compiler option: `-fdump-ada-spec`. We illustrate this by revisiting our previous example.

This was our C header file:

```

extern int func_cnt;

int my_func (int a);

```

To create Ada bindings, we'll call the compiler like this:

```
gcc -c -fdump-ada-spec -C ./test.h
```

The result is an Ada spec file called `test_h.ads`:

```

pragma Ada_2005;
pragma Style_Checks (Off);

with Interfaces.C; use Interfaces.C;

package test_h is

    func_cnt : aliased int; -- ./test.h:3
    pragma Import (C, func_cnt, "func_cnt");

    function my_func (arg1 : int) return int; -- ./test.h:5
    pragma Import (C, my_func, "my_func");

end test_h;

```

Now we simply refer to this `test_h` package in our Ada application:

```

with Interfaces.C; use Interfaces.C;
with Ada.Text_IO; use Ada.Text_IO;
with test_h; use test_h;

```

(continues on next page)

(continued from previous page)

```

procedure Show_C_Func is
  V : int;
begin
  V := my_func (1);
  V := my_func (2);
  V := my_func (3);
  Put_Line ("Result is " & int'Image (V));

  Put_Line ("Function was called " & int'Image (func_cnt) & " times");
end Show_C_Func;

```

You can specify the name of the parent unit for the bindings you're creating as the operand to `fdump-ada-spec`:

```
gcc -c -fdump-ada-spec -fada-spec-parent=Ext_C_Code -C ./test.h
```

This creates the file `ext_c_code-test_h.ads`:

```

package Ext_C_Code.test_h is
  -- automatic generated bindings...
end Ext_C_Code.test_h;

```

14.5.1 Adapting bindings

The compiler does the best it can when creating bindings for a C header file. However, sometimes it has to guess about the translation and the generated bindings don't always match our expectations. For example, this can happen when creating bindings for functions that have pointers as arguments. In this case, the compiler may use `System.Address` as the type of one or more pointers. Although this approach works fine (as we'll see later), this is usually not how a human would interpret the C header file. The following example illustrates this issue.

Let's start with this C header file:

```

/*% filename: test.h */

struct test;

struct test * test_create(void);
void test_destroy(struct test *t);
void test_reset(struct test *t);
void test_set_name(struct test *t, char *name);
void test_set_address(struct test *t, char *address);
void test_display(const struct test *t);

```

And the corresponding C implementation:

```

#include <stdlib.h>
#include <string.h>
#include <stdio.h>

#include "test.h"

```

(continues on next page)

(continued from previous page)

```

struct test {
    char name[80];
    char address[120];
};

static size_t
strcpy(char *dst, const char *src, size_t dstsize)
{
    size_t len = strlen(src);
    if (dstsize) {
        size_t bl = (len < dstsize-1 ? len : dstsize-1);
        ((char*)memcpy(dst, src, bl))[bl] = 0;
    }
    return len;
}

struct test * test_create(void)
{
    return malloc (sizeof (struct test));
}

void test_destroy(struct test *t)
{
    if (t != NULL) {
        free(t);
    }
}

void test_reset(struct test *t)
{
    t->name[0] = '\0';
    t->address[0] = '\0';
}

void test_set_name(struct test *t, char *name)
{
    strcpy(t->name, name, sizeof(t->name));
}

void test_set_address(struct test *t, char *address)
{
    strcpy(t->address, address, sizeof(t->address));
}

void test_display(const struct test *t)
{
    printf("Name:   %s\n", t->name);
    printf("Address: %s\n", t->address);
}

```

Next, we'll create our bindings:

```
gcc -c -fdump-ada-spec -C ./test.h
```

This creates the following specification in test_h.ads:

```

pragma Ada_2005;
pragma Style_Checks (Off);

with Interfaces.C; use Interfaces.C;

```

(continues on next page)

(continued from previous page)

```

with System;
with Interfaces.C.Strings;

package test_h is

  -- skipped empty struct test

  function test_create return System.Address; -- ./test.h:5
  pragma Import (C, test_create, "test_create");

  procedure test_destroy (arg1 : System.Address); -- ./test.h:7
  pragma Import (C, test_destroy, "test_destroy");

  procedure test_reset (arg1 : System.Address); -- ./test.h:9
  pragma Import (C, test_reset, "test_reset");

  procedure test_set_name (arg1 : System.Address; arg2 : Interfaces.C.Strings.
↳chars_ptr); -- ./test.h:11
  pragma Import (C, test_set_name, "test_set_name");

  procedure test_set_address (arg1 : System.Address; arg2 : Interfaces.C.Strings.
↳chars_ptr); -- ./test.h:13
  pragma Import (C, test_set_address, "test_set_address");

  procedure test_display (arg1 : System.Address); -- ./test.h:15
  pragma Import (C, test_display, "test_display");

end test_h;

```

As we can see, the binding generator completely ignores the declaration `struct test` and all references to the `test` struct are replaced by addresses (`System.Address`). Nevertheless, these bindings are good enough to allow us to create a test application in Ada:

```

with Interfaces.C;           use Interfaces.C;
with Interfaces.C.Strings;  use Interfaces.C.Strings;
with Ada.Text_IO;          use Ada.Text_IO;
with test_h;                use test_h;

with System;

procedure Show_Automatic_C_Struct_Bindings is

  Name      : constant chars_ptr := New_String ("John Doe");
  Address   : constant chars_ptr := New_String ("Small Town");

  T : System.Address := test_create;

begin
  test_reset (T);
  test_set_name (T, Name);
  test_set_address (T, Address);

  test_display (T);
  test_destroy (T);
end Show_Automatic_C_Struct_Bindings;

```

We can successfully bind our C code with Ada using the automatically-generated bindings, but they aren't ideal. Instead, we would prefer Ada bindings that match our (human) interpretation of the C header file. This requires manual analysis of the header file. The good news is that we can use the automatic generated bindings as a starting point and adapt them to our needs. For example, we can:

1. Define a Test type based on System.Address and use it in all relevant functions.
2. Remove the test_ prefix in all operations on the Test type.

This is the resulting specification:

```
with Interfaces.C; use Interfaces.C;
with System;
with Interfaces.C.Strings;

package adapted_test_h is

  type Test is new System.Address;

  function Create return Test;
  pragma Import (C, Create, "test_create");

  procedure Destroy (T : Test);
  pragma Import (C, Destroy, "test_destroy");

  procedure Reset (T : Test);
  pragma Import (C, Reset, "test_reset");

  procedure Set_Name (T      : Test;
                     Name  : Interfaces.C.Strings.chars_ptr); -- ./test.h:11
  pragma Import (C, Set_Name, "test_set_name");

  procedure Set_Address (T      : Test;
                        Address : Interfaces.C.Strings.chars_ptr);
  pragma Import (C, Set_Address, "test_set_address");

  procedure Display (T : Test); -- ./test.h:15
  pragma Import (C, Display, "test_display");

end adapted_test_h;
```

And this is the corresponding Ada body:

```
with Interfaces.C;          use Interfaces.C;
with Interfaces.C.Strings; use Interfaces.C.Strings;
with Ada.Text_IO;          use Ada.Text_IO;
with adapted_test_h;       use adapted_test_h;

with System;

procedure Show_Adapted_C_Struct_Bindings is

  Name      : constant chars_ptr := New_String ("John Doe");
  Address   : constant chars_ptr := New_String ("Small Town");

  T : Test := Create;

begin
  Reset (T);
  Set_Name (T, Name);
  Set_Address (T, Address);

  Display (T);
  Destroy (T);
end Show_Adapted_C_Struct_Bindings;
```

Now we can use the Test type and its operations in a clean, readable way.

OBJECT-ORIENTED PROGRAMMING

Object-oriented programming (OOP) is a large and ill-defined concept in programming languages and one that tends to encompass many different meanings because different languages often implement their own vision of it, with similarities and differences from the implementations in other languages.

However, one model mostly "won" the battle of what object-oriented means, if only by sheer popularity. It's the model used in the Java programming language, which is very similar to the one used by C++. Here are some defining characteristics:

- Type derivation and extension: Most object oriented languages allow the user to add fields to derived types.
- Subtyping: Objects of a type derived from a base type can, in some instances, be substituted for objects of the base type.
- Runtime polymorphism: Calling a subprogram, usually called a *method*, attached to an object type can dispatch at runtime depending on the exact type of the object.
- Encapsulation: Objects can hide some of their data.
- Extensibility: People from the "outside" of your package, or even your whole library, can derive from your object types and define their own behaviors.

Ada dates from before object-oriented programming was as popular as it is today. Some of the mechanisms and concepts from the above list were in the earliest version of Ada even before what we would call OOP was added:

- As we saw, encapsulation is not implemented at the type level in Ada, but instead at the package level.
- Subtyping can be implemented using, well, subtypes, which have a full and permissive static substitutability model. The substitution will fail at runtime if the dynamic constraints of the subtype are not fulfilled.
- Runtime polymorphism can be implemented using variant records.

However, this lists leaves out type extensions, if you don't consider variant records, and extensibility.

The 1995 revision of Ada added a feature filling the gaps, which allowed people to program following the object-oriented paradigm in an easier fashion. This feature is called *tagged types*.

Note: It's possible to program in Ada without ever creating tagged types. If that's your preferred style of programming or you have no specific use for tagged types, feel free to not use them, as is the case for many features of Ada.

However, they can be the best way to express solutions to certain problems and they may be the best way to solve your problem. If that's the case, read on!

15.1 Derived types

Before presenting tagged types, we should discuss a topic we have brushed on, but not really covered, up to now:

You can create one or more new types from every type in Ada. Type derivation is built into the language.

Listing 1: newtypes.ads

```

1 package Newtypes is
2   type Point is record
3     X, Y : Integer;
4   end record;
5
6   type New_Point is new Point;
7 end Newtypes;
```

Type derivation is useful to enforce strong typing because the type system treats the two types as incompatible.

But the benefits are not limited to that: you can inherit things from the type you derive from. You not only inherit the representation of the data, but you can also inherit behavior.

When you inherit a type you also inherit what are called *primitive operations*. A primitive operation (or just a *primitive*) is a subprogram attached to a type. Ada defines primitives as subprograms defined in the same scope as the type.

Attention: A subprogram will only become a primitive of the type if:

1. The subprogram is declared in the same scope as the type and
2. The type and the subprogram are declared in a package

Listing 2: primitives.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Primitives is
4   package Week is
5     type Days is (Monday, Tuesday, Wednesday, Thursday,
6                 Friday, Saturday, Sunday);
7
8     -- Print day is a primitive of the type Days
9     procedure Print_Day (D : Days);
10  end Week;
11
12  package body Week is
13    procedure Print_Day (D : Days) is
14      begin
15        Put_Line (Days'Image (D));
16      end Print_Day;
17  end Week;
18
19  use Week;
20  type Weekend_Days is new Days range Saturday .. Sunday;
21
22  -- A procedure Print_Day is automatically inherited here. It is as if
23  -- the procedure
24  --
```

(continues on next page)

(continued from previous page)

```

25  -- procedure Print_Day (D : Weekend_Days);
26  --
27  -- has been declared with the same body
28
29  Sat : Weekend_Days := Saturday;
30  begin
31  Print_Day (Sat);
32  end Primitives;

```

Build output

```

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

```

Runtime output

```
SATURDAY
```

This kind of inheritance can be very useful, and is not limited to record types (you can use it on discrete types, as in the example above), but it's only superficially similar to object-oriented inheritance:

- Records can't be extended using this mechanism alone. You also can't specify a new representation for the new type: it will **always** have the same representation as the base type.
- There's no facility for dynamic dispatch or polymorphism. Objects are of a fixed, static type.

There are other differences, but it's not useful to list them all here. Just remember that this is a kind of inheritance you can use if you only want to statically inherit behavior without duplicating code or using composition, but a kind you can't use if you want any dynamic features that are usually associated with OOP.

15.2 Tagged types

The 1995 revision of the Ada language introduced tagged types to fulfill the need for an unified solution that allows programming in an object-oriented style similar to the one described at the beginning of this chapter.

Tagged types are very similar to normal records except that some functionality is added:

- Types have a *tag*, stored inside each object, that identifies the *runtime type*¹⁹ of that object.
- Primitives can dispatch. A primitive on a tagged type is what you would call a *method* in Java or C++. If you derive a base type and override a primitive of it, you can often call it on an object with the result that which primitive is called depends on the exact runtime type of the object.
- Subtyping rules are introduced allowing a tagged type derived from a base type to be statically compatible with the base type.

Let's see our first tagged type declarations:

¹⁹ https://en.wikipedia.org/wiki/Run-time_type_information

Listing 3: p.ads

```

1 package P is
2   type My_Class is tagged null record;
3     -- Just like a regular record, but with tagged qualifier
4
5     -- Methods are outside of the type definition:
6
7   procedure Foo (Self : in out My_Class);
8     -- If you define a procedure taking a My_Class argument
9     -- in the same package, it will be a method.
10
11    -- Here's how you derive a tagged type:
12
13   type Derived is new My_Class with record
14     A : Integer;
15     -- You can add fields in derived types.
16   end record;
17
18   overriding procedure Foo (Self : in out Derived);
19     -- The "overriding" qualifier is optional, but if it is present,
20     -- it must be valid.
21 end P;
```

Listing 4: p.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 package body P is
4   procedure Foo (Self : in out My_Class) is
5     begin
6       Put_Line ("In My_Class.Foo");
7     end Foo;
8
9   procedure Foo (Self : in out Derived) is
10    begin
11      Put_Line ("In Derived.Foo, A = " & Integer'Image (Self.A));
12    end Foo;
13 end P;
```

15.3 Classwide types

To remain consistent with the rest of the language, a new notation needed to be introduced to say "This object is of this type or any descendent derives tagged type".

In Ada, we call this the *classwide type*. It's used in OOP as soon as you need polymorphism. For example, you can't do the following:

Listing 5: main.adb

```

1 with P; use P;
2
3 procedure Main is
4
5   O1 : My_Class;
6   -- Declaring an object of type My_Class
7
8   O2 : Derived := (A => 12);
```

(continues on next page)

(continued from previous page)

```

9   -- Declaring an object of type Derived
10
11  O3 : My_Class := O2;
12  -- INVALID: Trying to assign a value of type derived to a variable of
13  -- type My_Class.
14  begin
15      null;
16  end Main;

```

Compilation output

```

main.adb:11:21: expected type "My_Class" defined at p.ads:2
main.adb:11:21: found type "Derived" defined at p.ads:13

```

This is because an object of a type T is exactly of the type T, whether T is tagged or not. What you want to say as a programmer is "I want O3 to be able to hold an object of type My_Class or any type descending from My_Class". Here's how you do that:

Listing 6: main.adb

```

1  with P; use P;
2
3  procedure Main is
4      O1 : My_Class;
5      -- Declare an object of type My_Class
6
7      O2 : Derived := (A => 12);
8      -- Declare an object of type Derived
9
10     O3 : My_Class'Class := O2;
11     -- Now valid: My_Class'Class designates the classwide type for
12     -- My_Class, which is the set of all types descending from My_Class
13     -- (including My_Class).
14  begin
15      null;
16  end Main;

```

Build output

```

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

```

Attention: Because an object of a classwide type can be the size of any descendent of its base type, it has an unknown size. It's therefore an indefinite type, with the expected restrictions:

- It can't be stored as a field/component of a record
- An object of a classwide type needs to be initialized immediately (you can't specify the constraints of such a type in any way other than by initializing it).

15.4 Dispatching operations

We saw that you can override operations in types derived from another tagged type. The eventual goal of OOP is to make a dispatching call: a call to a primitive (method) that depends on the exact type of the object.

But, if you think carefully about it, a variable of type `My_Class` always contains an object of exactly that type. If you want to have a variable that can contain a `My_Class` or any derived type, it has to be of type `My_Class'Class`.

In other words, to make a dispatching call, you must first have an object that can be either of a type or any type derived from this type, namely an object of a classwide type.

Listing 7: main.adb

```
1 with P; use P;
2
3 procedure Main is
4   O1 : My_Class;
5   -- Declare an object of type My_Class
6
7   O2 : Derived := (A => 12);
8   -- Declare an object of type Derived
9
10  O3 : My_Class'Class := O2;
11
12  O4 : My_Class'Class := O1;
13 begin
14   Foo (O1);
15   -- Non dispatching: Calls My_Class.Foo
16   Foo (O2);
17   -- Non dispatching: Calls Derived.Foo
18   Foo (O3);
19   -- Dispatching: Calls Derived.Foo
20   Foo (O4);
21   -- Dispatching: Calls My_Class.Foo
22 end Main;
```

Build output

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

Runtime output

```
In My_Class.Foo
In Derived.Foo, A = 12
In Derived.Foo, A = 12
In My_Class.Foo
```

Attention

You can convert an object of type `Derived` to an object of type `My_Class`. This is called a *view conversion* in Ada parlance and is useful, for example, if you want to call a parent method.

In that case, the object really is converted to a `My_Class` object, which means its tag is changed.

Since tagged objects are always passed by reference, you can use this kind of conversion to modify the state of an object: changes to converted object will affect the original one.

Listing 8: main.adb

```

1  with P; use P;
2
3  procedure Main is
4      01 : Derived := (A => 12);
5      -- Declare an object of type Derived
6
7      02 : My_Class := My_Class (01);
8
9      03 : My_Class'Class := 02;
10 begin
11     Foo (01);
12     -- Non dispatching: Calls Derived.Foo
13     Foo (02);
14     -- Non dispatching: Calls My_Class.Foo
15
16     Foo (03);
17     -- Dispatching: Calls My_Class.Foo
18 end Main;

```

Build output

```

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

```

Runtime output

```

In Derived.Foo, A = 12
In My_Class.Foo
In My_Class.Foo

```

15.5 Dot notation

You can also call primitives of tagged types with a notation that's more familiar to object oriented programmers. Given the Foo primitive above, you can also write the above program this way:

Listing 9: main.adb

```

1  with P; use P;
2
3  procedure Main is
4      01 : My_Class;
5      -- Declare an object of type My_Class
6
7      02 : Derived := (A => 12);
8      -- Declare an object of type Derived
9
10     03 : My_Class'Class := 02;

```

(continues on next page)

(continued from previous page)

```

11
12     04 : My_Class'Class := 01;
13 begin
14     01.Foo;
15     -- Non dispatching: Calls My_Class.Foo
16     02.Foo;
17     -- Non dispatching: Calls Derived.Foo
18     03.Foo;
19     -- Dispatching: Calls Derived.Foo
20     04.Foo;
21     -- Dispatching: Calls My_Class.Foo
22 end Main;

```

Build output

```

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

```

Runtime output

```

In My_Class.Foo
In Derived.Foo, A = 12
In Derived.Foo, A = 12
In My_Class.Foo

```

If the dispatching parameter of a primitive is the first parameter, which is the case in our examples, you can call the primitive using the dot notation. Any remaining parameter are passed normally:

Listing 10: main.adb

```

1  with P; use P;
2
3  procedure Main is
4      package Extend is
5          type D2 is new Derived with null record;
6
7          procedure Bar (Self : in out D2; Val : Integer);
8      end Extend;
9
10     package body Extend is
11         procedure Bar (Self : in out D2; Val : Integer) is
12             begin
13                 Self.A := Self.A + Val;
14             end Bar;
15         end Extend;
16
17     use Extend;
18
19     Obj : D2 := (A => 15);
20 begin
21     Obj.Bar (2);
22     Obj.Foo;
23 end Main;

```

Build output


```

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

```

Runtime output

```
In Derived.Foo, A = 17
```

15.6 Private & Limited

We've seen previously (in the *Privacy* (page 101) chapter) that types can be declared limited or private. These encapsulation techniques can also be applied to tagged types, as we'll see in this section.

This is an example of a tagged private type:

Listing 11: p.ads

```

1 package P is
2   type T is tagged private;
3 private
4   type T is tagged record
5     E : Integer;
6   end record;
7 end P;

```

This is an example of a tagged limited type:

Listing 12: p.ads

```

1 package P is
2   type T is tagged limited record
3     E : Integer;
4   end record;
5 end P;

```

Naturally, you can combine both *limited* and *private* types and declare a tagged limited private type:

Listing 13: p.ads

```

1 package P is
2   type T is tagged limited private;
3
4   procedure Init (A : in out T);
5 private
6   type T is tagged limited record
7     E : Integer;
8   end record;
9 end P;

```

Listing 14: p.adb

```

1 package body P is
2
3   procedure Init (A : in out T) is
4     begin
5       A.E := 0;
6     end Init;
7
8 end P;
```

Listing 15: main.adb

```

1 with P; use P;
2
3 procedure Main is
4   T1, T2 : T;
5   begin
6     T1.Init;
7     T2.Init;
8
9     -- The following line doesn't work because type T is private:
10    -- T1.E := 0;
11
12    -- The following line doesn't work because type T is limited:
13    -- T2 := T1;
14 end Main;
```

Build output

```

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

Note that the code in the Main procedure above presents two assignments that trigger compilation errors because type T is limited private. In fact, you cannot:

- assign to T1.E directly because type T is private;
- assign T1 to T2 because type T is limited.

In this case, there's no distinction between tagged and non-tagged types: these compilation errors would also occur for non-tagged types.

15.7 Classwide access types

In this section, we'll discuss an useful pattern for object-oriented programming in Ada: classwide access type. Let's start with an example where we declare a tagged type T and a derived type T_New:

Listing 16: p.ads

```

1 package P is
2   type T is tagged null record;
3
```

(continues on next page)

(continued from previous page)

```

4  procedure Show (Dummy : T);
5
6  type T_New is new T with null record;
7
8  procedure Show (Dummy : T_New);
9  end P;

```

Listing 17: p.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2
3  package body P is
4
5      procedure Show (Dummy : T) is
6      begin
7          Put_Line ("Using type " & T'External_Tag);
8      end Show;
9
10     procedure Show (Dummy : T_New) is
11     begin
12         Put_Line ("Using type " & T_New'External_Tag);
13     end Show;
14
15 end P;

```

Note that we're using null records for both types T and T_New. Although these types don't actually have any component, we can still use them to demonstrate dispatching. Also note that the example above makes use of the 'External_Tag attribute in the implementation of the Show procedure to get a string for the corresponding tagged type.

As we've seen before, we must use a classwide type to create objects that can make dispatching calls. In other words, objects of type T'Class will dispatch. For example:

Listing 18: dispatching_example.adb

```

1  with P; use P;
2
3  procedure Dispatching_Example is
4      T2          :          T_New;
5      T_Dispatch : constant T'Class := T2;
6  begin
7      T_Dispatch.Show;
8  end Dispatching_Example;

```

Build output

```

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

```

Runtime output

```
Using type P.T_NEW
```

A more useful application is to declare an array of objects that can dispatch. For example, we'd like to declare an array T_Arr, loop over this array and dispatch according to the actual type of each

individual element:

```
for I in T_Arr'Range loop
  T_Arr (I).Show;
  -- Call Show procedure according to actual type of T_Arr (I)
end loop;
```

However, it's not possible to declare an array of type T'Class directly:

Listing 19: classwide_compilation_error.adb

```
1 with P; use P;
2
3 procedure Classwide_Compilation_Error is
4   T_Arr : array (1 .. 2) of T'Class;
5   --           ^ Compilation Error!
6 begin
7   for I in T_Arr'Range loop
8     T_Arr (I).Show;
9   end loop;
10 end Classwide_Compilation_Error;
```

Compilation output

```
classwide_compilation_error.adb:4:31: unconstrained element type in array_
↳declaration
```

In fact, it's impossible for the compiler to know which type would actually be used for each element of the array. However, if we use dynamic allocation via access types, we can allocate objects of different types for the individual elements of an array T_Arr. We do this by using classwide access types, which have the following format:

```
type T_Class is access T'Class;
```

We can rewrite the previous example using the T_Class type. In this case, dynamically allocated objects of this type will dispatch according to the actual type used during the allocation. Also, let's introduce an Init procedure that won't be overridden for the derived T_New type. This is the adapted code:

Listing 20: p.ads

```
1 package P is
2   type T is tagged record
3     E : Integer;
4   end record;
5
6   type T_Class is access T'Class;
7
8   procedure Init (A : in out T);
9
10  procedure Show (Dummy : T);
11
12  type T_New is new T with null record;
13
14  procedure Show (Dummy : T_New);
15
16 end P;
```

Listing 21: p.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
```

(continues on next page)

(continued from previous page)

```

3 package body P is
4
5     procedure Init (A : in out T) is
6     begin
7         Put_Line ("Initializing type T...");
8         A.E := 0;
9     end Init;
10
11    procedure Show (Dummy : T) is
12    begin
13        Put_Line ("Using type " & T'External_Tag);
14    end Show;
15
16    procedure Show (Dummy : T_New) is
17    begin
18        Put_Line ("Using type " & T_New'External_Tag);
19    end Show;
20
21 end P;

```

Listing 22: main.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with P;           use P;
3
4 procedure Main is
5     T_Arr : array (1 .. 2) of T_Class;
6 begin
7     T_Arr (1) := new T;
8     T_Arr (2) := new T_New;
9
10    for I in T_Arr'Range loop
11        Put_Line ("Element # " & Integer'Image (I));
12
13        T_Arr (I).Init;
14        T_Arr (I).Show;
15
16        Put_Line ("-----");
17    end loop;
18 end Main;

```

Build output

```

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

```

Runtime output

```

Element # 1
Initializing type T...
Using type P.T
-----
Element # 2
Initializing type T...
Using type P.T_NEW

```

(continues on next page)

(continued from previous page)

In this example, the first element (`T_Arr (1)`) is of type `T`, while the second element is of type `T_New`. When running the example, the `Init` procedure of type `T` is called for both elements of the `T_Arr` array, while the call to the `Show` procedure selects the corresponding procedure according to the type of each element of `T_Arr`.

STANDARD LIBRARY: CONTAINERS

In previous chapters, we've used arrays as the standard way to group multiple objects of a specific data type. In many cases, arrays are good enough for manipulating those objects. However, there are situations that require more flexibility and more advanced operations. For those cases, Ada provides support for containers — such as vectors and sets — in its standard library.

We present an introduction to containers here. For a list of all containers available in Ada, see [Appendix B](#) (page 257).

16.1 Vectors

In the following sections, we present a general overview of vectors, including instantiation, initialization, and operations on vector elements and vectors.

16.1.1 Instantiation

Here's an example showing the instantiation and declaration of a vector *V*:

Listing 1: show_vector_inst.adb

```
1 with Ada.Containers.Vectors;  
2  
3 procedure Show_Vector_Inst is  
4  
5     package Integer_Vectors is new Ada.Containers.Vectors  
6         (Index_Type => Natural,  
7          Element_Type => Integer);  
8  
9     V : Integer_Vectors.Vector;  
10 begin  
11     null;  
12 end Show_Vector_Inst;
```

Containers are based on generic packages, so we can't simply declare a vector as we would declare an array of a specific type:

```
A : array (1 .. 10) of Integer;
```

Instead, we first need to instantiate one of those packages. We with the container package (`Ada.Containers.Vectors` in this case) and instantiate it to create an instance of the generic package for the desired type. Only then can we declare the vector using the type from the instantiated package. This instantiation needs to be done for any container type from the standard library.

In the instantiation of `Integer_Vectors`, we indicate that the vector contains elements of `Integer` type by specifying it as the `Element_Type`. By setting `Index_Type` to `Natural`, we specify

that the allowed range includes all natural numbers. We could have used a more restrictive range if desired.

16.1.2 Initialization

One way to initialize a vector is from a concatenation of elements. We use the & operator, as shown in the following example:

Listing 2: show_vector_init.adb

```
1 with Ada.Containers; use Ada.Containers;
2 with Ada.Containers.Vectors;
3
4 with Ada.Text_IO; use Ada.Text_IO;
5
6 procedure Show_Vector_Init is
7
8     package Integer_Vectors is new Ada.Containers.Vectors
9         (Index_Type => Natural,
10          Element_Type => Integer);
11
12     use Integer_Vectors;
13
14     V : Vector := 20 & 10 & 0 & 13;
15 begin
16     Put_Line ("Vector has "
17             & Count_Type'Image (V.Length) & " elements");
18 end Show_Vector_Init;
```

Build output

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

Runtime output

```
Vector has 4 elements
```

We specify `use Integer_Vectors`, so we have direct access to the types and operations from the instantiated package. Also, the example introduces another operation on the vector: `Length`, which retrieves the number of elements in the vector. We can use the dot notation because `Vector` is a tagged type, allowing us to write either `V.Length` or `Length (V)`.

16.1.3 Appending and prepending elements

You add elements to a vector using the `Prepend` and `Append` operations. As the names suggest, these operations add elements to the beginning or end of a vector, respectively. For example:

Listing 3: show_vector_append.adb

```
1 with Ada.Containers; use Ada.Containers;
2 with Ada.Containers.Vectors;
3
```

(continues on next page)

(continued from previous page)

```

4 with Ada.Text_IO; use Ada.Text_IO;
5
6 procedure Show_Vector_Append is
7
8   package Integer_Vectors is new Ada.Containers.Vectors
9     (Index_Type => Natural,
10      Element_Type => Integer);
11
12   use Integer_Vectors;
13
14   V : Vector;
15 begin
16   Put_Line ("Appending some elements to the vector...");
17   V.Append (20);
18   V.Append (10);
19   V.Append (0);
20   V.Append (13);
21   Put_Line ("Finished appending.");
22
23   Put_Line ("Prepending some elements to the vector...");
24   V.Prepend (30);
25   V.Prepend (40);
26   V.Prepend (100);
27   Put_Line ("Finished prepending.");
28
29   Put_Line ("Vector has "
30           & Count_Type'Image (V.Length) & " elements");
31 end Show_Vector_Append;

```

Build output

```

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

```

Runtime output

```

Appending some elements to the vector...
Finished appending.
Prepending some elements to the vector...
Finished prepending.
Vector has 7 elements

```

This example puts elements into the vector in the following sequence: (100, 40, 30, 20, 10, 0, 13).

The Reference Manual specifies that the worst-case complexity must be:

- $O(\log N)$ for the Append operation, and
- $O(N \log N)$ for the Prepend operation.

16.1.4 Accessing first and last elements

We access the first and last elements of a vector using the `First_Element` and `Last_Element` functions. For example:

Listing 4: `show_vector_first_last_element.adb`

```

1 with Ada.Containers; use Ada.Containers;
2 with Ada.Containers.Vectors;
3
4 with Ada.Text_IO; use Ada.Text_IO;
5
6 procedure Show_Vector_First_Last_Element is
7
8     package Integer_Vectors is new Ada.Containers.Vectors
9         (Index_Type => Natural,
10          Element_Type => Integer);
11
12     use Integer_Vectors;
13
14     function Img (I : Integer) return String renames Integer'Image;
15     function Img (I : Count_Type) return String renames Count_Type'Image;
16
17     V : Vector := 20 & 10 & 0 & 13;
18 begin
19     Put_Line ("Vector has " & Img (V.Length) & " elements");
20
21     -- Using V.First_Element to retrieve first element
22     Put_Line ("First element is " & Img (V.First_Element));
23
24     -- Using V.Last_Element to retrieve last element
25     Put_Line ("Last element is " & Img (V.Last_Element));
26 end Show_Vector_First_Last_Element;
```

Build output

```

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

Runtime output

```

Vector has 4 elements
First element is 20
Last element is 13
```

You can swap elements by calling the procedure `Swap` and retrieving a reference (a *cursor*) to the first and last elements of the vector by calling `First` and `Last`. A cursor allows us to iterate over a container and process individual elements from it.

With these operations, we're able to write code to swap the first and last elements of a vector:

Listing 5: `show_vector_first_last_element.adb`

```

1 with Ada.Containers; use Ada.Containers;
2 with Ada.Containers.Vectors;
3
4 with Ada.Text_IO; use Ada.Text_IO;
```

(continues on next page)

(continued from previous page)

```

5
6 procedure Show_Vector_First_Last_Element is
7
8   package Integer_Vectors is new Ada.Containers.Vectors
9     (Index_Type => Natural,
10      Element_Type => Integer);
11
12   use Integer_Vectors;
13
14   function Img (I : Integer) return String renames Integer'Image;
15
16   V : Vector := 20 & 10 & 0 & 13;
17 begin
18   -- We use V.First and V.Last to retrieve cursor for first and
19   -- last elements.
20   -- We use V.Swap to swap elements.
21   V.Swap (V.First, V.Last);
22
23   Put_Line ("First element is now " & Img (V.First_Element));
24   Put_Line ("Last element is now " & Img (V.Last_Element));
25 end Show_Vector_First_Last_Element;
```

Build output

```

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

Runtime output

```

First element is now 13
Last element is now 20
```

16.1.5 Iterating

The easiest way to iterate over a container is to use a `for E of Our_Container` loop. This gives us a reference (E) to the element at the current position. We can then use E directly. For example:

Listing 6: show_vector_iteration.adb

```

1 with Ada.Containers.Vectors;
2
3 with Ada.Text_IO; use Ada.Text_IO;
4
5 procedure Show_Vector_Iteration is
6
7   package Integer_Vectors is new Ada.Containers.Vectors
8     (Index_Type => Natural,
9      Element_Type => Integer);
10
11   use Integer_Vectors;
12
13   function Img (I : Integer) return String renames Integer'Image;
14
15   V : Vector := 20 & 10 & 0 & 13;
```

(continues on next page)

(continued from previous page)

```

16 begin
17   Put_Line ("Vector elements are: ");
18
19   --
20   -- Using for ... of loop to iterate:
21   --
22   for E of V loop
23     Put_Line ("- " & Img (E));
24   end loop;
25
26 end Show_Vector_Iteration;

```

Build output

```

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

```

Runtime output

```

Vector elements are:
- 20
- 10
- 0
- 13

```

This code displays each element from the vector V.

Because we're given a reference, we can display not only the value of an element but also modify it. For example, we could easily write a loop to add one to each element of vector V:

```

for E of V loop
  E := E + 1;
end loop;

```

We can also use indices to access vector elements. The format is similar to a loop over array elements: we use a `for I in <range> loop`. The range is provided by `V.First_Index` and `V.Last_Index`. We can access the current element by using it as an array index: `V (I)`. For example:

Listing 7: show_vector_index_iteration.adb

```

1 with Ada.Containers.Vectors;
2
3 with Ada.Text_IO; use Ada.Text_IO;
4
5 procedure Show_Vector_Index_Iteration is
6
7   package Integer_Vectors is new Ada.Containers.Vectors
8     (Index_Type => Natural,
9      Element_Type => Integer);
10
11   use Integer_Vectors;
12
13   V : Vector := 20 & 10 & 0 & 13;
14 begin
15   Put_Line ("Vector elements are: ");

```

(continues on next page)

(continued from previous page)

```

16
17  --
18  -- Using indices in a "for I in ..." loop to iterate:
19  --
20  for I in V.First_Index .. V.Last_Index loop
21    -- Displaying current index I
22    Put (" - ["
23         & Extended_Index'Image (I)
24         & "]" );
25
26    Put (Integer'Image (V (I)));
27
28    -- We could also use the V.Element (I) function to retrieve the
29    -- element at the current index I
30
31    New_Line;
32  end loop;
33
34 end Show_Vector_Index_Iteration;

```

Build output

```

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

```

Runtime output

```

Vector elements are:
- [ 0] 20
- [ 1] 10
- [ 2] 0
- [ 3] 13

```

Here, in addition to displaying the vector elements, we're also displaying each index, *I*, just like what we can do for array indices. Also, we can access the element by using either the short form *V (I)* or the longer form *V.Element (I)* but not *V.I*.

As mentioned in the previous section, you can use cursors to iterate over containers. For this, use the function *Iterate*, which retrieves a cursor for each position in the vector. The corresponding loop has the format *for C in V.Iterate loop*. Like the previous example using indices, you can again access the current element by using the cursor as an array index: *V (C)*. For example:

Listing 8: show_vector_cursor_iteration.adb

```

1  with Ada.Containers.Vectors;
2
3  with Ada.Text_IO; use Ada.Text_IO;
4
5  procedure Show_Vector_Cursor_Iteration is
6
7    package Integer_Vectors is new Ada.Containers.Vectors
8      (Index_Type => Natural,
9       Element_Type => Integer);
10
11   use Integer_Vectors;
12

```

(continues on next page)

```

13   V : Vector := 20 & 10 & 0 & 13;
14 begin
15   Put_Line ("Vector elements are: ");
16
17   --
18   -- Use a cursor to iterate in a loop:
19   --
20   for C in V.Iterate loop
21     -- Using To_Index function to retrieve index
22     -- for the cursor position
23     Put ("- ["
24         & Extended_Index'Image (To_Index (C))
25         & "] ");
26
27     Put (Integer'Image (V (C)));
28
29     -- We could use Element (C) to retrieve the vector
30     -- element for the cursor position
31
32     New_Line;
33   end loop;
34
35   -- Alternatively, we could iterate with a while-loop:
36   --
37   -- declare
38   --   C : Cursor := V.First;
39   -- begin
40   --   while C /= No_Element loop
41   --     some processing here...
42   --
43   --     C := Next (C);
44   --   end loop;
45   -- end;
46
47 end Show_Vector_Cursor_Iteration;

```

Build output

```

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

```

Runtime output

```

Vector elements are:
- [ 0] 20
- [ 1] 10
- [ 2]  0
- [ 3] 13

```

Instead of accessing an element in the loop using `V (C)`, we could also have used the longer form `Element (C)`. In this example, we're using the function `To_Index` to retrieve the index corresponding to the current cursor.

As shown in the comments after the loop, we could also use a `while ... loop` to iterate over the vector. In this case, we would start with a cursor for the first element (retrieved by calling `V.First`) and then call `Next (C)` to retrieve a cursor for subsequent elements. `Next (C)` returns `No_Element` when the cursor reaches the end of the vector.

You can directly modify the elements using a reference. This is what it looks like when using both indices and cursors:

```
-- Modify vector elements using index
for I in V.First_Index .. V.Last_Index loop
  V (I) := V (I) + 1;
end loop;

-- Modify vector elements using cursor
for C in V.Iterate loop
  V (C) := V (C) + 1;
end loop;
```

The Reference Manual requires that the worst-case complexity for accessing an element be $O(\log N)$.

Another way of modifying elements of a vector is using a *process procedure*, which takes an individual element and does some processing on it. You can call `Update_Element` and pass both a cursor and an access to the process procedure. For example:

Listing 9: show_vector_update.adb

```
1 with Ada.Containers.Vectors;
2
3 with Ada.Text_IO; use Ada.Text_IO;
4
5 procedure Show_Vector_Update is
6
7   package Integer_Vectors is new Ada.Containers.Vectors
8     (Index_Type => Natural,
9      Element_Type => Integer);
10
11  use Integer_Vectors;
12
13  procedure Add_One (I : in out Integer) is
14  begin
15    I := I + 1;
16  end Add_One;
17
18  V : Vector := 20 & 10 & 12;
19 begin
20   --
21   -- Use V.Update_Element to process elements
22   --
23   for C in V.Iterate loop
24     V.Update_Element (C, Add_One'Access);
25   end loop;
26
27 end Show_Vector_Update;
```

Build output

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

16.1.6 Finding and changing elements

You can locate a specific element in a vector by retrieving its index. `Find_Index` retrieves the index of the first element matching the value you're looking for. Alternatively, you can use `Find` to retrieve a cursor referencing that element. For example:

Listing 10: show_find_vector_element.adb

```

1 with Ada.Containers.Vectors;
2
3 with Ada.Text_IO; use Ada.Text_IO;
4
5 procedure Show_Find_Vector_Element is
6
7     package Integer_Vectors is new Ada.Containers.Vectors
8         (Index_Type => Natural,
9          Element_Type => Integer);
10
11     use Integer_Vectors;
12
13     V : Vector := 20 & 10 & 0 & 13;
14     Idx : Extended_Index;
15     C : Cursor;
16 begin
17     -- Using Find_Index to retrieve index of element with value 10
18     Idx := V.Find_Index (10);
19     Put_Line ("Index of element with value 10 is "
20             & Extended_Index'Image (Idx));
21
22     -- Using Find to retrieve cursor for element with value 13
23     C := V.Find (13);
24     Idx := To_Index (C);
25     Put_Line ("Index of element with value 13 is "
26             & Extended_Index'Image (Idx));
27 end Show_Find_Vector_Element;
```

Build output

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

Runtime output

```

Index of element with value 10 is 1
Index of element with value 13 is 3
```

As we saw in the previous section, we can directly access vector elements by using either an index or cursor. However, an exception is raised if we try to access an element with an invalid index or cursor, so we must check whether the index or cursor is valid before using it to access an element. In our example, `Find_Index` or `Find` might not have found the element in the vector. We check for this possibility by comparing the index to `No_Index` or the cursor to `No_Element`. For example:

```

-- Modify vector element using index
if Idx /= No_Index then
    V (Idx) := 11;
end if;
```

(continues on next page)

(continued from previous page)

```
-- Modify vector element using cursor
if C /= No_Element then
  V (C) := 14;
end if;
```

Instead of writing `V (C) := 14`, we could use the longer form `V.Replace_Element (C, 14)`.

16.1.7 Inserting elements

In the previous sections, we've seen examples of how to add elements to a vector:

- using the concatenation operator (&) at the vector declaration, or
- calling the Prepend and Append procedures.

You may want to insert an element at a specific position, e.g. before a certain element in the vector. You do this by calling `Insert`. For example:

Listing 11: show_vector_insert.adb

```
1 with Ada.Containers; use Ada.Containers;
2 with Ada.Containers.Vectors;
3
4 with Ada.Text_IO; use Ada.Text_IO;
5
6 procedure Show_Vector_Insert is
7
8   package Integer_Vectors is new Ada.Containers.Vectors
9     (Index_Type => Natural,
10      Element_Type => Integer);
11
12   use Integer_Vectors;
13
14   procedure Show_Elements (V : Vector) is
15   begin
16     New_Line;
17     Put_Line ("Vector has "
18              & Count_Type'Image (V.Length) & " elements");
19
20     if not V.Is_Empty then
21       Put_Line ("Vector elements are: ");
22       for E of V loop
23         Put_Line ("- " & Integer'Image (E));
24       end loop;
25     end if;
26   end Show_Elements;
27
28   V : Vector := 20 & 10 & 12;
29   C : Cursor;
30 begin
31   Show_Elements (V);
32
33   New_Line;
34   Put_Line ("Adding element with value 9 (before 10)...");
35
36   --
37   -- Using V.Insert to insert element into vector
38   --
39   C := V.Find (10);
40   if C /= No_Element then
```

(continues on next page)

(continued from previous page)

```

41     V.Insert (C, 9);
42   end if;
43
44   Show_Elements (V);
45
46 end Show_Vector_Insert;
```

Build output

```

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

Runtime output

```

Vector has 3 elements
Vector elements are:
- 20
- 10
- 12

Adding element with value 9 (before 10)...

Vector has 4 elements
Vector elements are:
- 20
- 9
- 10
- 12
```

In this example, we're looking for an element with the value of 10. If we find it, we insert an element with the value of 9 before it.

16.1.8 Removing elements

You can remove elements from a vector by passing either a valid index or cursor to the `Delete` procedure. If we combine this with the functions `Find_Index` and `Find` from the previous section, we can write a program that searches for a specific element and deletes it, if found:

Listing 12: show_remove_vector_element.adb

```

1  with Ada.Containers.Vectors;
2
3  with Ada.Text_IO; use Ada.Text_IO;
4
5  procedure Show_Remove_Vector_Element is
6    package Integer_Vectors is new Ada.Containers.Vectors
7      (Index_Type => Natural,
8       Element_Type => Integer);
9
10   use Integer_Vectors;
11
12   V : Vector := 20 & 10 & 0 & 13 & 10 & 13;
13   Idx : Extended_Index;
```

(continues on next page)

(continued from previous page)

```

14   C   : Cursor;
15 begin
16   -- Use Find_Index to retrieve index of element with value 10
17   Idx := V.Find_Index (10);
18
19   -- Checking whether index is valid
20   if Idx /= No_Index then
21     -- Removing element using V.Delete
22     V.Delete (Idx);
23   end if;
24
25   -- Use Find to retrieve cursor for element with value 13
26   C := V.Find (13);
27
28   -- Check whether index is valid
29   if C /= No_Element then
30     -- Remove element using V.Delete
31     V.Delete (C);
32   end if;
33
34 end Show_Remove_Vector_Element;

```

Build output

```

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

```

We can extend this approach to delete all elements matching a certain value. We just need to keep searching for the element in a loop until we get an invalid index or cursor. For example:

Listing 13: show_remove_vector_elements.adb

```

1 with Ada.Containers; use Ada.Containers;
2 with Ada.Containers.Vectors;
3
4 with Ada.Text_IO; use Ada.Text_IO;
5
6 procedure Show_Remove_Vector_Elements is
7
8   package Integer_Vectors is new Ada.Containers.Vectors
9     (Index_Type => Natural,
10      Element_Type => Integer);
11
12   use Integer_Vectors;
13
14   procedure Show_Elements (V : Vector) is
15   begin
16     New_Line;
17     Put_Line ("Vector has " & Count_Type'Image (V.Length) & " elements");
18
19     if not V.Is_Empty then
20       Put_Line ("Vector elements are: ");
21       for E of V loop
22         Put_Line ("- " & Integer'Image (E));
23       end loop;
24     end if;

```

(continues on next page)

(continued from previous page)

```

25   end Show_Elements;
26
27   V : Vector := 20 & 10 & 0 & 13 & 10 & 14 & 13;
28   begin
29     Show_Elements (V);
30
31     --
32     -- Remove elements using an index
33     --
34     declare
35       E : constant Integer := 10;
36       I : Extended_Index;
37     begin
38       New_Line;
39       Put_Line ("Removing all elements with value of "
40               & Integer'Image (E) & "...");
41       loop
42         I := V.Find_Index (E);
43         exit when I = No_Index;
44         V.Delete (I);
45       end loop;
46     end;
47
48     --
49     -- Remove elements using a cursor
50     --
51     declare
52       E : constant Integer := 13;
53       C : Cursor;
54     begin
55       New_Line;
56       Put_Line ("Removing all elements with value of "
57               & Integer'Image (E) & "...");
58       loop
59         C := V.Find (E);
60         exit when C = No_Element;
61         V.Delete (C);
62       end loop;
63     end;
64
65     Show_Elements (V);
66   end Show_Remove_Vector_Elements;

```

Build output

```

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

```

Runtime output

```

Vector has 7 elements
Vector elements are:
- 20
- 10
- 0

```

(continues on next page)

(continued from previous page)

```

- 13
- 10
- 14
- 13

Removing all elements with value of 10...

Removing all elements with value of 13...

Vector has 3 elements
Vector elements are:
- 20
- 0
- 14

```

In this example, we remove all elements with the value 10 from the vector by retrieving their index. Likewise, we remove all elements with the value 13 by retrieving their cursor.

16.1.9 Other Operations

We've seen some operations on vector elements. Here, we'll see operations on the vector as a whole. The most prominent is the concatenation of multiple vectors, but we'll also see operations on vectors, such as sorting and sorted merging operations, that view the vector as a sequence of elements and operate on the vector considering the element's relations to each other.

We do vector concatenation using the & operator on vectors. Let's consider two vectors V1 and V2. We can concatenate them by doing `V := V1 & V2`. V contains the resulting vector.

The generic package `Generic_Sorting` is a child package of `Ada.Containers.Vectors`. It contains sorting and merging operations. Because it's a generic package, you can't use it directly, but have to instantiate it. In order to use these operations on a vector of integer values (`Integer_Vectors`, in our example), you need to instantiate it directly as a child of `Integer_Vectors`. The next example makes it clear how to do this.

After instantiating `Generic_Sorting`, we make all the operations available to us with the `use` statement. We can then call `Sort` to sort the vector and `Merge` to merge one vector into another.

The following example presents code that manipulates three vectors (V1, V2, V3) using the concatenation, sorting and merging operations:

Listing 14: show_vector_ops.adb

```

1 with Ada.Containers; use Ada.Containers;
2 with Ada.Containers.Vectors;
3
4 with Ada.Text_IO; use Ada.Text_IO;
5
6 procedure Show_Vector_Ops is
7
8     package Integer_Vectors is new Ada.Containers.Vectors
9         (Index_Type => Natural,
10          Element_Type => Integer);
11
12     package Integer_Vectors_Sorting is new Integer_Vectors.Generic_Sorting;
13
14     use Integer_Vectors;
15     use Integer_Vectors_Sorting;
16
17     procedure Show_Elements (V : Vector) is
18     begin

```

(continues on next page)

(continued from previous page)

```

19     New_Line;
20     Put_Line ("Vector has " & Count_Type'Image (V.Length) & " elements");
21
22     if not V.Is_Empty then
23         Put_Line ("Vector elements are: ");
24         for E of V loop
25             Put_Line ("- " & Integer'Image (E));
26         end loop;
27     end if;
28 end Show_Elements;
29
30 V, V1, V2, V3 : Vector;
31 begin
32     V1 := 10 & 12 & 18;
33     V2 := 11 & 13 & 19;
34     V3 := 15 & 19;
35
36     New_Line;
37     Put_Line ("---- V1 ----");
38     Show_Elements (V1);
39
40     New_Line;
41     Put_Line ("---- V2 ----");
42     Show_Elements (V2);
43
44     New_Line;
45     Put_Line ("---- V3 ----");
46     Show_Elements (V3);
47
48     New_Line;
49     Put_Line ("Concatenating V1, V2 and V3 into V:");
50
51     V := V1 & V2 & V3;
52
53     Show_Elements (V);
54
55     New_Line;
56     Put_Line ("Sorting V:");
57
58     Sort (V);
59
60     Show_Elements (V);
61
62     New_Line;
63     Put_Line ("Merging V2 into V1:");
64
65     Merge (V1, V2);
66
67     Show_Elements (V1);
68
69 end Show_Vector_Ops;

```

Build output

```

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

```

Runtime output

```
---- V1 ----  
  
Vector has 3 elements  
Vector elements are:  
- 10  
- 12  
- 18  
  
---- V2 ----  
  
Vector has 3 elements  
Vector elements are:  
- 11  
- 13  
- 19  
  
---- V3 ----  
  
Vector has 2 elements  
Vector elements are:  
- 15  
- 19  
  
Concatenating V1, V2 and V3 into V:  
  
Vector has 8 elements  
Vector elements are:  
- 10  
- 12  
- 18  
- 11  
- 13  
- 19  
- 15  
- 19  
  
Sorting V:  
  
Vector has 8 elements  
Vector elements are:  
- 10  
- 11  
- 12  
- 13  
- 15  
- 18  
- 19  
- 19  
  
Merging V2 into V1:  
  
Vector has 6 elements  
Vector elements are:  
- 10  
- 11  
- 12  
- 13  
- 18  
- 19
```

The Reference Manual requires that the worst-case complexity of a call to `Sort` be $O(N^2)$ and the average complexity be better than $O(N^2)$.

16.2 Sets

Sets are another class of containers. While vectors allow duplicated elements to be inserted, sets ensure that no duplicated elements exist.

In the following sections, we'll see operations you can perform on sets. However, since many of the operations on vectors are similar to the ones used for sets, we'll cover them more quickly here. Please refer back to the section on vectors for a more detailed discussion.

16.2.1 Initialization and iteration

To initialize a set, you can call the `Insert` procedure. However, if you do, you need to ensure no duplicate elements are being inserted: if you try to insert a duplicate, you'll get an exception. If you have less control over the elements to be inserted so that there may be duplicates, you can use another option instead:

- a version of `Insert` that returns a Boolean value indicating whether the insertion was successful;
- the `Include` procedure, which silently ignores any attempt to insert a duplicated element.

To iterate over a set, you can use a `for E of S` loop, as you saw for vectors. This gives you a reference to each element in the set.

Let's see an example:

Listing 15: `show_set_init.adb`

```

1  with Ada.Containers; use Ada.Containers;
2  with Ada.Containers.Ordered_Sets;
3
4  with Ada.Text_IO; use Ada.Text_IO;
5
6  procedure Show_Set_Init is
7
8      package Integer_Sets is new Ada.Containers.Ordered_Sets
9          (Element_Type => Integer);
10
11     use Integer_Sets;
12
13     S : Set;
14     -- Same as: S : Integer_Sets.Set;
15     C : Cursor;
16     Ins : Boolean;
17 begin
18     S.Insert (20);
19     S.Insert (10);
20     S.Insert (0);
21     S.Insert (13);
22
23     -- Calling S.Insert(0) now would raise Constraint_Error
24     -- because this element is already in the set.
25     -- We instead call a version of Insert that doesn't raise an
26     -- exception but instead returns a Boolean indicating the status
27
28     S.Insert (0, C, Ins);

```

(continues on next page)

(continued from previous page)

```

29  if not Ins then
30      Put_Line ("Inserting 0 into set was not successful");
31  end if;
32
33  -- We can also call S.Include instead
34  -- If the element is already present, the set remains unchanged
35  S.Include (0);
36  S.Include (13);
37  S.Include (14);
38
39  Put_Line ("Set has " & Count_Type'Image (S.Length) & " elements");
40
41  --
42  -- Iterate over set using for .. of loop
43  --
44  Put_Line ("Elements:");
45  for E of S loop
46      Put_Line ("- " & Integer'Image (E));
47  end loop;
48  end Show_Set_Init;

```

Build output

```

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

```

Runtime output

```

Inserting 0 into set was not successful
Set has 5 elements
Elements:
- 0
- 10
- 13
- 14
- 20

```

16.2.2 Operations on elements

In this section, we briefly explore the following operations on sets:

- Delete and Exclude to remove elements;
- Contains and Find to verify the existence of elements.

To delete elements, you call the procedure `Delete`. However, analogously to the `Insert` procedure above, `Delete` raises an exception if the element to be deleted isn't present in the set. If you want to permit the case where an element might not exist, you can call `Exclude`, which silently ignores any attempt to delete a non-existent element.

`Contains` returns a Boolean value indicating whether a value is contained in the set. `Find` also looks for an element in a set, but returns a cursor to the element or `No_Element` if the element doesn't exist. You can use either function to search for elements in a set.

Let's look at an example that makes use of these operations:

Listing 16: show_set_element_ops.adb

```

1  with Ada.Containers; use Ada.Containers;
2  with Ada.Containers.Ordered_Sets;
3
4  with Ada.Text_IO; use Ada.Text_IO;
5
6  procedure Show_Set_Element_Ops is
7
8      package Integer_Sets is new Ada.Containers.Ordered_Sets
9          (Element_Type => Integer);
10
11     use Integer_Sets;
12
13     procedure Show_Elements (S : Set) is
14     begin
15         New_Line;
16         Put_Line ("Set has " & Count_Type'Image (S.Length) & " elements");
17         Put_Line ("Elements:");
18         for E of S loop
19             Put_Line ("- " & Integer'Image (E));
20         end loop;
21     end Show_Elements;
22
23     S : Set;
24     begin
25         S.Insert (20);
26         S.Insert (10);
27         S.Insert (0);
28         S.Insert (13);
29
30         S.Delete (13);
31
32         -- Calling S.Delete (13) again raises Constraint_Error
33         -- because the element is no longer present
34         -- in the set, so it can't be deleted.
35         -- We can call V.Exclude instead:
36         S.Exclude (13);
37
38         if S.Contains (20) then
39             Put_Line ("Found element 20 in set");
40         end if;
41
42         -- Alternatively, we could use S.Find instead of S.Contains
43         if S.Find (0) /= No_Element then
44             Put_Line ("Found element 0 in set");
45         end if;
46
47         Show_Elements (S);
48     end Show_Set_Element_Ops;

```

Build output

```

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

```

Runtime output

```
Found element 20 in set
Found element 0 in set
```

```
Set has 3 elements
Elements:
- 0
- 10
- 20
```

In addition to ordered sets used in the examples above, the standard library also offers hashed sets. The Reference Manual requires the following average complexity of each operation:

Operations	Ordered_Sets	Hashed_Sets
<ul style="list-style-type: none"> • Insert • Include • Replace • Delete • Exclude • Find 	$O((\log N)^2)$ or better	$O(\log N)$
Subprogram using cursor	O(1)	O(1)

16.2.3 Other Operations

The previous sections mostly dealt with operations on individual elements of a set. But Ada also provides typical set operations: union, intersection, difference and symmetric difference. In contrast to some vector operations we've seen before (e.g. Merge), here you can use built-in operators, such as -. The following table lists the operations and its associated operator:

Set Operation	Operator
Union	or
Intersection	and
Difference	-
Symmetric difference	xor

The following example makes use of these operators:

Listing 17: show_set_ops.adb

```
1 with Ada.Containers; use Ada.Containers;
2 with Ada.Containers.Ordered_Sets;
3
4 with Ada.Text_IO; use Ada.Text_IO;
5
6 procedure Show_Set_Ops is
7
8     package Integer_Sets is new Ada.Containers.Ordered_Sets
9         (Element_Type => Integer);
10
11     use Integer_Sets;
12
13     procedure Show_Elements (S : Set) is
14     begin
15         Put_Line ("Elements:");
16         for E of S loop
17             Put_Line ("- " & Integer'Image (E));
```

(continues on next page)

```

18     end loop;
19 end Show_Elements;
20
21 procedure Show_Op (S      : Set;
22                  Op_Name : String) is
23 begin
24     New_Line;
25     Put_Line (Op_Name & "(set #1, set #2) has "
26             & Count_Type'Image (S.Length) & " elements");
27 end Show_Op;
28
29 S1, S2, S3 : Set;
30 begin
31     S1.Insert (0);
32     S1.Insert (10);
33     S1.Insert (13);
34
35     S2.Insert (0);
36     S2.Insert (10);
37     S2.Insert (14);
38
39     S3.Insert (0);
40     S3.Insert (10);
41
42     New_Line;
43     Put_Line ("---- Set #1 ----");
44     Show_Elements (S1);
45
46     New_Line;
47     Put_Line ("---- Set #2 ----");
48     Show_Elements (S2);
49
50     New_Line;
51     Put_Line ("---- Set #3 ----");
52     Show_Elements (S3);
53
54     New_Line;
55     if S3.Is_Subset (S1) then
56         Put_Line ("S3 is a subset of S1");
57     else
58         Put_Line ("S3 is not a subset of S1");
59     end if;
60
61     S3 := S1 and S2;
62     Show_Op (S3, "Intersection");
63     Show_Elements (S3);
64
65     S3 := S1 or S2;
66     Show_Op (S3, "Union");
67     Show_Elements (S3);
68
69     S3 := S1 - S2;
70     Show_Op (S3, "Difference");
71     Show_Elements (S3);
72
73     S3 := S1 xor S2;
74     Show_Op (S3, "Symmetric difference");
75     Show_Elements (S3);
76
77 end Show_Set_Ops;

```

Build output

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

Runtime output

```
---- Set #1 ----
Elements:
- 0
- 10
- 13

---- Set #2 ----
Elements:
- 0
- 10
- 14

---- Set #3 ----
Elements:
- 0
- 10

S3 is a subset of S1

Intersection(set #1, set #2) has 2 elements
Elements:
- 0
- 10

Union(set #1, set #2) has 4 elements
Elements:
- 0
- 10
- 13
- 14

Difference(set #1, set #2) has 1 elements
Elements:
- 13

Symmetric difference(set #1, set #2) has 2 elements
Elements:
- 13
- 14
```

16.3 Indefinite maps

The previous sections presented containers for elements of definite types. Although most examples in those sections presented Integer types as element type of the containers, containers can also be used with indefinite types, an example of which is the String type. However, indefinite types require a different kind of containers designed specially for them.

We'll also be exploring a different class of containers: maps. They associate a key with a specific value. An example of a map is the one-to-one association between a person and their age. If we consider a person's name to be the key, the value is the person's age.

16.3.1 Hashed maps

Hashed maps are maps that make use of a hash as a key. The hash itself is calculated by a function you provide.

In other languages

Hashed maps are similar to dictionaries in Python and hashes in Perl. One of the main differences is that these scripting languages allow using different types for the values contained in a single map, while in Ada, both the type of key and value are specified in the package instantiation and remains constant for that specific map. You can't have a map where two elements are of different types or two keys are of different types. If you want to use multiple types, you must create a different map for each and use only one type in each map.

When instantiating a hashed map from `Ada.Containers.Indefinite_Hashed_Maps`, we specify following elements:

- `Key_Type`: type of the key
- `Element_Type`: type of the element
- `Hash`: hash function for the `Key_Type`
- `Equivalent_Keys`: an equality operator (e.g. `=`) that indicates whether two keys are to be considered equal.
 - If the type specified in `Key_Type` has a standard operator, you can use it, which you do by specifying that operator as the value of `Equivalent_Keys`.

In the next example, we'll use a string as a key type. We'll use the `Hash` function provided by the standard library for strings (in the `Ada.Strings` package) and the standard equality operator.

You add elements to a hashed map by calling `Insert`. If an element is already contained in a map `M`, you can access it directly by using its key. For example, you can change the value of an element by calling `M ("My_Key") := 10`. If the key is not found, an exception is raised. To verify if a key is available, use the function `Contains` (as we've seen above in the section on sets).

Let's see an example:

Listing 18: `show_hashed_map.adb`

```
1 with Ada.Containers.Indefinite_Hashed_Maps;  
2 with Ada.Strings.Hash;  
3  
4 with Ada.Text_IO; use Ada.Text_IO;  
5  
6 procedure Show_Hashed_Map is  
7  
8   package Integer_Hashed_Maps is new
```

(continues on next page)

(continued from previous page)

```

9      Ada.Containers.Indefinite_Hashed_Maps
10     (Key_Type      => String,
11      Element_Type  => Integer,
12      Hash          => Ada.Strings.Hash,
13      Equivalent_Keys => "=");
14
15     use Integer_Hashed_Maps;
16
17     M : Map;
18     -- Same as: M : Integer_Hashed_Maps.Map;
19 begin
20     M.Include ("Alice", 24);
21     M.Include ("John", 40);
22     M.Include ("Bob", 28);
23
24     if M.Contains ("Alice") then
25         Put_Line ("Alice's age is "
26                 & Integer'Image (M ("Alice")));
27     end if;
28
29     -- Update Alice's age
30     -- Key must already exist in M.
31     -- Otherwise an exception is raised.
32     M ("Alice") := 25;
33
34     New_Line; Put_Line ("Name & Age:");
35     for C in M.Iterate loop
36         Put_Line (Key (C) & ": " & Integer'Image (M (C)));
37     end loop;
38
39 end Show_Hashed_Map;

```

Build output

```

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

```

Runtime output

```

Alice's age is 24

Name & Age:
John: 40
Bob: 28
Alice: 25

```

16.3.2 Ordered maps

Ordered maps share many features with hashed maps. The main differences are:

- A hash function isn't needed. Instead, you must provide an ordering function (< operator), which the ordered map will use to order elements and allow fast access, $O(\log n)$, using a binary search.
 - If the type specified in Key_Type has a standard < operator, you can use it in a similar way as we did for Equivalent_Keys above for hashed maps.

Let's see an example:

Listing 19: show_ordered_map.adb

```

1 with Ada.Containers.Indefinite_Ordered_Maps;
2
3 with Ada.Text_IO; use Ada.Text_IO;
4
5 procedure Show_Ordered_Map is
6
7     package Integer_Ordered_Maps is new
8         Ada.Containers.Indefinite_Ordered_Maps
9             (Key_Type      => String,
10              Element_Type => Integer);
11
12     use Integer_Ordered_Maps;
13
14     M : Map;
15 begin
16     M.Include ("Alice", 24);
17     M.Include ("John", 40);
18     M.Include ("Bob", 28);
19
20     if M.Contains ("Alice") then
21         Put_Line ("Alice's age is "
22                 & Integer'Image (M ("Alice")));
23     end if;
24
25     -- Update Alice's age
26     -- Key must already exist in M
27     M ("Alice") := 25;
28
29     New_Line; Put_Line ("Name & Age:");
30     for C in M.Iterate loop
31         Put_Line (Key (C) & ": " & Integer'Image (M (C)));
32     end loop;
33
34 end Show_Ordered_Map;
```

Build output

```

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

Runtime output


```
Alice's age is 24
```

```
Name & Age:
Alice: 25
Bob: 28
John: 40
```

You can see a great similarity between the examples above and from the previous section. In fact, since both kinds of maps share many operations, we didn't need to make extensive modifications when we changed our example to use ordered maps instead of hashed maps. The main difference is seen when we run the examples: the output of a hashed map is usually unordered, but the output of an ordered map is always ordered, as implied by its name.

16.3.3 Complexity

Hashed maps are generally the fastest data structure available to you in Ada if you need to associate heterogeneous keys to values and search for them quickly. In most cases, they are slightly faster than ordered maps. So if you don't need ordering, use hashed maps.

The Reference Manual requires the following average complexity of operations:

Operations	Ordered_Maps	Hashed_Maps
<ul style="list-style-type: none"> • Insert • Include • Replace • Delete • Exclude • Find 	$O((\log N)^2)$ or better	$O(\log N)$
Subprogram using cursor	$O(1)$	$O(1)$

STANDARD LIBRARY: DATES & TIMES

The standard library supports processing of dates and times using two approaches:

- *Calendar* approach, which is suitable for handling dates and times in general;
- *Real-time* approach, which is better suited for real-time applications that require enhanced precision — for example, by having access to an absolute clock and handling time spans. Note that this approach only supports times, not dates.

The following sections present these two approaches.

17.1 Date and time handling

The `Ada.Calendar` package supports handling of dates and times. Let's look at a simple example:

Listing 1: `display_current_time.adb`

```
1 with Ada.Calendar;           use Ada.Calendar;
2 with Ada.Calendar.Formatting; use Ada.Calendar.Formatting;
3 with Ada.Text_IO;           use Ada.Text_IO;
4
5 procedure Display_Current_Time is
6   Now : Time := Clock;
7 begin
8   Put_Line ("Current time: " & Image (Now));
9 end Display_Current_Time;
```

Build output

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

Runtime output

```
Current time: 2021-05-14 18:17:10
```

This example displays the current date and time, which is retrieved by a call to the `Clock` function. We call the function `Image` from the `Ada.Calendar.Formatting` package to get a `String` for the current date and time. We could instead retrieve each component using the `Split` function. For example:

Listing 2: display_current_year.adb

```

1  with Ada.Calendar;           use Ada.Calendar;
2  with Ada.Text_IO;           use Ada.Text_IO;
3
4  procedure Display_Current_Year is
5      Now      : Time := Clock;
6
7      Now_Year  : Year_Number;
8      Now_Month : Month_Number;
9      Now_Day   : Day_Number;
10     Now_Seconds : Day_Duration;
11  begin
12     Split (Now,
13           Now_Year,
14           Now_Month,
15           Now_Day,
16           Now_Seconds);
17
18     Put_Line ("Current year is: " & Year_Number'Image (Now_Year));
19     Put_Line ("Current month is: " & Month_Number'Image (Now_Month));
20     Put_Line ("Current day is: " & Day_Number'Image (Now_Day));
21  end Display_Current_Year;

```

Build output

```

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

```

Runtime output

```

Current year is: 2021
Current month is: 5
Current day is: 14

```

Here, we're retrieving each element and displaying it separately.

17.1.1 Delaying using date

You can delay an application so that it restarts at a specific date and time. We saw something similar in the chapter on tasking. You do this using a `delay until` statement. For example:

Listing 3: display_delay_next_specific_time.adb

```

1  with Ada.Calendar;           use Ada.Calendar;
2  with Ada.Calendar.Formatting; use Ada.Calendar.Formatting;
3  with Ada.Calendar.Time_Zones; use Ada.Calendar.Time_Zones;
4  with Ada.Text_IO;           use Ada.Text_IO;
5
6  procedure Display_Delay_Next_Specific_Time is
7      TZ : Time_Offset := UTC_Time_Offset;
8      Next : Time      := Ada.Calendar.Formatting.Time_Of
9          (Year   => 2018,
10         Month  => 5,
11         Day    => 1,

```

(continues on next page)

(continued from previous page)

```

12     Hour      => 15,
13     Minute   => 0,
14     Second   => 0,
15     Sub_Second => 0.0,
16     Leap_Second => False,
17     Time_Zone => TZ);
18
19     -- Next = 2018-05-01 15:00:00.00 (local time-zone)
20 begin
21     Put_Line ("Let's wait until...");
22     Put_Line (Image (Next, True, TZ));
23
24     delay until Next;
25
26     Put_Line ("Enough waiting!");
27 end Display_Delay_Next_Specific_Time;

```

Build output

```

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

```

Runtime output

```

Let's wait until...
2018-05-01 15:00:00.00
Enough waiting!

```

In this example, we specify the date and time by initializing `Next` using a call to `Time_Of`, a function taking the various components of a date (year, month, etc) and returning an element of the `Time` type. Because the date specified is in the past, the `delay until` statement won't produce any noticeable effect. However, if we passed a date in the future, the program would wait until that specific date and time arrived.

Here we're converting the time to the local timezone. If we don't specify a timezone, *Coordinated Universal Time* (abbreviated to UTC) is used by default. By retrieving the time offset to UTC with a call to `UTC_Time_Offset` from the `Ada.Calendar.Time_Zones` package, we can initialize `TZ` and use it in the call to `Time_Of`. This is all we need do to make the information provided to `Time_Of` relative to the local time zone.

We could achieve a similar result by initializing `Next` with a `String`. We can do this with a call to `Value` from the `Ada.Calendar.Formatting` package. This is the modified code:

Listing 4: `display_delay_next_specific_time.adb`

```

1  with Ada.Calendar;           use Ada.Calendar;
2  with Ada.Calendar.Formatting; use Ada.Calendar.Formatting;
3  with Ada.Calendar.Time_Zones; use Ada.Calendar.Time_Zones;
4  with Ada.Text_IO;           use Ada.Text_IO;
5
6  procedure Display_Delay_Next_Specific_Time is
7      TZ : Time_Offset := UTC_Time_Offset;
8      Next : Time      := Ada.Calendar.Formatting.Value
9          ("2018-05-01 15:00:00.00", TZ);
10
11     -- Next = 2018-05-01 15:00:00.00 (local time-zone)

```

(continues on next page)

(continued from previous page)

```

12 begin
13   Put_Line ("Let's wait until...");
14   Put_Line (Image (Next, True, TZ));
15
16   delay until Next;
17
18   Put_Line ("Enough waiting!");
19 end Display_Delay_Next_Specific_Time;

```

Build output

```

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

```

Runtime output

```

Let's wait until...
2018-05-01 15:00:00.00
Enough waiting!

```

In this example, we're again using TZ in the call to Value to adjust the input time to the current time zone.

In the examples above, we were delaying to a specific date and time. Just like we saw in the tasking chapter, we could instead specify the delay relative to the current time. For example, we could delay by 5 seconds, using the current time:

Listing 5: display_delay_next.adb

```

1  with Ada.Calendar;          use Ada.Calendar;
2  with Ada.Text_IO;          use Ada.Text_IO;
3
4  procedure Display_Delay_Next is
5     D : Duration := 5.0;      -- seconds
6     Now : Time := Clock;
7     Next : Time := Now + D;  -- use duration to
8                               -- specify next point in time
9
10  begin
11     Put_Line ("Let's wait "
12              & Duration'Image (D) & " seconds...");
13     delay until Next;
14     Put_Line ("Enough waiting!");
15 end Display_Delay_Next;

```

Build output

```

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

```

Runtime output

```
Let's wait 5.000000000 seconds...
Enough waiting!
```

Here, we're specifying a duration of 5 seconds in `D`, adding it to the current time from `Now`, and storing the sum in `Next`. We then use it in the `delay until` statement.

17.2 Real-time

In addition to `Ada.Calendar`, the standard library also supports time operations for real-time applications. These are included in the `Ada.Real_Time` package. This package also include a `Time` type. However, in the `Ada.Real_Time` package, the `Time` type is used to represent an absolute clock and handle a time span. This contrasts with the `Ada.Calendar`, which uses the `Time` type to represent dates and times.

In the previous section, we used the `Time` type from the `Ada.Calendar` and the `delay until` statement to delay an application by 5 seconds. We could have used the `Ada.Real_Time` package instead. Let's modify that example:

Listing 6: `display_delay_next_real_time.adb`

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Real_Time; use Ada.Real_Time;
3
4 procedure Display_Delay_Next_Real_Time is
5   D      : Time_Span := Seconds (5);
6   Next   : Time      := Clock + D;
7 begin
8   Put_Line ("Let's wait "
9             & Duration'Image (To_Duration (D)) & " seconds...");
10  delay until Next;
11  Put_Line ("Enough waiting!");
12 end Display_Delay_Next_Real_Time;
```

Build output

```

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

Runtime output

```
Let's wait 5.000000000 seconds...
Enough waiting!
```

The main difference is that `D` is now a variable of type `Time_Span`, defined in the `Ada.Real_Time` package. We call the function `Seconds` to initialize `D`, but could have gotten a finer granularity by calling `Nanoseconds` instead. Also, we need to first convert `D` to the `Duration` type using `To_Duration` before we can display it.

17.2.1 Benchmarking

One interesting application using the `Ada.Real_Time` package is benchmarking. We've used that package before in a previous section when discussing tasking. Let's look at an example of benchmarking:

Listing 7: `display_benchmarking.adb`

```

1 with Ada.Text_IO;   use Ada.Text_IO;
2 with Ada.Real_Time; use Ada.Real_Time;
3
4 procedure Display_Benchmarking is
5
6     procedure Computational_Intensive_App is
7     begin
8         delay 5.0;
9     end Computational_Intensive_App;
10
11     Start_Time, Stop_Time : Time;
12     Elapsed_Time         : Time_Span;
13
14 begin
15     Start_Time := Clock;
16
17     Computational_Intensive_App;
18
19     Stop_Time := Clock;
20     Elapsed_Time := Stop_Time - Start_Time;
21
22     Put_Line ("Elapsed time: "
23             & Duration'Image (To_Duration (Elapsed_Time))
24             & " seconds");
25 end Display_Benchmarking;

```

Build output

```

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

```

Runtime output

```
Elapsed time: 5.000169982 seconds
```

This example defines a dummy `Computational_Intensive_App` implemented using a simple delay statement. We initialize `Start_Time` and `Stop_Time` from the then-current clock and calculate the elapsed time. By running this program, we see that the time is roughly 5 seconds, which is expected due to the delay statement.

A similar application is benchmarking of CPU time. We can implement this using the `Execution_Time` package. Let's modify the previous example to measure CPU time:

Listing 8: `display_benchmarking_cpu_time.adb`

```

1 with Ada.Text_IO;   use Ada.Text_IO;
2 with Ada.Real_Time; use Ada.Real_Time;
3 with Ada.Execution_Time; use Ada.Execution_Time;
4

```

(continues on next page)

(continued from previous page)

```

5 procedure Display_Benchmarking_CPU_Time is
6
7     procedure Computational_Intensive_App is
8     begin
9         delay 5.0;
10    end Computational_Intensive_App;
11
12    Start_Time, Stop_Time : CPU_Time;
13    Elapsed_Time         : Time_Span;
14
15 begin
16     Start_Time := Clock;
17
18     Computational_Intensive_App;
19
20     Stop_Time := Clock;
21     Elapsed_Time := Stop_Time - Start_Time;
22
23     Put_Line ("CPU time: "
24              & Duration'Image (To_Duration (Elapsed_Time))
25              & " seconds");
26 end Display_Benchmarking_CPU_Time;

```

Build output

```

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

```

Runtime output

```
CPU time: 0.000062500 seconds
```

In this example, `Start_Time` and `Stop_Time` are of type `CPU_Time` instead of `Time`. However, we still call the `Clock` function to initialize both variables and calculate the elapsed time in the same way as before. By running this program, we see that the CPU time is significantly lower than the 5 seconds we've seen before. This is because the `delay` statement doesn't require much CPU time. The results will be different if we change the implementation of `Computational_Intensive_App` to use a mathematical functions in a long loop. For example:

Listing 9: `display_benchmarking_math.adb`

```

1 with Ada.Text_IO;      use Ada.Text_IO;
2 with Ada.Real_Time;    use Ada.Real_Time;
3 with Ada.Execution_Time; use Ada.Execution_Time;
4
5 with Ada.Numerics.Generic_Elementary_Functions;
6
7 procedure Display_Benchmarking_Math is
8
9     procedure Computational_Intensive_App is
10     package Funcs is new Ada.Numerics.Generic_Elementary_Functions
11     (Float_Type => Long_Long_Float);
12     use Funcs;
13
14     X : Long_Long_Float;

```

(continues on next page)

```

15  begin
16  for I in 0 .. 1_000_000 loop
17      X := Tan (Arctan
18          (Tan (Arctan
19              (Tan (Arctan
20                  (Tan (Arctan
21                      (Tan (Arctan
22                          (Tan (Arctan
23                              (0.577)))))))))));
24  end loop;
25  end Computational_Intensive_App;
26
27  procedure Benchm_Elapsed_Time is
28      Start_Time, Stop_Time : Time;
29      Elapsed_Time          : Time_Span;
30
31  begin
32      Start_Time := Clock;
33
34      Computational_Intensive_App;
35
36      Stop_Time := Clock;
37      Elapsed_Time := Stop_Time - Start_Time;
38
39      Put_Line ("Elapsed time: "
40          & Duration'Image (To_Duration (Elapsed_Time))
41          & " seconds");
42  end Benchm_Elapsed_Time;
43
44  procedure Benchm_CPU_Time is
45      Start_Time, Stop_Time : CPU_Time;
46      Elapsed_Time          : Time_Span;
47
48  begin
49      Start_Time := Clock;
50
51      Computational_Intensive_App;
52
53      Stop_Time := Clock;
54      Elapsed_Time := Stop_Time - Start_Time;
55
56      Put_Line ("CPU time: "
57          & Duration'Image (To_Duration (Elapsed_Time))
58          & " seconds");
59  end Benchm_CPU_Time;
60  begin
61      Benchm_Elapsed_Time;
62      Benchm_CPU_Time;
63  end Display_Benchmarking_Math;

```

Build output

```

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

```

Runtime output

```
Elapsed time: 1.247201114 seconds  
CPU time: 1.226435056 seconds
```

Now that our dummy `Computational_Intensive_App` involves mathematical operations requiring significant CPU time, the measured elapsed and CPU time are much closer to each other than before.

STANDARD LIBRARY: STRINGS

In previous chapters, we've seen source-code examples using the `String` type, which is a fixed-length string type — essentially, it's an array of characters. In many cases, this data type is good enough to deal with textual information. However, there are situations that require more advanced text processing. Ada offers alternative approaches for these cases:

- *Bounded strings*: similar to fixed-length strings, bounded strings have a maximum length, which is set at its instantiation. However, bounded strings are not arrays of characters. At any time, they can contain a string of varied length — provided this length is below or equal to the maximum length.
- *Unbounded strings*: similar to bounded strings, unbounded strings can contain strings of varied length. However, in addition to that, they don't have a maximum length. In this sense, they are very flexible.

The following sections present an overview of the different string types and common operations for string types.

18.1 String operations

Operations on standard (fixed-length) strings are available in the `Ada.Strings.Fixed` package. As mentioned previously, standard strings are arrays of elements of `Character` type with a *fixed-length*. That's why this child package is called `Fixed`.

One of the simplest operations provided is counting the number of substrings available in a string (`Count`) and finding their corresponding indices (`Index`). Let's look at an example:

Listing 1: `show_find_substring.adb`

```
1 with Ada.Strings.Fixed; use Ada.Strings.Fixed;
2 with Ada.Text_IO;      use Ada.Text_IO;
3
4 procedure Show_Find_Substring is
5
6     S : String := "Hello" & 3 * " World";
7     P : constant String := "World";
8     Idx : Natural;
9     Cnt : Natural;
10 begin
11     Cnt := Ada.Strings.Fixed.Count
12         (Source => S,
13          Pattern => P);
14
15     Put_Line ("String: " & S);
16     Put_Line ("Count for '" & P & "': " & Natural'Image (Cnt));
17
18     Idx := 0;
```

(continues on next page)

(continued from previous page)

```

19   for I in 1 .. Cnt loop
20       Idx := Index
21         (Source => S,
22          Pattern => P,
23          From   => Idx + 1);
24
25       Put_Line ("Found instance of '" & P & "' at position: "
26                & Natural'Image (Idx));
27   end loop;
28
29 end Show_Find_Substring;

```

Build output

```

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

```

Runtime output

```

String: Hello World World World
Count for 'World': 3
Found instance of 'World' at position: 7
Found instance of 'World' at position: 13
Found instance of 'World' at position: 19

```

We initialize the string `S` using a multiplication. Writing `"Hello" & 3 * " World"` creates the string `Hello World World World`. We then call the function `Count` to get the number of instances of the word `World` in `S`. Next we call the function `Index` in a loop to find the index of each instance of `World` in `S`.

That example looked for instances of a specific substring. In the next example, we retrieve all the words in the string. We do this using `Find-Token` and specifying whitespaces as separators. For example:

Listing 2: show_find_words.adb

```

1  with Ada.Strings;          use Ada.Strings;
2  with Ada.Strings.Fixed;   use Ada.Strings.Fixed;
3  with Ada.Strings.Maps;    use Ada.Strings.Maps;
4  with Ada.Text_IO;        use Ada.Text_IO;
5
6  procedure Show_Find_Words is
7
8     S : String := "Hello" & 3 * " World";
9     F : Positive;
10    L : Natural;
11    I : Natural := 1;
12
13    Whitespace : constant Character_Set :=
14      To_Set (' ');
15  begin
16    Put_Line ("String: " & S);
17    Put_Line ("String length: " & Integer'Image (S'Length));
18
19    while I in S'Range loop
20      Find-Token

```

(continues on next page)

(continued from previous page)

```

21     (Source => S,
22      Set    => Whitespace,
23      From   => I,
24      Test   => Outside,
25      First  => F,
26      Last   => L);
27
28     exit when L = 0;
29
30     Put_Line ("Found word instance at position "
31              & Natural'Image (F)
32              & ": '" & S (F .. L) & "'");
33     -- & "-" & F'Img & "-" & L'Img
34
35     I := L + 1;
36 end loop;
37 end Show_Find_Words;

```

Build output

```

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

```

Runtime output

```

String: Hello World World World
String length: 23
Found word instance at position 1: 'Hello'
Found word instance at position 7: 'World'
Found word instance at position 13: 'World'
Found word instance at position 19: 'World'

```

We pass a set of characters to be used as delimiters to the procedure `Find-Token`. This set is a member of the `Character_Set` type from the `Ada.Strings.Maps` package. We call the `To_Set` function (from the same package) to initialize the set to `Whitespace` and then call `Find-Token` to loop over each valid index and find the starting index of each word. We pass `Outside` to the `Test` parameter of the `Find-Token` procedure to indicate that we're looking for indices that are outside the `Whitespace` set, i.e. actual words. The `First` and `Last` parameters of `Find-Token` are output parameters that indicate the valid range of the substring. We use this information to display the string (`S (F .. L)`).

The operations we've looked at so far read strings, but don't modify them. We next discuss operations that change the content of strings:

Operation	Description
Insert	Insert substring in a string
Overwrite	Overwrite a string with a substring
Delete	Delete a substring
Trim	Remove whitespaces from a string

All these operations are available both as functions or procedures. Functions create a new string but procedures perform the operations in place. The procedure will raise an exception if the constraints of the string are not satisfied. For example, if we have a string `S` containing 10 characters, inserting a string with two characters (e.g. `"!!"`) into it produces a string containing 12 characters. Since it has a fixed length, we can't increase its size. One possible solution in this case is to

specify that truncation should be applied while inserting the substring. This keeps the length of S fixed. Let's see an example that makes use of both function and procedure versions of Insert, Overwrite, and Delete:

Listing 3: show_adapted_strings.adb

```

1  with Ada.Strings;           use Ada.Strings;
2  with Ada.Strings.Fixed;    use Ada.Strings.Fixed;
3  with Ada.Text_IO;         use Ada.Text_IO;
4
5  procedure Show_Adapted_Strings is
6
7      S : String := "Hello World";
8      P : constant String := "World";
9      N : constant String := "Beautiful";
10
11     procedure Display_Adapted_String
12     (Source : String;
13      Before : Positive;
14      New_Item : String;
15      Pattern : String)
16     is
17         S_Ins_In : String := Source;
18         S_Ovr_In : String := Source;
19         S_Del_In : String := Source;
20
21         S_Ins : String := Insert (Source, Before, New_Item & " ");
22         S_Ovr : String := Overwrite (Source, Before, New_Item);
23         S_Del : String := Trim (Delete (Source,
24                                     Before,
25                                     Before + Pattern'Length - 1),
26                                 Ada.Strings.Right);
27     begin
28         Insert (S_Ins_In, Before, New_Item, Right);
29         Overwrite (S_Ovr_In, Before, New_Item, Right);
30         Delete (S_Del_In, Before, Before + Pattern'Length - 1);
31
32         Put_Line ("Original: " & Source & "");
33
34         Put_Line ("Insert: " & S_Ins & "");
35         Put_Line ("Overwrite: " & S_Ovr & "");
36         Put_Line ("Delete: " & S_Del & "");
37
38         Put_Line ("Insert (in-place): " & S_Ins_In & "");
39         Put_Line ("Overwrite (in-place): " & S_Ovr_In & "");
40         Put_Line ("Delete (in-place): " & S_Del_In & "");
41     end Display_Adapted_String;
42
43     Idx : Natural;
44     begin
45         Idx := Index
46             (Source => S,
47              Pattern => P);
48
49         if Idx > 0 then
50             Display_Adapted_String (S, Idx, N, P);
51         end if;
52     end Show_Adapted_Strings;

```

Build output

Compile

(continues on next page)

(continued from previous page)

```

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

```

Runtime output

```

Original: 'Hello World'
Insert:   'Hello Beautiful World'
Overwrite: 'Hello Beautiful'
Delete:   'Hello'
Insert (in-place): 'Hello Beaut'
Overwrite (in-place): 'Hello Beaut'
Delete (in-place): 'Hello '

```

In this example, we look for the index of the substring `World` and perform operations on this substring within the outer string. The procedure `Display_Adapted_String` uses both versions of the operations. For the procedural version of `Insert` and `Overwrite`, we apply truncation to the right side of the string (`Right`). For the `Delete` procedure, we specify the range of the substring, which is replaced by whitespaces. For the function version of `Delete`, we also call `Trim` which trims the trailing whitespace.

18.2 Limitation of fixed-length strings

Using fixed-length strings is usually good enough for strings that are initialized when they are declared. However, as seen in the previous section, procedural operations on strings cause difficulties when done on fixed-length strings because fixed-length strings are arrays of characters. The following example shows how cumbersome the initialization of fixed-length strings can be when it's not performed in the declaration:

Listing 4: `show_char_array.adb`

```

1  with Ada.Text_IO;          use Ada.Text_IO;
2
3  procedure Show_Char_Array is
4    S : String (1 .. 15);
5    -- Strings are arrays of Character
6  begin
7    S := "Hello          ";
8    -- Alternatively:
9    --
10   -- #1:
11   --   S (1 .. 5)      := "Hello";
12   --   S (6 .. S'Last) := (others => ' ');
13   --
14   -- #2:
15   --   S := ('H', 'e', 'l', 'l', 'o', others => ' ');
16
17   Put_Line ("String: " & S);
18   Put_Line ("String Length: " & Integer'Image (S'Length));
19 end Show_Char_Array;

```

Build output

```

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

```

Runtime output

```

String: Hello
String Length: 15

```

In this case, we can't simply write `S := "Hello"` because the resulting array of characters for the `Hello` constant has a different length than the `S` string. Therefore, we need to include trailing whitespaces to match the length of `S`. As shown in the example, we could use an exact range for the initialization (`S (1 .. 5)`) or use an explicit array of individual characters.

When strings are initialized or manipulated at run-time, it's usually better to use bounded or unbounded strings. An important feature of these types is that they aren't arrays, so the difficulties presented above don't apply. Let's start with bounded strings.

18.3 Bounded strings

Bounded strings are defined in the `Ada.Strings.Bounded.Generic_Bounded_Length` package. Because this is a generic package, you need to instantiate it and set the maximum length of the bounded string. You can then declare bounded strings of the `Bounded_String` type.

Both bounded and fixed-length strings have a maximum length that they can hold. However, bounded strings are not arrays, so initializing them at run-time is much easier. For example:

Listing 5: `show_bounded_string.adb`

```

1  with Ada.Strings;          use Ada.Strings;
2  with Ada.Strings.Bounded;
3  with Ada.Text_IO;        use Ada.Text_IO;
4
5  procedure Show_Bounded_String is
6    package B_Str is new
7      Ada.Strings.Bounded.Generic_Bounded_Length (Max => 15);
8    use B_Str;
9
10   S1, S2 : Bounded_String;
11
12   procedure Display_String_Info (S : Bounded_String) is
13   begin
14     Put_Line ("String: " & To_String (S));
15     Put_Line ("String Length: " & Integer'Image (Length (S)));
16     -- String:          S'Length => ok
17     -- Bounded_String: S'Length => compilation error
18     --                  bounded strings are not arrays!
19
20     Put_Line ("Max. Length: " & Integer'Image (Max_Length));
21   end Display_String_Info;
22 begin
23   S1 := To_Bounded_String ("Hello");
24   Display_String_Info (S1);
25
26   S2 := To_Bounded_String ("Hello World");

```

(continues on next page)

(continued from previous page)

```

27   Display_String_Info (S2);
28
29   S1 := To_Bounded_String ("Something longer to say here...",
30                           Right);
31   Display_String_Info (S1);
32 end Show_Bounded_String;

```

Build output

```

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

```

Runtime output

```

String: Hello
String Length: 5
Max. Length: 15
String: Hello World
String Length: 11
Max. Length: 15
String: Something longe
String Length: 15
Max. Length: 15

```

By using bounded strings, we can easily assign to `S1` and `S2` multiple times during execution. We use the `To_Bounded_String` and `To_String` functions to convert, in the respective direction, between fixed-length and bounded strings. A call to `To_Bounded_String` raises an exception if the length of the input string is greater than the maximum capacity of the bounded string. To avoid this, we can use the truncation parameter (`Right` in our example).

Bounded strings are not arrays, so we can't use the `'Length` attribute as we did for fixed-length strings. Instead, we call the `Length` function, which returns the length of the bounded string. The `Max_Length` constant represents the maximum length of the bounded string that we set when we instantiated the package.

After initializing a bounded string, we can manipulate it. For example, we can append a string to a bounded string using `Append` or concatenate bounded strings using the `&` operator. Like so:

Listing 6: `show_bounded_string_op.adb`

```

1  with Ada.Strings;           use Ada.Strings;
2  with Ada.Strings.Bounded;
3  with Ada.Text_IO;         use Ada.Text_IO;
4
5  procedure Show_Bounded_String_Op is
6    package B_Str is new
7      Ada.Strings.Bounded.Generic_Bounded_Length (Max => 30);
8    use B_Str;
9
10   S1, S2 : Bounded_String;
11 begin
12   S1 := To_Bounded_String ("Hello");
13   -- Alternatively: A := Null_Bounded_String & "Hello";
14   Append (S1, " World");
15   -- Alternatively: Append (A, " World", Right);
16   Put_Line ("String: " & To_String (S1));

```

(continues on next page)

(continued from previous page)

```

17
18     S2 := To_Bounded_String ("Hello!");
19     S1 := S1 & " " & S2;
20     Put_Line ("String: " & To_String (S1));
21 end Show_Bounded_String_Op;

```

Build output

```

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

```

Runtime output

```

String: Hello World
String: Hello World Hello!

```

We can initialize a bounded string with an empty string using the `Null_Bounded_String` constant. Also, we can use the `Append` procedure and specify the truncation mode like we do with the `To_Bounded_String` function.

18.4 Unbounded strings

Unbounded strings are defined in the `Ada.Strings.Unbounded` package. This is *not* a generic package, so we don't need to instantiate it before using the `Unbounded_String` type. As you may recall from the previous section, bounded strings require a package instantiation.

Unbounded strings are similar to bounded strings. The main difference is that they can hold strings of any size and adjust according to the input string: if we assign, e.g., a 10-character string to an unbounded string and later assign a 50-character string, internal operations in the container ensure that memory is allocated to store the new string. In most cases, developers don't need to worry about these operations. Also, no truncation is necessary.

Initialization of unbounded strings is very similar to bounded strings. Let's look at an example:

Listing 7: `show_unbounded_string.adb`

```

1 with Ada.Strings;           use Ada.Strings;
2 with Ada.Strings.Unbounded; use Ada.Strings.Unbounded;
3 with Ada.Text_IO;         use Ada.Text_IO;
4
5 procedure Show_Unbounded_String is
6   S1, S2 : Unbounded_String;
7
8   procedure Display_String_Info (S : Unbounded_String) is
9     begin
10      Put_Line ("String: " & To_String (S));
11      Put_Line ("String Length: " & Integer'Image (Length (S)));
12    end Display_String_Info;
13 begin
14   S1 := To_Unbounded_String ("Hello");
15   -- Alternatively: A := Null_Unbounded_String & "Hello";
16   Display_String_Info (S1);
17

```

(continues on next page)

(continued from previous page)

```

18   S2 := To_Unbounded_String ("Hello World");
19   Display_String_Info (S2);
20
21   S1 := To_Unbounded_String ("Something longer to say here...");
22   Display_String_Info (S1);
23 end Show_Unbounded_String;

```

Build output

```

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

```

Runtime output

```

String: Hello
String Length: 5
String: Hello World
String Length: 11
String: Something longer to say here...
String Length: 31

```

Like bounded strings, we can assign to S1 and S2 multiple times during execution and use the `To_Unbounded_String` and `To_String` functions to convert back-and-forth between fixed-length strings and unbounded strings. However, in this case, truncation is not needed.

And, just like for bounded strings, you can use the `Append` function and the `&` operator for unbounded strings. For example:

Listing 8: show_unbounded_string_op.adb

```

1 with Ada.Strings.Unbounded; use Ada.Strings.Unbounded;
2 with Ada.Text_IO;          use Ada.Text_IO;
3
4 procedure Show_Unbounded_String_Op is
5   S1, S2 : Unbounded_String := Null_Unbounded_String;
6 begin
7   S1 := S1 & "Hello";
8   S2 := S2 & "Hello!";
9
10  Append (S1, " World");
11  Put_Line ("String: " & To_String (S1));
12
13  S1 := S1 & " " & S2;
14  Put_Line ("String: " & To_String (S1));
15 end Show_Unbounded_String_Op;

```


STANDARD LIBRARY: FILES AND STREAMS

Ada provides different approaches for file input/output (I/O):

- *Text I/O*, which supports file I/O in text format, including the display of information on the console.
- *Sequential I/O*, which supports file I/O in binary format written in a sequential fashion for a specific data type.
- *Direct I/O*, which supports file I/O in binary format for a specific data type, but also supporting access to any position of a file.
- *Stream I/O*, which supports I/O of information for multiple data types, including objects of unbounded types, using files in binary format.

This table presents a summary of the features we've just seen:

File I/O option	Format	Random access	Data types
Text I/O	text		string type
Sequential I/O	binary		single type
Direct I/O	binary	X	single type
Stream I/O	binary	X	multiple types

In the following sections, we discuss details about these I/O approaches.

19.1 Text I/O

In most parts of this course, we used the `Put_Line` procedure to display information on the console. However, this procedure also accepts a `File_Type` parameter. For example, you can select between standard output and standard error by setting this parameter explicitly:

Listing 1: show_std_text_out.adb

```
1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Std_Text_Out is
4 begin
5   Put_Line (Standard_Output, "Hello World #1");
6   Put_Line (Standard_Error, "Hello World #2");
7 end Show_Std_Text_Out;
```

Build output

```
Compile
  [Ada]      show_std_text_out.adb
Bind
```

(continues on next page)

(continued from previous page)

```

[gprbind]    show_std_text_out.bexch
[Ada]       show_std_text_out.ali
Link
[link]      show_std_text_out.adb

```

Runtime output

```

Hello World #1
Hello World #2

```

You can also use this parameter to write information to any text file. To create a new file for writing, use the `Create` procedure, which initializes a `File_Type` element that you can later pass to `Put_Line` (instead of, e.g., `Standard_Output`). After you finish writing information, you can close the file by calling the `Close` procedure.

You use a similar method to read information from a text file. However, when opening the file, you must specify that it's an input file (`In_File`) instead of an output file. Also, instead of calling the `Put_Line` procedure, you call the `Get_Line` function to read information from the file.

Let's see an example that writes information into a new text file and then reads it back from the same file:

Listing 2: show_simple_text_file_io.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Simple_Text_File_IO is
4     F      : File_Type;
5     File_Name : constant String := "simple.txt";
6 begin
7     Create (F, Out_File, File_Name);
8     Put_Line (F, "Hello World #1");
9     Put_Line (F, "Hello World #2");
10    Put_Line (F, "Hello World #3");
11    Close (F);
12
13    Open (F, In_File, File_Name);
14    while not End_Of_File (F) loop
15        Put_Line (Get_Line (F));
16    end loop;
17    Close (F);
18 end Show_Simple_Text_File_IO;

```

Build output

```

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

```

Runtime output

```

Hello World #1
Hello World #2
Hello World #3

```

In addition to the `Create` and `Close` procedures, the standard library also includes a `Reset` procedure, which, as the name implies, resets (erases) all the information from the file. For example:

Listing 3: show_text_file_reset.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Text_File_Reset is
4   F      : File_Type;
5   File_Name : constant String := "simple.txt";
6 begin
7   Create (F, Out_File, File_Name);
8   Put_Line (F, "Hello World #1");
9   Reset (F);
10  Put_Line (F, "Hello World #2");
11  Close (F);
12
13  Open (F, In_File, File_Name);
14  while not End_Of_File (F) loop
15    Put_Line (Get_Line (F));
16  end loop;
17  Close (F);
18 end Show_Text_File_Reset;

```

Build output

```

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

```

Runtime output

```
Hello World #2
```

By running this program, we notice that, although we've written the first string (Hello World #1) to the file, it has been erased because of the call to Reset.

In addition to opening a file for reading or writing, you can also open an existing file and append to it. Do this by calling the Open procedure with the Append_File option.

When calling the Open procedure, an exception is raised if the specified file isn't found. Therefore, you should handle exceptions in that context. The following example deletes a file and then tries to open the same file for reading:

Listing 4: show_text_file_input_except.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 procedure Show_Text_File_Input_Except is
4   F      : File_Type;
5   File_Name : constant String := "simple.txt";
6 begin
7   -- Open output file and delete it
8   Create (F, Out_File, File_Name);
9   Delete (F);
10
11  -- Try to open deleted file
12  Open (F, In_File, File_Name);
13  Close (F);
14 exception
15  when Name_Error =>

```

(continues on next page)

(continued from previous page)

```

16   Put_Line ("File does not exist");
17   when others =>
18     Put_Line ("Error while processing input file");
19 end Show_Text_File_Input_Except;

```

Build output

```

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

```

Runtime output

```
File does not exist
```

In this example, we create the file by calling `Create` and then delete it by calling `Delete`. After the call to `Delete`, we can no longer use the `File_Type` element. After deleting the file, we try to open the non-existent file, which raises a `Name_Error` exception.

19.2 Sequential I/O

The previous section presented details about text file I/O. Here, we discuss doing file I/O in binary format. The first package we'll explore is the `Ada.Sequential_IO` package. Because this package is a generic package, you need to instantiate it for the data type you want to use for file I/O. Once you've done that, you can use the same procedures we've seen in the previous section: `Create`, `Open`, `Close`, `Reset` and `Delete`. However, instead of calling the `Get_Line` and `Put_Line` procedures, you'd call the `Read` and `Write` procedures.

In the following example, we instantiate the `Ada.Sequential_IO` package for floating-point types:

Listing 5: show_seq_float_io.adb

```

1  with Ada.Text_IO;
2  with Ada.Sequential_IO;
3
4  procedure Show_Seq_Float_IO is
5     package Float_IO is new Ada.Sequential_IO (Float);
6     use Float_IO;
7
8     F          : Float_IO.File_Type;
9     File_Name : constant String := "float_file.bin";
10 begin
11   Create (F, Out_File, File_Name);
12   Write (F, 1.5);
13   Write (F, 2.4);
14   Write (F, 6.7);
15   Close (F);
16
17   declare
18     Value : Float;
19   begin
20     Open (F, In_File, File_Name);
21     while not End_Of_File (F) loop
22       Read (F, Value);

```

(continues on next page)

(continued from previous page)

```

23     Ada.Text_IO.Put_Line (Float'Image (Value));
24     end loop;
25     Close (F);
26     end;
27 end Show_Seq_Float_IO;
```

Build output

```

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

Runtime output

```

1.50000E+00
2.40000E+00
6.70000E+00
```

We use the same approach to read and write complex information. The following example uses a record that includes a Boolean and a floating-point value:

Listing 6: show_seq_rec_io.adb

```

1  with Ada.Text_IO;
2  with Ada.Sequential_IO;
3
4  procedure Show_Seq_Rec_IO is
5      type Num_Info is record
6          Valid : Boolean := False;
7          Value : Float;
8      end record;
9
10     procedure Put_Line (N : Num_Info) is
11     begin
12         if N.Valid then
13             Ada.Text_IO.Put_Line ("(ok,      " & Float'Image (N.Value) & ")");
14         else
15             Ada.Text_IO.Put_Line ("(not ok, -----)");
16         end if;
17     end Put_Line;
18
19     package Num_Info_IO is new Ada.Sequential_IO (Num_Info);
20     use Num_Info_IO;
21
22     F          : Num_Info_IO.File_Type;
23     File_Name : constant String := "float_file.bin";
24 begin
25     Create (F, Out_File, File_Name);
26     Write (F, (True, 1.5));
27     Write (F, (False, 2.4));
28     Write (F, (True, 6.7));
29     Close (F);
30
31     declare
32         Value : Num_Info;
33     begin
34         Open (F, In_File, File_Name);
```

(continues on next page)

(continued from previous page)

```

35     while not End_Of_File (F) loop
36         Read (F, Value);
37         Put_Line (Value);
38     end loop;
39     Close (F);
40 end;
41 end Show_Seq_Rec_IO;
```

Build output

```

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

Runtime output

```

(ok,      1.50000E+00)
(not ok,  -----)
(ok,      6.70000E+00)
```

As the example shows, we can use the same approach we used for floating-point types to perform file I/O for this record. Once we instantiate the `Ada.Sequential_IO` package for the record type, file I/O operations are performed the same way.

19.3 Direct I/O

Direct I/O is available in the `Ada.Direct_IO` package. This mechanism is similar to the sequential I/O approach just presented, but allows us to access any position in the file. The package instantiation and most operations are very similar to sequential I/O. To rewrite the `Show_Seq_Float_IO` application presented in the previous section to use the `Ada.Direct_IO` package, we just need to replace the instances of the `Ada.Sequential_IO` package by the `Ada.Direct_IO` package. This is the new source code:

Listing 7: show_dir_float_io.adb

```

1  with Ada.Text_IO;
2  with Ada.Direct_IO;
3
4  procedure Show_Dir_Float_IO is
5      package Float_IO is new Ada.Direct_IO (Float);
6      use Float_IO;
7
8      F          : Float_IO.File_Type;
9      File_Name  : constant String := "float_file.bin";
10 begin
11     Create (F, Out_File, File_Name);
12     Write (F, 1.5);
13     Write (F, 2.4);
14     Write (F, 6.7);
15     Close (F);
16
17     declare
18         Value : Float;
```

(continues on next page)

(continued from previous page)

```

19  begin
20      Open (F, In_File, File_Name);
21      while not End_Of_File (F) loop
22          Read (F, Value);
23          Ada.Text_IO.Put_Line (Float'Image (Value));
24      end loop;
25      Close (F);
26  end;
27  end Show_Dir_Float_IO;

```

Build output

```

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

```

Runtime output

```

1.50000E+00
2.40000E+00
6.70000E+00

```

Unlike sequential I/O, direct I/O allows you to access any position in the file. However, it doesn't offer an option to append information to a file. Instead, it provides an `Inout_File` mode allowing reading and writing to a file via the same `File_Type` element.

To access any position in the file, call the `Set_Index` procedure to set the new position / index. You can use the `Index` function to retrieve the current index. Let's see an example:

Listing 8: show_dir_float_in_out_file.adb

```

1  with Ada.Text_IO;
2  with Ada.Direct_IO;
3
4  procedure Show_Dir_Float_In_Out_File is
5      package Float_IO is new Ada.Direct_IO (Float);
6      use Float_IO;
7
8      F          : Float_IO.File_Type;
9      File_Name  : constant String := "float_file.bin";
10  begin
11      -- Open file for input / output
12      Create (F, Inout_File, File_Name);
13      Write (F, 1.5);
14      Write (F, 2.4);
15      Write (F, 6.7);
16
17      -- Set index to previous position and overwrite value
18      Set_Index (F, Index (F) - 1);
19      Write (F, 7.7);
20
21  declare
22      Value : Float;
23  begin
24      -- Set index to start of file
25      Set_Index (F, 1);
26

```

(continues on next page)

(continued from previous page)

```

27   while not End_Of_File (F) loop
28       Read (F, Value);
29       Ada.Text_IO.Put_Line (Float'Image (Value));
30   end loop;
31   Close (F);
32   end;
33 end Show_Dir_Float_In_Out_File;
```

Build output

```

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

Runtime output

```

1.50000E+00
2.40000E+00
7.70000E+00
```

By running this example, we see that the file contains 7.7, rather than the previous 6.7 that we wrote. We overwrote the value by changing the index to the previous position before doing another write.

In this example we used the `Inout_File` mode. Using that mode, we just changed the index back to the initial position before reading from the file (`Set_Index (F, 1)`) instead of closing the file and reopening it for reading.

19.4 Stream I/O

All the previous approaches for file I/O in binary format (sequential and direct I/O) are specific for a single data type (the one we instantiate them with). You can use these approaches to write objects of a single data type that may be an array or record (potentially with many fields), but if you need to create and process files that include different data types, or any objects of an unbounded type, these approaches are not sufficient. Instead, you should use stream I/O.

Stream I/O shares some similarities with the previous approaches. We still use the `Create`, `Open` and `Close` procedures. However, instead of accessing the file directly via a `File_Type` element, you use a `Stream_Access` element. To read and write information, you use the `'Read` or `'Write` attributes of the data types you're reading or writing.

Let's look at a version of the `Show_Dir_Float_IO` procedure from the previous section that makes use of stream I/O instead of direct I/O:

Listing 9: show_float_stream.adb

```

1 with Ada.Text_IO;
2 with Ada.Streams.Stream_IO; use Ada.Streams.Stream_IO;
3
4 procedure Show_Float_Stream is
5   F      : File_Type;
6   S      : Stream_Access;
7   File_Name : constant String := "float_file.bin";
8 begin
```

(continues on next page)

(continued from previous page)

```

9   Create (F, Out_File, File_Name);
10  S := Stream (F);
11
12  Float'Write (S, 1.5);
13  Float'Write (S, 2.4);
14  Float'Write (S, 6.7);
15
16  Close (F);
17
18  declare
19      Value : Float;
20  begin
21      Open (F, In_File, File_Name);
22      S := Stream (F);
23
24      while not End_Of_File (F) loop
25          Float'Read (S, Value);
26          Ada.Text_IO.Put_Line (Float'Image (Value));
27      end loop;
28      Close (F);
29  end;
30 end Show_Float_Stream;

```

Build output

```

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

```

Runtime output

```

1.50000E+00
2.40000E+00
6.70000E+00

```

After the call to `Create`, we retrieve the corresponding `Stream_Access` element by calling the `Stream` function. We then use this stream to write information to the file via the `'Write` attribute of the `Float` type. After closing the file and reopening it for reading, we again retrieve the corresponding `Stream_Access` element and processed to read information from the file via the `'Read` attribute of the `Float` type.

You can use streams to create and process files containing different data types within the same file. You can also read and write unbounded data types such as strings. However, when using unbounded data types you must call the `'Input` and `'Output` attributes of the unbounded data type: these attributes write information about bounds or discriminants in addition to the object's actual data.

The following example shows file I/O that mixes both strings of different lengths and floating-point values:

Listing 10: show_string_stream.adb

```

1  with Ada.Text_IO;
2  with Ada.Streams.Stream_IO; use Ada.Streams.Stream_IO;
3
4  procedure Show_String_Stream is
5      F      : File_Type;

```

(continues on next page)

(continued from previous page)

```

6   S       : Stream_Access;
7   File_Name : constant String := "float_file.bin";
8
9   procedure Output (S : Stream_Access;
10                  FV : Float;
11                  SV : String) is
12   begin
13       String'Output (S, SV);
14       Float'Output (S, FV);
15   end Output;
16
17   procedure Input_Display (S : Stream_Access) is
18       SV : String := String'Input (S);
19       FV : Float  := Float'Input (S);
20   begin
21       Ada.Text_IO.Put_Line (Float'Image (FV) & " --- " & SV);
22   end Input_Display;
23
24   begin
25       Create (F, Out_File, File_Name);
26       S := Stream (F);
27
28       Output (S, 1.5, "Hi!!");
29       Output (S, 2.4, "Hello world!");
30       Output (S, 6.7, "Something longer here...");
31
32       Close (F);
33
34       Open (F, In_File, File_Name);
35       S := Stream (F);
36
37       while not End_Of_File (F) loop
38           Input_Display (S);
39       end loop;
40       Close (F);
41
42   end Show_String_Stream;

```

Build output

```

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

```

Runtime output

```

1.50000E+00 --- Hi!!
2.40000E+00 --- Hello world!
6.70000E+00 --- Something longer here...

```

When you use Stream I/O, no information is written into the file indicating the type of the data that you wrote. If a file contains data from different types, you must reference types in the same order when reading a file as when you wrote it. If not, the information you get will be corrupted. Unfortunately, strong data typing doesn't help you in this case. Writing simple procedures for file I/O (as in the example above) may help ensuring that the file format is consistent.

Like direct I/O, stream I/O supports also allows you to access any location in the file. However, when doing so, you need to be extremely careful that the position of the new index is consistent

with the data types you're expecting.

STANDARD LIBRARY: NUMERICS

The standard library provides support for common numeric operations on floating-point types as well as on complex types and matrices. In the sections below, we present a brief introduction to these numeric operations.

20.1 Elementary Functions

The `Ada.Numerics.Elementary_Functions` package provides common operations for floating-point types, such as square root, logarithm, and the trigonometric functions (e.g., `sin`, `cos`). For example:

Listing 1: `show_elem_math.adb`

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Numerics; use Ada.Numerics;
3
4 with Ada.Numerics.Elementary_Functions;
5 use Ada.Numerics.Elementary_Functions;
6
7 procedure Show_Elem_Math is
8   X : Float;
9 begin
10  X := 2.0;
11  Put_Line ("Square root of " & Float'Image (X)
12           & " is " & Float'Image (Sqrt (X)));
13
14  X := e;
15  Put_Line ("Natural log of " & Float'Image (X)
16           & " is " & Float'Image (Log (X)));
17
18  X := 10.0 ** 6.0;
19  Put_Line ("Log_10      of " & Float'Image (X)
20           & " is " & Float'Image (Log (X, 10.0)));
21
22  X := 2.0 ** 8.0;
23  Put_Line ("Log_2      of " & Float'Image (X)
24           & " is " & Float'Image (Log (X, 2.0)));
25
26  X := Pi;
27  Put_Line ("Cos        of " & Float'Image (X)
28           & " is " & Float'Image (Cos (X)));
29
30  X := -1.0;
31  Put_Line ("Arccos    of " & Float'Image (X)
32           & " is " & Float'Image (Arccos (X)));
33 end Show_Elem_Math;
```

Build output

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

Runtime output

```
Square root of 2.00000E+00 is 1.41421E+00
Natural log of 2.71828E+00 is 1.00000E+00
Log_10      of 1.00000E+06 is 6.00000E+00
Log_2      of 2.56000E+02 is 8.00000E+00
Cos        of 3.14159E+00 is -1.00000E+00
Arccos     of -1.00000E+00 is 3.14159E+00
```

Here we use the standard `e` and `Pi` constants from the `Ada.Numerics` package.

The `Ada.Numerics.Elementary_Functions` package provides operations for the `Float` type. Similar packages are available for `Long_Float` and `Long_Long_Float` types. For example, the `Ada.Numerics.Long_Elementary_Functions` package offers the same set of operations for the `Long_Float` type. In addition, the `Ada.Numerics.Generic_Elementary_Functions` package is a generic version of the package that you can instantiate for custom floating-point types. In fact, the `Elementary_Functions` package can be defined as follows:

```
package Elementary_Functions is new
  Ada.Numerics.Generic_Elementary_Functions (Float);
```

20.2 Random Number Generation

The `Ada.Numerics.Float_Random` package provides a simple random number generator for the range between 0.0 and 1.0. To use it, declare a generator `G`, which you pass to `Random`. For example:

Listing 2: `show_float_random_num.adb`

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Numerics.Float_Random; use Ada.Numerics.Float_Random;
3
4 procedure Show_Float_Random_Num is
5   G : Generator;
6   X : Uniformly_Distributed;
7 begin
8   Reset (G);
9
10  Put_Line ("Some random numbers between "
11           & Float'Image (Uniformly_Distributed'First) & " and "
12           & Float'Image (Uniformly_Distributed'Last) & ":");
13  for I in 1 .. 15 loop
14    X := Random (G);
15    Put_Line (Float'Image (X));
16  end loop;
17 end Show_Float_Random_Num;
```

Build output

```

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

```

Runtime output

```

Some random numbers between 0.00000E+00 and 1.00000E+00:
1.48510E-01
6.80098E-02
2.53527E-01
5.64306E-01
6.62029E-01
9.94627E-01
2.19372E-01
1.26992E-01
3.31311E-01
9.32065E-01
3.87250E-01
7.29419E-01
6.04756E-02
3.31727E-01
5.55823E-01

```

The standard library also includes a random number generator for discrete numbers, which is part of the `Ada.Numerics.Discrete_Random` package. Since it's a generic package, you have to instantiate it for the desired discrete type. This allows you to specify a range for the generator. In the following example, we create an application that displays random integers between 1 and 10:

Listing 3: show_discrete_random_num.adb

```

1  with Ada.Text_IO; use Ada.Text_IO;
2  with Ada.Numerics.Discrete_Random;
3
4  procedure Show_Discrete_Random_Num is
5
6      subtype Random_Range is Integer range 1 .. 10;
7
8      package R is new Ada.Numerics.Discrete_Random (Random_Range);
9      use R;
10
11     G : Generator;
12     X : Random_Range;
13 begin
14     Reset (G);
15
16     Put_Line ("Some random numbers between "
17             & Integer'Image (Random_Range'First) & " and "
18             & Integer'Image (Random_Range'Last) & ":");
19     for I in 1 .. 15 loop
20         X := Random (G);
21         Put_Line (Integer'Image (X));
22     end loop;
23 end Show_Discrete_Random_Num;

```

Build output

```

Compile
  [Ada]          show_discrete_random_num.adb

```

(continues on next page)

(continued from previous page)

```

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

```

Runtime output

```

Some random numbers between 1 and 10:
2
3
5
6
10
3
9
2
1
3
1
3
7
6
10

```

Here, package R is instantiated with the `Random_Range` type, which has a constrained range between 1 and 10. This allows us to control the range used for the random numbers. We could easily modify the application to display random integers between 0 and 20 by changing the specification of the `Random_Range` type. We can also use floating-point or fixed-point types.

20.3 Complex Types

The `Ada.Numerics.Complex_Types` package provides support for complex number types and the `Ada.Numerics.Complex_Elementary_Functions` package provides support for common operations on complex number types, similar to the `Ada.Numerics.Elementary_Functions` package. Finally, you can use the `Ada.Text_IO.Complex_IO` package to perform I/O operations on complex numbers. In the following example, we declare variables of the `Complex` type and initialize them using an aggregate:

Listing 4: show_elem_math.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2 with Ada.Numerics; use Ada.Numerics;
3
4 with Ada.Numerics.Complex_Types;
5 use Ada.Numerics.Complex_Types;
6
7 with Ada.Numerics.Complex_Elementary_Functions;
8 use Ada.Numerics.Complex_Elementary_Functions;
9
10 with Ada.Text_IO.Complex_IO;
11
12 procedure Show_Elem_Math is
13
14     package C_IO is new Ada.Text_IO.Complex_IO (Complex_Types);
15     use C_IO;
16

```

(continues on next page)

(continued from previous page)

```

17   X, Y : Complex;
18   R, Th : Float;
19   begin
20   X := (2.0, -1.0);
21   Y := (3.0, 4.0);
22
23   Put (X);
24   Put (" * ");
25   Put (Y);
26   Put (" is ");
27   Put (X * Y);
28   New_Line;
29   New_Line;
30
31   R := 3.0;
32   Th := Pi / 2.0;
33   X := Compose_From_Polar (R, Th);
34   -- Alternatively:
35   -- X := R * Exp ((0.0, Th));
36   -- X := R * e ** Complex'(0.0, Th);
37
38   Put ("Polar form:      "
39       & Float'Image (R) & " * e**(i * "
40       & Float'Image (Th) & ")");
41   New_Line;
42
43   Put ("Modulus      of ");
44   Put (X);
45   Put (" is ");
46   Put (Float'Image (abs (X)));
47   New_Line;
48
49   Put ("Argument     of ");
50   Put (X);
51   Put (" is ");
52   Put (Float'Image (Argument (X)));
53   New_Line;
54   New_Line;
55
56   Put ("Sqrt          of ");
57   Put (X);
58   Put (" is ");
59   Put (Sqrt (X));
60   New_Line;
61   end Show_Elem_Math;

```

Build output

```

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

```

Runtime output

```

( 2.00000E+00,-1.00000E+00) * ( 3.00000E+00, 4.00000E+00) is ( 1.00000E+01, 5.
↪00000E+00)

```

(continues on next page)

(continued from previous page)

```

Polar form:      3.00000E+00 * e**(i * 1.57080E+00)
Modulus      of (-1.31134E-07, 3.00000E+00) is 3.00000E+00
Argument     of (-1.31134E-07, 3.00000E+00) is 1.57080E+00

Sqrt         of (-1.31134E-07, 3.00000E+00) is ( 1.22474E+00, 1.22474E+00)

```

As we can see from this example, all the common operators, such as `*` and `+`, are available for complex types. You also have typical operations on complex numbers, such as `Argument` and `Exp`. In addition to initializing complex numbers in the cartesian form using aggregates, you can do so from the polar form by calling the `Compose_From_Polar` function.

The `Ada.Numerics.Complex_Types` and `Ada.Numerics.Complex_Elementary_Functions` packages provide operations for the `Float` type. Similar packages are available for `Long_Float` and `Long_Long_Float` types. In addition, the `Ada.Numerics.Generic_Complex_Types` and `Ada.Numerics.Generic_Complex_Elementary_Functions` packages are generic versions that you can instantiate for custom or pre-defined floating-point types. For example:

```

with Ada.Numerics.Generic_Complex_Types;
with Ada.Numerics.Generic_Complex_Elementary_Functions;
with Ada.Text_IO.Complex_IO;

procedure Show_Elem_Math is

  package Complex_Types is new
    Ada.Numerics.Generic_Complex_Types (Float);
  use Complex_Types;

  package Elementary_Functions is new
    Ada.Numerics.Generic_Complex_Elementary_Functions (Complex_Types);
  use Elementary_Functions;

  package C_IO is new Ada.Text_IO.Complex_IO (Complex_Types);
  use C_IO;

  X, Y : Complex;
  R, Th : Float;

```

20.4 Vector and Matrix Manipulation

The `Ada.Numerics.Real_Arrays` package provides support for vectors and matrices. It includes common matrix operations such as inverse, determinant, eigenvalues in addition to simpler operators such as matrix addition and multiplication. You can declare vectors and matrices using the `Real_Vector` and `Real_Matrix` types, respectively.

The following example uses some of the operations from the `Ada.Numerics.Real_Arrays` package:

Listing 5: show_matrix.adb

```

1 with Ada.Text_IO; use Ada.Text_IO;
2
3 with Ada.Numerics.Real_Arrays;
4 use Ada.Numerics.Real_Arrays;
5
6 procedure Show_Matrix is
7
8   procedure Put_Vector (V : Real_Vector) is
9     begin

```

(continues on next page)

(continued from previous page)

```

10     Put ("    ");
11     for I in V'Range loop
12         Put (Float'Image (V (I)) & " ");
13     end loop;
14     Put_Line ("");
15 end Put_Vector;
16
17 procedure Put_Matrix (M : Real_Matrix) is
18 begin
19     for I in M'Range (1) loop
20         Put ("    ");
21         for J in M'Range (2) loop
22             Put (Float'Image (M (I, J)) & " ");
23         end loop;
24         Put_Line ("");
25     end loop;
26 end Put_Matrix;
27
28 V1      : Real_Vector := (1.0, 3.0);
29 V2      : Real_Vector := (75.0, 11.0);
30
31 M1      : Real_Matrix :=
32         ((1.0, 5.0, 1.0),
33          (2.0, 2.0, 1.0));
34 M2      : Real_Matrix :=
35         ((31.0, 11.0, 10.0),
36          (34.0, 16.0, 11.0),
37          (32.0, 12.0, 10.0),
38          (31.0, 13.0, 10.0));
39 M3      : Real_Matrix := ((1.0, 2.0),
40                          (2.0, 3.0));
41 begin
42     Put_Line ("V1");
43     Put_Vector (V1);
44     Put_Line ("V2");
45     Put_Vector (V2);
46     Put_Line ("V1 * V2 =");
47     Put_Line ("    "
48             & Float'Image (V1 * V2));
49     Put_Line ("V1 * V2 =");
50     Put_Matrix (V1 * V2);
51     New_Line;
52
53     Put_Line ("M1");
54     Put_Matrix (M1);
55     Put_Line ("M2");
56     Put_Matrix (M2);
57     Put_Line ("M2 * Transpose(M1) =");
58     Put_Matrix (M2 * Transpose (M1));
59     New_Line;
60
61     Put_Line ("M3");
62     Put_Matrix (M3);
63     Put_Line ("Inverse (M3) =");
64     Put_Matrix (Inverse (M3));
65     Put_Line ("abs Inverse (M3) =");
66     Put_Matrix (abs Inverse (M3));
67     Put_Line ("Determinant (M3) =");
68     Put_Line ("    "
69             & Float'Image (Determinant (M3)));
70     Put_Line ("Solve (M3, V1) =");

```

(continues on next page)

(continued from previous page)

```

71   Put_Vector (Solve (M3, V1));
72   Put_Line ("Eigenvalues (M3) =");
73   Put_Vector (Eigenvalues (M3));
74   New_Line;
75 end Show_Matrix;

```

Build output

```

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

```

Runtime output

```

V1
  ( 1.00000E+00  3.00000E+00 )
V2
  ( 7.50000E+01  1.10000E+01 )
V1 * V2 =
  1.08000E+02
V1 * V2 =
  ( 7.50000E+01  1.10000E+01 )
  ( 2.25000E+02  3.30000E+01 )

M1
  ( 1.00000E+00  5.00000E+00  1.00000E+00 )
  ( 2.00000E+00  2.00000E+00  1.00000E+00 )
M2
  ( 3.10000E+01  1.10000E+01  1.00000E+01 )
  ( 3.40000E+01  1.60000E+01  1.10000E+01 )
  ( 3.20000E+01  1.20000E+01  1.00000E+01 )
  ( 3.10000E+01  1.30000E+01  1.00000E+01 )
M2 * Transpose(M1) =
  ( 9.60000E+01  9.40000E+01 )
  ( 1.25000E+02  1.11000E+02 )
  ( 1.02000E+02  9.80000E+01 )
  ( 1.06000E+02  9.80000E+01 )

M3
  ( 1.00000E+00  2.00000E+00 )
  ( 2.00000E+00  3.00000E+00 )
Inverse (M3) =
  (-3.00000E+00  2.00000E+00 )
  ( 2.00000E+00 -1.00000E+00 )
abs Inverse (M3) =
  ( 3.00000E+00  2.00000E+00 )
  ( 2.00000E+00  1.00000E+00 )
Determinant (M3) =
  -1.00000E+00
Solve (M3, V1) =
  ( 3.00000E+00 -1.00000E+00 )
Eigenvalues (M3) =
  ( 4.23607E+00 -2.36068E-01 )

```

Matrix dimensions are automatically determined from the aggregate used for initialization when you don't specify them. You can, however, also use explicit ranges. For example:

```
M1      : Real_Matrix (1 .. 2, 1 .. 3) :=  
          ((1.0, 5.0, 1.0),  
           (2.0, 2.0, 1.0));
```

The `Ada.Numerics.Real_Arrays` package implements operations for the `Float` type. Similar packages are available for `Long_Float` and `Long_Long_Float` types. In addition, the `Ada.Numerics.Generic_Real_Arrays` package is a generic version that you can instantiate with custom floating-point types. For example, the `Real_Arrays` package can be defined as follows:

```
package Real_Arrays is new  
  Ada.Numerics.Generic_Real_Arrays (Float);
```


21.1 Appendix A: Generic Formal Types

The following tables contain examples of available formal types for generics:

Formal type	Actual type
Incomplete type Format: type T;	Any type
Discrete type Format: type T is (<>);	Any integer, modular or enumeration type
Range type Format: type T is range <>;	Any signed integer type
Modular type Format: type T is mod <>;	Any modular type
Floating-point type Format: type T is digits <>;	Any floating-point type
Binary fixed-point type Format: type T is delta <>;	Any binary fixed-point type
Decimal fixed-point type Format: type T is delta <> digits <>;	Any decimal fixed-point type
Definite nonlimited private type Format: type T is private;	Any nonlimited, definite type
Nonlimited Private type with discriminant Format: type T (D : DT) is private;	Any nonlimited type with discriminant
Access type Format: type A is access T;	Any access type for type T
Definite derived type Format: type T is new B;	Any concrete type derived from base type B
Limited private type Format: type T is limited private;	Any definite type, limited or not
Incomplete tagged type Format: type T is tagged;	Any concrete, definite, tagged type
Definite tagged private type Format: type T is tagged private;	Any concrete, definite, tagged type
Definite tagged limited private type Format: type T is tagged limited private;	Any concrete definite tagged type, limited or not
Definite abstract tagged private type Format: type T is abstract tagged private;	Any nonlimited, definite tagged type, abstract or concrete
Definite abstract tagged limited private type Format: type T is abstract tagged limited private;	Any definite tagged type, limited or not, abstract or concrete

continues on next page

Table 1 – continued from previous page

Formal type	Actual type
Definite derived tagged type Format: type T is new B with private;	Any concrete tagged type derived from base type B
Definite abstract derived tagged type Format: type T is abstract new B with private;	Any tagged type derived from base type B abstract or concrete
Array type Format: type A is array (R) of T;	Any array type with range R containing elements of type T
Interface type Format: type T is interface;	Any interface type T
Limited interface type Format: type T is limited interface;	Any limited interface type T
Task interface type Format: type T is task interface;	Any task interface type T
Synchronized interface type Format: type T is synchronized interface;	Any synchronized interface type T
Protected interface type Format: type T is protected interface;	Any protected interface type T
Derived interface type Format: type T is new B and I with private;	Any type T derived from base type B and interface I
Derived type with multiple interfaces Format: type T is new B and I1 and I2 with private;	Any type T derived from base type B and interfaces I1 and I2
Abstract derived interface type Format: type T is abstract new B and I with private;	Any type T derived from abstract base type B and interface I
Limited derived interface type Format: type T is limited new B and I with private;	Any type T derived from limited base type B and limited interface I
Abstract limited derived interface type Format: type T is abstract limited new B and I with private;	Any type T derived from abstract limited base type B and limited interface I
Synchronized interface type Format: type T is synchronized new SI with private;	Any type T derived from synchronized interface SI
Abstract synchronized interface type Format: type T is abstract synchronized new SI with private;	Any type T derived from synchronized interface SI

21.1.1 Indefinite version

Many of the examples above can be used for formal indefinite types:

Formal type	Actual type
Indefinite incomplete type Format: type T (<>);	Any type
Indefinite nonlimited private type Format: type T (<>) is private;	Any nonlimited type indefinite or definite
Indefinite limited private type Format: type T (<>) is limited private;	Any type, limited or not, indefinite or definite
Incomplete indefinite tagged private type Format: type T (<>) is tagged;	Any concrete tagged type, indefinite or definite
Indefinite tagged private type Format: type T (<>) is tagged private;	Any concrete, limited tagged type, indefinite or definite
Indefinite tagged limited private type Format: type T (<>) is tagged limited private;	Any concrete tagged type, limited or not, indefinite or definite
Indefinite abstract tagged private type Format: type T (<>) is abstract tagged private;	Any nonlimited tagged type, indefinite or definite, abstract or concrete
Indefinite abstract tagged limited private type Format: type T (<>) is abstract tagged limited private;	Any tagged type, limited or not, indefinite or definite abstract or concrete
Indefinite derived tagged type Format: type T (<>) is new B with private;	Any tagged type derived from base type B, indefinite or definite
Indefinite abstract derived tagged type Format: type T (<>) is abstract new B with private;	Any tagged type derived from base type B, indefinite or definite abstract or concrete

The same examples could also contain discriminants. In this case, (<>) is replaced by a list of discriminants, e.g.: (D: DT).

21.2 Appendix B: Containers

The following table shows all containers available in Ada, including their versions (standard, bounded, unbounded, indefinite):

Category	Container	Std	Bounded	Unbounded	Indefinite
Vector	Vectors	Y	Y		Y
List	Doubly_Linked_Lists	Y	Y		Y
Map	Hashed_Maps	Y	Y		Y
Map	Ordered_Maps	Y	Y		Y
Set	Hashed_Sets	Y	Y		Y
Set	Ordered_Sets	Y	Y		Y
Tree	Multiway_Trees	Y	Y		Y
Generic	Holders				Y
Queue	Synchronized_Queue_Interfaces	Y			
Queue	Synchronized_Queues		Y	Y	
Queue	Priority_Queues		Y	Y	

The following table presents the prefixing applied to the container name that depends on its version. As indicated in the table, the standard version does not have a prefix associated with it.

Version	Naming prefix
Std	
Bounded	Bounded_
Unbounded	Unbounded_
Indefinite	Indefinite_