
GNATColl Documentation

Release 27.0w

AdaCore

Apr 28, 2026

CONTENTS

1	Introduction to the GNAT Component Collection	3
2	Building GNATColl	5
2.1	Configuring the build environment	5
2.2	Building GNATColl	7
2.3	Installing GNATColl	7
3	Scripts: Embedding script languages	9
3.1	Supported languages	9
3.1.1	The Shell language	10
3.1.2	The Python language	11
3.1.3	Classes exported to all languages	12
3.2	Scripts API	12
3.2.1	Initializing the scripting module	13
3.2.2	Creating interactive consoles	15
3.2.3	Exporting classes and methods	16
3.2.4	Executing startup scripts	24
3.2.5	Multithreading applications and scripts	24
3.2.6	Debugging scripts	24
4	Traces: Logging information	25
4.1	Configuring traces	25
4.2	Using the traces module	27
4.2.1	Logging unexpected exceptions	28
4.2.2	Checking whether the handle is active	29
4.3	Log decorators	29
4.4	Defining custom stream types	31
4.5	Logging to syslog	32
4.6	Dynamically disabling features	33
5	Strings: high-performance strings	35
5.1	Small string optimization	35
5.2	Character types	36
5.3	Configuring the size of small strings	36
5.4	Task safety	37
5.5	Copy on write	37
5.6	Growth strategy	37
5.7	Substrings	38
5.8	API	39
6	Memory: Monitoring memory usage	41

7	Mmap: Reading and Writing Files	43
8	Boyer-Moore: Searching strings	47
9	Paragraph filling: formatting text	49
10	Templates: generating text	51
11	Email: Processing email messages	53
11.1	Message formats	53
11.2	Parsing messages	54
11.3	Parsing mailboxes	54
11.4	Creating messages	55
12	Ravenscar: patterns for multitasking	57
12.1	Tasks	57
12.2	Servers	57
12.3	Timers	57
13	Storage Pools: controlling memory management	59
14	VFS: Manipulating Files	61
14.1	Filesystems abstraction	61
14.1.1	file names encoding	62
14.2	Remote filesystems	62
14.2.1	Filesystem factory	63
14.2.2	Transport layer	63
14.3	Virtual files	64
15	Tribooleans: Three state logic	65
16	Geometry: primitive geometric operations	67
17	Projects: manipulating gpr files	69
18	Refcount: Reference counting	71
19	Readline: interactive command line	75
20	Config: Parsing configuration files	77
21	Pools: Controlling access to resources	79
22	JSON: handling JSON data	81
23	Iconv: Converting between character encodings	83
23.1	Compiling	83
23.2	Using GNATCOLL.Iconv	83
23.3	API	84
24	SQL: Database interface	85
24.1	Database Abstraction Layers	86
24.2	Database example	87
24.3	Database schema	87
24.4	Connecting to the database	90
24.5	Loading initial data in the database	92
24.6	Writing queries	93

24.7	Executing queries	95
24.8	Prepared queries	97
24.9	Getting results	99
24.10	Creating your own SQL types	100
24.11	Query logs	101
24.12	Writing your own cursors	102
24.13	The Object-Relational Mapping layer (ORM)	103
24.13.1	reverse relationships	107
24.14	Modifying objects in the ORM	108
24.15	Object factories in ORM	109
25	Xref: Cross-referencing source code	113
26	Xref: gnatsinspect	117
27	Terminal: controlling the console	121
27.1	Colors	121
27.2	Cursors	122
28	Promises: deferring work	123
29	Indices and tables	125
	Index	127

Warning: Now that the GNATcoll project has been split in several sub-projects, this documentation for GNATcoll is obsolete:

- for core components, see [GNATcoll: GNAT Component Collection – Core packages](#)
- for database components, see [GNATcoll – Database packages](#)
- for the iconv binding, see [GNATcoll – Iconv](#)
- for the syslog binding, see [GNATcoll – Syslog](#)
- for the readline binding, see [GNATcoll – Readline](#).

INTRODUCTION TO THE GNAT COMPONENT COLLECTION

The reusable library known as the GNAT Component Collection (GNATColl) is based on one main principle: general-purpose packages that are part of the GNAT technology should also be available to GNAT user application code. The compiler front end, the GNAT Programming Studio (GPS) Interactive Development Environment, and the GNAT Tracker web-based interface all served as sources for the components.

The GNATColl components complement the predefined Ada and GNAT libraries and deal with a range of common programming issues including string and text processing, memory management, and file handling. Several of the components are especially useful in enterprise applications.

The next chapter explains how to build and install GNATColl. The remainder of this document then describes the various components that make up the collection:

- *Scripts*: Embedding script languages
- *Traces*: Logging information
- *Memory*: Monitoring memory usage
- *Mmap*: Reading and writing files
- *Boyer-Moore*: Searching strings
- *Paragraph filling*: Formatting text
- *Templates*: Generating text
- *Email*: Processing email messages
- *Ravenscar*: Patterns for multitasking
- *Storage Pools*: Controlling memory management
- *VFS*: Manipulating files
- *Tribooleans*: Three-state logic
- *Geometry*: Primitive geometric operations
- *Projects*: Manipulating gpr files
- *Refcount*: Reference counting
- *Config*: Parsing configuration files
- *Pools*: Controlling access to resources
- *JSON*: Handling JavaScript Object Notation
- *SQL*: Database interface

BUILDING GNATCOLL

The build process is extremely flexible, allowing you to choose which modules to build, the features they should have, and various other properties. In the instructions detailed below, it is assumed that you have unpacked the GNATColl package in a temporary directory and that *installdir* is the directory in which you would like to install the selected components.

★ GNATColl requires the same version of GNAT it was released with, since it is sharing some sources with GNAT's *libgnatutil* library. If you do not have such a compiler, please contact info@adacore.com

2.1 Configuring the build environment

The first step is to configure the build environment. This is done by running the *configure* command in the root directory of the GNATColl tree.

On Windows, this requires a properly setup Cygwin environment, to provide Unix-like tools.

Some GNATColl components need access to a subset of the GNAT source files. An example is the *GNATCOLL.Projects* module, which reuses the same parser as the GNAT tools.

GNATColl will locate the needed source files in one of the following ways:

- If you have a copy of the GNAT sources, create a link called *gnat_src* that points to the directory containing those sources. This link should be created in the root GNATColl directory.
- Otherwise, recent versions of GNAT are distributed with a *libgpr* library that contains the project parser. GNATCOLL will automatically make use of it. You must use the same version of GNAT that GNATColl was released with, otherwise the sources might not be compatible. If you have an older version of GNAT, you could also chose to install *libgpr* independently.

If neither of the above is satisfied, GNATColl will not include support for *GNATCOLL.Projects*. You can also explicitly disable project support by configuring with *-disable-project*.

The *configure* command accepts a variety of arguments; the following are likely to be the most useful:

-prefix= `installdir`

This specifies the directory in which GNATColl should be installed.

-enable-shared* and *-disable-shared

If neither of these switches is specified, GNATColl will try to build both static and shared libraries (if the latter are supported on your system). The compilation needs to be done twice, since the compilation options might not be the same in both cases.

If you intend to always use static libraries, you can specify *-disable-shared*.

When you link GNATColl with your own application, the default is to link with the static libraries. You can change this default, which becomes shared libraries if you explicitly specify `-enable-shared`. However, even if the default is static libraries, you can still override this (see below the `LIBRARY_TYPE` variable).

`-with-python= 'directory '` and `-without-python`

This specifies where GNATColl should find python. For example, if the python executable is in `/usr/bin`, the *directory* to specify is `/usr`. In most cases, however, `configure` will be able to detect this automatically, so this option is only useful when python is installed some special directory. If you specify the second option, support for python will not be built in.

The switch `-with-python-exec` can also be used to specify an alternative python executable. If you pass the value `'python3'`, it will force GNATCOLL to build with support for python3, not python2.

`-enable-shared-python`

This specifies the location of the python library as *directory/lib*, which will in general be a shared library. By default, `configure` will search in a different directory of the python installation, and is more likely to find the static library instead (which makes distributing your application easier). However, whether shared or static libraries are used depends on how python was installed on your system.

`-disable-syslog`

If this switch is specified, then support for syslog (*Logging to syslog*) will not be build. This support allows sending the traces from all or part of your application to the system logger, rather than to files or *stdout*.

`-with-postgresql=<dir>` and `-without-postgresql`

GNATColl embeds a set of packages to query a database engine. The `configure` command attempts to find which systems are installed on your system, and then builds the needed support. But you can also explicitly disable such support.

If the directory in which PostgreSQL is installed contains spaces, you should use a syntax like:

```
./configure --with-postgres="/Program Files/PostgreSQL/8.4"
```

Generally speaking, we do not recommend using paths with spaces, since such a setup often introduces complications.

It is possible to link with a static library for postgres, by specifying the full path to `libpq.a`, as in:

```
./configure --with-postgres="/usr/local/lib/libpq.a"
```

However, that library depends on shared libraries `ssl` and `crypto`, so your application **is** still **not** fully linked statically.

`-with-sqlite=<dir>` and `-without-sqlite`

GNATCOLL embeds a set of packages to access sqlite database. This requires a fairly recent version of sqlite. These switches can be used to point to the sqlite install on your system. By default, GNATCOLL will recompile its own embedded version of sqlite and link statically with it, which avoids issues with shared libraries and makes sure the version is compatible with GNATCOLL's needs.

`configure` will look for the sqlite libraries in `dir/lib/`, `dir/lib64` or `dir`, in that order, and using the first directory that exists.

`-enable-gpl`

GNATCOLL provides interfaces to libraries that are licensed under the Full GNU Public License. This means that, should you choose to distribute your application to anyone, it must be free software and have a GPL-compatible license.

To avoid ambiguities, these interfaces are disabled by default, unless you provide the `-enable-gpl` switch to `configure`.

This currently only impacts *GNATCOLL.Readline*.

If all goes well (i.e. all required dependencies are found on the system), `configure` will generate a number of files, including `Makefile`, `Makefile.conf` and `gnatcoll_shared.gpr`.

2.2 Building GNATColl

If `configure` has run successfully, it generates a *Makefile* to allow you to build the rest of GNATColl. This is done by simply typing the following command:

```
make
```

Depending on the switches passed to `configure`, this will either build both static and shared libraries, or static only (see the `-disable-shared` `configure` switch).

Optionally, you can also build the examples and/or the automatic test suite, with the following commands:

```
make examples
make test
```

The latter will do a local installation of `gnatcoll` in a subdirectory called `local_install`, and use this to run the tests. This checks whether the installation of `gnatcoll` was successful.

2.3 Installing GNATColl

Installing the library is done with the following command:

```
make install
```

Note that this command does not try to recompile GNATColl, so you must build it first. This command will install both the shared and the static libraries if both were built.

As mentioned in the description of the `configure` switches, your application will by default be linked with the static library, unless you specified the `-enable-shared` switch.

However, you can always choose later which kind of library to use for GNATColl by setting the environment variable `LIBRARY_TYPE` to either `"relocatable"` or `"static"`.

Your application can now use the GNATColl code through a project file, by adding a `with` clause to `gnatcoll.gpr` or `gnatcoll_python.gpr`.

If you wish to install in a different location than was specified at `configure` time, you can override the `"prefix"` variable from the command line, for instance:

```
make prefix=/alternate/directory install
```

This does not require any recompilation.

SCRIPTS: EMBEDDING SCRIPT LANGUAGES

In a lot of contexts, you want to give the possibility to users to extend your application. This can be done in several ways: define an Ada API from which they can build dynamically loadable modules, provide the whole source code to your application and let users recompile it, interface with a simpler scripting languages,...

Dynamically loadable modules can be loaded on demand, as their name indicate. However, they generally require a relatively complex environment to build, and are somewhat less portable. But when your users are familiar with Ada, they provide a programming environment in which they are comfortable. As usual, changing the module requires recompilation, re-installation,...

Providing the source code to your application is generally even more complex for users. This requires an even more complex setup, your application is generally too big for users to dive into, and modifications done by one users are hard to provide to other users, or will be lost when you distribute a new version of your application.

The third solution is to embed one or more scripting languages in your application, and export some functions to it. This often requires your users to learn a new language, but these languages are generally relatively simple, and since they are interpreted they are easier to learn in an interactive console. The resulting scripts can easily be redistributed to other users or even distributed with future versions of your application.

The module in GNATColl helps you implement the third solution. It was used extensively in the GPS programming environment for its python interface.

 Each of the scripting language is optional

This module can be compiled with any of these languages as an optional dependency (except for the shell language, which is always built-in, but is extremely minimal, and doesn't have to be loaded at run time anyway). If the necessary libraries are found on the system, GNATColl will be build with support for the corresponding language, but your application can chose at run time whether or not to activate the support for a specific language.

 Use a scripting language to provide an automatic testing framework for your application.

The GPS environment uses python command for its *automatic test suite*, including graphical tests such as pressing on a button, selecting a menu,...

3.1 Supported languages

The module provides built-in support for several scripting languages, and other languages can “easily” be added. Your application does not change when new languages are added, since the interface to export subprograms and classes to the scripting languages is language-neutral, and will automatically export to all known scripting languages.

Support is provided for the following languages:

Shell

This is a very simple-minded scripting language, which doesn't provide flow-control instructions (*The Shell language*).

Python

Python (<http://www.python.org>) is an advanced scripting language that comes with an extensive library. It is fully object-oriented (*The Python language*).

3.1.1 The Shell language

The shell language was initially developed in the context of the GPS programming environment, as a way to embed scripting commands in XML configuration files.

In this language, you can execute any of the commands exported by the application, passing any number of arguments they need. Arguments to function calls can, but need not, be quoted. Quoting is only mandatory when they contain spaces, newline characters, or double-quotes (""). To quote an argument, surround it by double-quotes, and precede each double-quote it contains by a backslash character. Another way of quoting is similar to what python provides, which is to triple-quote the argument, i.e. surround it by "" on each side. In such a case, any special character (in particular other double-quotes or backslashes) lose their special meaning and are just taken as part of the argument. This is in particular useful when you do not know in advance the contents of the argument you are quoting:

```
Shell> function_name arg1 "arg 2" ""arg 3""
```

Commands are executed as if on a stack machine: the result of a command is pushed on the stack, and later commands can reference it using % following by a number. By default, the number of previous results that are kept is set to 9, and this can only be changed by modifying the source code for GNATColl. The return values are also modified by commands executed internally by your application, and that might have no visible output from the user's point of view. As a result, you should never assume you know what %1,... contain unless you just executed a command in the same script:

```
Shell> function_name arg1
Shell> function2_name %1
```

In particular, the %1 syntax is used when emulating object-oriented programming in the shell. A method of a class is just a particular function that contains a '.' in its name, and whose first implicit argument is the instance on which it applies. This instance is generally the result of calling a constructor in an earlier call. Assuming, for instance, that we have exported a class "Base" to the shell from our Ada core, we could use the following code:

```
Shell> Base arg1 arg2
Shell> Base.method %1 arg1 arg2
```

to create an instance and call one of its methods. Of course, the shell is not the best language for object-oriented programming, and better languages should be used instead.

When an instance has associated properties (which you can export from Ada using *Set_Property*), you access the properties by prefixing its name with "@":

```
Shell> Base arg1 arg2      # Build new instance
Shell> @id %1              # Access its "id" field
Shell> @id %1 5            # Set its "id" field
```

Some commands are automatically added to the shell when this scripting language is added to the application. These are

Function *load* (file)

Loads the content of *file* from the disk, and execute each of its lines as a Shell command. This can for instance be used to load scripts when your application is loaded

Function *echo* (arg...)

This function takes any number of argument, and prints them in the console associated with the language. By default, when in an interactive console, the output of commands is automatically printed to the console. But when you execute a script through *load* above, you need to explicitly call *echo* to make some output visible.

Function *clear_cache*

This frees the memory used to store the output of previous commands. Calling *%l* afterward will not make sense until further commands are executed.

3.1.2 The Python language

Python is an interpreted, object-oriented language. See <http://www.python.org> for more information, including tutorials, on this language.



Python support is optional in GNATColl. If it hasn't been installed on your system, GNATColl will be compiled without it, but that will not impact applications using GNATColl, since the same packages (and the same API therein) are provided in both cases. Of course, if python support wasn't compiled in, these packages will do nothing.

In addition to the API common to all languages (*Scripts API*), GNATColl also comes with a low-level interface to the python library. This interface is available in the `GNATCOLL.Python` package. In general, it is much simpler to use the common API rather than this specialized one, though, since otherwise you will need to take care of lots of details like memory management, conversion to and from python types,...



All functions exported to python are available in a specific namespace

All functions exported to python through GNATColl are available in a single python module, whose name you must specify when adding support for python. This is done to avoid namespace pollution. You can further organize the subprograms through python classes to provide more logical namespaces.

As in Ada, python lets you use named parameters in subprogram calls, and thus let's you change the order of arguments on the command line. This is fully supported by GNATColl, although your callbacks will need to specify the name of the parameters for this to work fine:

```
>>> func_name (arg1, arg2)
>>> func_name (arg2=arg2, arg1=arg1)`
```

Some commands and types are always exported by GNATColl, since they are needed by most application, or even internally by GNATColl itself.

Exception Unexpected_Exception

Exception Exception

Exception Missing_Arguments

Exception Invalid_Argument

A number of exceptions are added automatically, so that the internal state of your application is reflected in python. These are raised on unexpected uncaught Ada exceptions, when your callbacks return explicit errors, or when a function call is missing some arguments.

Function exec_in_console command

This function can be used in your script when you need to modify the contents of the python interpreter itself.

When you run a python script, all its commands (including the global variables) are within the context of the script. Therefore, you cannot affect variables which are used for instance in the rest of your application or in the python console. With this function, *command* will be executed as if it had been typed in the python console:

```
exec_in_console ("sys.ps1 = 'foo'")
=> foo> # Prompt was changed in the console
```

3.1.3 Classes exported to all languages

In addition to the functions exported by each specific scripting language, as described above, GNATColl exports the following to all the scripting languages. These are exported when your Ada code calls the Ada procedure *GNATCOLL.Scripts.Register_Standard_Classes*, which should be done after you have loaded all the scripting languages.

Class Console

Console is a name that you can choose yourself when you call the above Ada procedure. It will be assumed to be *Console* in the rest of this document.

This class provides an interface to consoles. A console is an input/output area in your application (whether it is a text area in a graphical application, or simply standard text I/O in text mode). In particular, the python standard output streams *sys.stdin*, *sys.stdout* and *sys.stderr* are redirected to an instance of that class. If you want to see python's error messages or usual output in your application, you must register that class, and define a default console for your scripting language through calls to *GNATCOLL.Scripts.Set_Default_Console*.

You can later add new methods to this class, which would be specific to your application. Or you can derive this class into a new class to achieve a similar goal.

Console.write(text)

This method writes *text* to the console associated with the class instance. See the examples delivered with GNATColl for examples on how to create a graphical window and make it into a *Console*.

Console.clear()

Clears the contents of the console.

Console.flush()

Does nothing currently, but is needed for compatibility with python. Output through *Console* instances is not buffered anyway.

Console.isatty(): Boolean

Whether the console is a pseudo-terminal. This is always wrong in the case of GNATColl.

Console.read([size]): string

Reads at most *size* bytes from the console, and returns the resulting string.

Console.readline([size]): string

Reads at most *size* lines from the console, and returns them as a single string.

3.2 Scripts API

This section will give an overview of the API used in the scripts module. The reference documentation for this API is in the source files themselves. In particular, each *.ads* file fully documents all its public API.

As described above, GNATColl contains several levels of API. In particular, it provides a low-level interface to python, in the packages *GNATCOLL.Python*. This interface is used by the rest of GNATColl, but is likely too low-level to really be convenient in your applications, since you need to take care of memory management and type conversions by yourself.

Instead, GNATColl provides a language-neutral Ada API. Using this API, it is transparent for your application whether you are talking to the Shell, to python, or to another language integrated in GNATColl. The code remains exactly the same, and new scripting languages can be added in later releases of GNATColl without requiring a change in your application. This flexibility is central to the design of GNATColl.

In exchange for that flexibility, however, there are language-specific features that cannot be performed through the GNATColl API. At present, this includes for instance exporting functions that return hash tables. But GNATColl doesn't try to export the greatest set of features common to all languages. On the contrary, it tries to fully support all the languages, and provide reasonable fallback for languages that do not support that feature. For instance, named parameters (which are a part of the python language) are fully supported, although the shell language doesn't support them. But that's an implementation detail transparent to your own application.

Likewise, your application might decide to always load the python scripting language. If GNATColl wasn't compiled with python support, the corresponding Ada function still exists (and thus your code still compiles), although of course it does nothing. But since the rest of the code is independent of python, this is totally transparent for your application.

 GNATColl comes with some examples, which you can use as a reference when building your own application. See the `<prefix>/share/examples/gnatcoll` directory.

Interfacing your application with the scripting module is a multistep process:

- You *must* **initialize** GNATColl and decide which features to load
- You *can* create an **interactive console** for the various languages, so that users can perform experiments interactively. This is optional, and you could decide to keep the scripting language has a hidden implementation detail (or just for automatic testing purposes for instance)
- You *can* **export** some classes and methods. This is optional, but it doesn't really make sense to just embed a scripting language and export nothing to it. In such a case, you might as well spawn a separate executable.
- You *can* load **start up scripts** or plug-ins that users have written to extend your application.

3.2.1 Initializing the scripting module

GNATColl must be initialized properly in order to provide added value to your application. This cannot be done automatically simply by depending on the library, since this initialization requires multiple-step that must be done at specific moments in the initialization of your whole application.

This initialization does not depend on whether you have build support for python in GNATColl. The same packages and subprograms are available in all cases, and therefore you do not need conditional compilation in your application to support the various cases.

Create the scripts repository

The type *GNATCOLL.Scripts.Scripts_Repository* will contain various variables common to all the scripting languages, as well as a list of the languages that were activated. This is the starting point for all other types, since from there you have access to everything. You will have only one variable of this type in your application, but it should generally be available from all the code that interfaces with the scripting language.

Like the rest of GNATColl, this is a tagged type, which you can extend in your own code. For instance, the GPS programming environment is organized as a kernel and several optional modules. The kernel provides the core functionality of GPS, and should be available from most functions that interface with the scripting languages. Since these functions have very specific profiles, we cannot pass additional arguments to them. One way to work around this limitation is to store the additional arguments (in this case a pointer to the kernel) in a class derived from *Scripts_Repository_Data*.

As a result, the code would look like:

```
with GNATCOLL.Scripts;  
Repo : Scripts_Repository := new Scripts_Repository_Record;
```

or, in the more complex case of GPS described above:

```
type Kernel_Scripts_Repository is new  
  Scripts_Repository_Data with record  
    Kernel : ...;  
end record;  
Repo : Scripts_Repository := new Kernel_Scripts_Repository'  
  (Scripts_Repository_Data with Kernel => ...);
```

Loading the scripting language

The next step is to decide which scripting languages should be made available to users. This must be done before any function is exported, since only functions exported after a language has been loaded will be made available in that language.

 If for instance python support was build into GNATColl, and if you decide not to make it available to users, your application will still be linked with libpython. It is therefore recommended although not mandatory to only build those languages that you will use.

This is done through a simple call to one or more subprograms. The following example registers both the shell and python languages:

```
with GNATCOLL.Scripts.Python;  
with GNATCOLL.Scripts.Shell;  
Register_Shell_Scripting (Repo);  
Register_Python_Scripting (Repo, "MyModule");
```

Procedure Register_Shell_Scripting (Repo)

This adds support for the shell language. Any class or function that is now exported through GNATColl will be made available in the shell

Procedure Register_Python_Scripting (Repo, Module_Name)

This adds support for the python language. Any class or function exported from now on will be made available in python, in the module specified by *Module_Name*

Exporting standard classes

To be fully functional, GNATColl requires some predefined classes to be exported to all languages (*Classes exported to all languages*). For instance, the *Console* class is needed for proper interactive with the consoles associated with each language.

These classes are created with the following code:

```
Register_Standard_Classes (Repo, "Console");
```

This must be done only after all the scripting languages were loaded in the previous step, since otherwise the new classes would not be visible in the other languages.

Procedure Register_Standard_Classes(Repo,Console_Class)

The second parameter *Console_Class* is the name of the class that is bound to a console, and thus provides

input/output support. You can chose this name so that it matches the classes you intend to export later on from your application.

3.2.2 Creating interactive consoles

The goal of the scripting module in GNATColl is to work both in text-only applications and graphical applications. However, in both cases applications will need a way to capture the output of scripting languages and display them to the user (at least for errors, to help debugging scripts), and possibly emulate input when a script is waiting for such input.

GNATColl solved this problem by using an abstract class *GNATCOLL.Scripts.Virtual_Console_Record* that defines an API for these consoles. This API is used throughout *GNATCOLL.Scripts* whenever input or output has to be performed.

 The `examples/` directory in the GNATColl package shows how to implement a console in text mode and in graphical mode.

If you want to provide feedback or interact with users, you will need to provide an actual implementation for these *Virtual_Console*, specific to your application. This could be a graphical text window, or based on *Ada.Text_IO*. The full API is fully documented in `gnatcoll-scripts.ads`, but here is a list of the main subprograms that need to be overridden.

Virtual_Console.Insert_Text (Txt)

Virtual_Console.Insert_Log (Txt)

Virtual_Console.Insert_Error (Txt)

These are the various methods for doing output. Error messages could for instance be printed in a different color. Log messages should in general be directed elsewhere, and not be made visible to users unless in special debugging modes.

Virtual_Console.Insert_Prompt (Txt)

This method must display a prompt so that the user knows input is expected. Graphical consoles will in general need to remember where the prompt ended so that they also know where the user input starts

Virtual_Console.Set_As_Default_Console (Script)

This method is called when the console becomes the default console for a scripting language. They should in general keep a pointer on that language, so that when the user presses `enter` they know which language must execute the command

Virtual_Console.Read (Size, Whole_Line) : String

Read either several characters or whole lines from the console. This is called when the user scripts read from their stdin.

Virtual_Console.Set_Data_Primitive (Instance)

Virtual_Console.Get_Instance : Console

These two methods are responsible for storing an instance of *Console* into a *GNATCOLL.Scripts.Class_Instance*. Such an instance is what the user manipulates from his scripting language. But when he executes a method, the Ada callback must know how to get the associated *Virtual_Console* back to perform actual operations on it.

These methods are implemented using one of the *GNATCOLL.Scripts.Set_Data* and *GNATCOLL.Scripts.Get_Data* operations when in text mode.

Once you have created one or more of these console, you can set them as the default console for each of the scripting languages. This way, any input/output done by scripts in this language will interact with that console, instead of being discarded. This is done through code similar to:

```
Console := GtkConsole.Create (...);
Set_Default_Console
  (Lookup_Scripting_Language (Repo, "python"),
   Virtual_Console (Console));
```

Creating a new instance of *Console*, although allowed, will by default create an unusable console. Indeed, depending on your application, you might want to create a new window, reuse an existing one, or do many other things when the user does:

```
c = Console()
```

As a result, GNATColl does not try to guess the correct behavior, and thus does not export a constructor for the console. So in the above python code, the default python constructor is used. But this constructor does not associate *c* with any actual *Virtual_Console*, and thus any call to a method of *c* will result in an error.

To make it possible for users to create their own consoles, you need to export a *Constructor_Method* (see below) for the *Console* class. In addition to your own processing, this constructor needs also to call:


```
declare
  Inst : constant Class_Instance := Nth_Arg (Data, 1);
begin
  C := new My_Console_Record;  -- or your own type
  GNATCOLL.Scripts.Set_Data (Inst, C);
end
```

3.2.3 Exporting classes and methods

Once all scripting languages have been loaded, you can start exporting new classes and functions to all the scripting languages. It is important to realize that through a single Ada call, they are exported to all loaded scripting languages, without further work required on your part.

Classes diagram

The following diagram shows the dependencies between the major data types defined in `GNATCOLL.Scripts`. Most of these are abstract classes that are implemented by the various scripting languages. Here is a brief description of the role of each type:



Class *Scripts_Repository*

As we have seen before, this is a type of which there is a single instance in your whole application, and whose main role is to give access to each of the scripting languages (*Lookup_Scripting_Language* function), and to make it possible to register each exported function only once (it then takes care of exporting it to each scripting language).

Class *Scripting_Language*

Instances of this type represent a specific language. It provides various operations to export subprograms, execute commands, create the other types described below,... There should exist a single instance of this class per supported language.

This class interacts with the script interpreter (for instance python), and all code executed in python goes through this type, which then executes your Ada callbacks to perform the actual operation.

It is also associated with a default console, as described above, so that all input and output of the scripts can be made visible to the user.

Class *Callback_Data*

This type is an opaque tagged type that provides a language-independent interface to the scripting language. It gives for instance access to the various parameters passed to your subprogram (*Nth_Arg* functions), allows you to set the return value (*Set_Return_Value* procedure), or raise exceptions (*Set_Error_Msg* procedure),...

Record Class *Type*

This type is not tagged, and cannot be extended. It basically represents a class in any of the scripting languages, and is used to create new instances of that class from Ada.

Class *Class_Instance*

A class instance represents a specific instance of a class. In general, such an instance is strongly bound to an instance of an Ada type. For instance, if you have a *Foo* type in your application that you wish to export, you would create a *Class_Type* called "Foo", and then the user can create as many instances as he wants of that class, each of which is associated with different values of *Foo* in Ada.

Another more specific example is the predefined *Console* class. As we have seen before, this is a *Virtual_Console* in Ada. You could for instance have two graphical windows in your application, each of which is a *Virtual_Console*. In the scripting language, this is exported as a class named *Console*. The user can create two instances of those, each of which is associated with one of your graphical windows. This way, executing *Console.write* on these instances would print the string on their respective graphical window.

Some scripting languages, in particular python, allow you to store any data within the class instances. In the example above, the user could for instance store the time stamp of the last output in each of the instances. It is

therefore important that, as much as possible, you always return the same *Class_Instance* for a given Ada object. See the following python example:

```
myconsole = Console ("title") # Create new console
myconsole.mydata = "20060619" # Any data, really
myconsole = Console ("title2") # Create another window
myconsole = Console ("title") # Must be same as first,
print myconsole.mydata # so that this prints "20060619"
```

Class Instance_Property

As we have seen above, a *Class_Instance* is associated in general with an Ada object. This *Instance_Property* tagged type should be extended for each Ada type you want to be able to store in a *Class_Instance*. You can then use the *Set_Data* and *Get_Data* methods of the *Class_Instance* to get and retrieve that associated Ada object.

Class Subprogram_Record

This class represents a callback in the scripting language, that is some code that can be executed when some conditions are met.

The exact semantic here depends on each of the programming languages. For instance, if you are programming in python, this is the name of a python method to execute. If you are programming in shell, this is any shell code.

The idea here is to blend in as smoothly as possible with the usual constructs of each language. For instance, in python one would prefer to write the second line rather than the third:

```
def on_exit():
    pass
set_on_exit_callback(on_exit) # Yes, python style
set_on_exit_callback("on_exit") # No
```

The last line (using a string as a parameter) would be extremely unusual in python, and would for instance force you to qualify the subprogram name with the name of its namespace (there would be no implicit namespace resolution).

To support this special type of parameters, the *Subprogram_Record* type was created in Ada.

Although the exact way they are all these types are created is largely irrelevant to your specific application in general, it might be useful for you to override part of the types to provide more advanced features. For instance, GPS redefines its own Shell language, that has basically the same behavior as the Shell language described above but whose *Subprogram_Record* in fact execute internal GPS actions rather than any shell code.

Exporting functions

All functions that you export to the scripting languages will result in a call to an Ada subprogram from your own application. This subprogram must have the following profile:

```
procedure Handler
  (Data    : in out Callback_Data'Class;
   Command : String);
```

The first parameter *Data* gives you access to the parameters of the subprogram as passed from the scripting language, and the second parameter *Command* is the name of the command to execute. The idea behind this second parameter is that a single Ada procedure might handle several different script function (for instance because they require common actions to be performed).

Register_Command (Repo,Command,Min_Args,Max_Args,Handler)

Each of the shell functions is then exported through a call to *Register_Command*. In its simplest form, this procedure takes the following arguments. *Repo* is the scripts repository, so that the command is exported to all

the scripting languages. *Command* is the name of the command. *Min_Args* and *Max_Args* are the minimum and maximum number of arguments. Most language allow option parameters, and this is how you specify them. *Handler* is the Ada procedure to call to execute the command.

Here is a simple example. It implements a function called *Add*, which takes two integers in parameter, and returns their sum:

```
Arg1_C : aliased constant String := "arg1";
Arg2_C : aliased constant String := "arg2";

procedure Sum
  (Data : in out Callback_Data'Class;
   Command : String)
is
  Arg1, Arg2 : Integer;
begin
  Name_Parameters ((1 => Arg1_C'Access, 2 => Arg2_C'Access));
  Arg1 := Nth_Arg (Data, 1);
  Arg2 := Nth_Arg (Data, 2);
  Set_Return_Value (Data, Arg1 + Arg2);
end Sum;

Register_Command (Repo, "sum", 2, 2, Sum'Access);
```

This is not the most useful function to export! Still, it illustrates a number of important concepts.

Automatic parameters types

When the command is registered, the number of arguments is specified. This means that GNATColl will check on its own whether the right number of arguments is provided. But the type of these arguments is not specified. Instead, your callback should proceed as if they were correct, and try to retrieve them through one of the numerous *Nth_Arg* functions. In the example above, we assume they are integer. But if one of them was passed as a string, an exception would be raised and sent back to the scripting language to display a proper error message to the user. You have nothing special to do here.

Support for named parameters

Some languages (especially python) support named parameters, ie parameters can be specified in any order on the command line, as long as they are properly identified (very similar to Ada's own capabilities). In the example above, the call to *Name_Parameters* is really optional, but adds this support for your own functions as well. You just have to specify the name of the parameters, and GNATColl will then ensure that when you call *Nth_Arg* the parameter number 1 is really "arg1". For scripting languages that do not support named parameters, this has no effect.

Your code can then perform as complex a code as needed, and finally return a value (or not) to the scripting language, through a call to *Set_Return_Value*.

After the above code has been executed, your users can go to the python console and type for instance:

```
from MyModule import *      # MyModule is the name we declared above
print sum (1,2)
=> 3
print sum ()
=> Error: Wrong number of parameters
```

(continues on next page)

(continued from previous page)

```
print sum ("1", 2)
=> Error: Parameter 1 should be an integer
print sum (arg2=2, arg1=1)
=> 3
```

Exporting classes

Whenever you want to make an Ada type accessible through the scripting languages, you should export it as a class. For object-oriented languages, this would map to the appropriate concept. For other languages, this provides a namespace, so that each method of the class now takes an additional first parameter which is the instance of the class, and the name of the method is prefixed by the class name.

Creating a new class is done through a call to *New_Class*, as shown in the example below:

```
MyClass : Class_Type;
MyClass := GNATCOLL.Scripts.New_Class (Repo, "MyClass");
```

At this stage, nothing is visible in the scripting language, but all the required setup has been done internally so that you can now add methods to this class.

You can then register the class methods in the same way that you registered functions. An additional parameter *Class* exists for *Register_Command*. A method is really just a standard function that has an implicit first parameter which is a *Class_Instance*. This extra parameter should not be taken into account in *Min_Args* and *Max_Args*. You can also declare the method as a static method, ie one that doesn't take this extra implicit parameter, and basically just uses the class as a namespace.

Some special method names are available. In particular, *Constructor_Method* should be used for the constructor of a class. It is a method that receives, as its first argument, a class instance that has just been created. It should associate that instance with the Ada object it represents.

Here is a simple example that exports a class. Each instance of this class is associated with a string, passed in parameter to the constructor. The class has a single method *print*, which prints its string parameter prefixed by the instance's string. To start with, here is a python example on what we want to achieve:

```
c1 = MyClass ("prefix1")
c1.print ("foo")
=> "prefix1 foo"
c2 = MyClass () # Using a default prefix
c2.print ("foo")
=> "default foo"
```

Here is the corresponding Ada code:

```
with GNATCOLL.Scripts.Impl;
procedure Handler
  (Data : **in out** Callback_Data'Class; Command : String)
is
  Inst : Class_Instance := Nth_Arg (Data, 1, MyClass);
begin
  if Command = Constructor_Method then
    Set_Data (Inst, MyClass, Nth_Arg (Data, 2, "default"));
  elsif Command = "print" then
    Insert_Text
```

(continues on next page)

(continued from previous page)

```

        (Get_Script (Data), null,
         String'(Get_Data (Inst)) & " " & Nth_Arg (Data, 2));
    end if;
end Handler;

Register_Command
  (Repo, Constructor_Method, 0, 1, Handler'Access, MyClass);
Register_Command
  (Repo, "print", 1, 1, Handler'Access, MyClass);

```

This example also demonstrates a few concepts: the constructor is declared as a method that takes one optional argument. The default value is in fact passed in the call to *Nth_Arg* and is set to “default”. In the handler, we know there is always a first argument which is the instance on which the method applies. The implementation for the constructor stores the prefix in the instance itself, so that several instances can have different prefixes (we can’t use global variables, of course, since we don’t know in advance how many instances will exist). The implementation for *print* inserts code in the default console for the script (we could of course use *Put_Line* or any other way to output data), and computes the string to output by concatenating the instance’s prefix and the parameter to *print*.

Note that *Set_Data* and *Get_Data* take the class in parameter, in addition to the class instance. This is needed for proper handling of multiple inheritance: say we have a class *C* that extends two classes *A* and *B*. The Ada code that deals with *A* associates an integer with the class instance, whereas the code that deals with *B* associates a string. Now, if you have an instance of *C* but call a method inherited from *A*, and if *Get_Data* didn’t specify the class, there would be a risk that a string would be returned instead of the expected integer. In fact, the proper solution here is that both *A* and *B* store their preferred data at the same time in the instances, but only fetch the one they actually need. Therefore instances of *C* are associated with two datas.

Here is a more advanced example that shows how to export an Ada object. Let’s assume we have the following Ada type that we want to make available to scripts:

```

type MyType is record
  Field : Integer;
end record;

```

As you can see, this is not a tagged type, but could certainly be. There is of course no procedure *Set_Data* in GNATCOLL. Scripts that enables us to store *MyType* in a *Class_Instance*. This example shows how to write such a procedure. The rest of the code would be similar to the first example, with a constructor that calls *Set_Data*, and methods that call *Get_Data*:

```

type MyPropsR is new Instance_Property_Record with record
  Val : MyType;
end record;
type MyProps is access all MyPropsR'Class;

procedure Set_Data
  (Inst : Class_Instance; Val : MyType)
is
begin
  Set_Data (Inst, Get_Name (MyClass), MyPropsR'(Val => Val));
end Set_Data;

function Get_Data (Inst : Class_Instance) return MyType is
  Data : MyProps := MyProps (Instance_Property'
    (Get_Data (Inst, Get_Name (MyClass))));

```

(continues on next page)

(continued from previous page)

```
begin
  return Data.Val;
end Get_Data;
```

Several aspects worth noting in this example. Each data is associated with a name, not a class as in the previous example. That's in fact the same thing, and mostly for historical reasons. We have to create our own instance of *Instance_Property_Record* to store the data, but the implementation presents no special difficulty. In fact, we don't absolutely need to create *Set_Data* and *Get_Data* and could do everything inline in the method implementation, but it is cleaner this way and easier to reuse.

GNATColl is fully responsible for managing the lifetime of the data associated with the class instances and you can override the procedure *Destroy* if you need special memory management.

Reusing class instances

We mentioned above that it is more convenient for users of your exported classes if you always return the same class instance for the same Ada object (for instance a graphical window should always be associated with the same class instance), so that users can associate their own internal data with them.

GNATColl provides a few types to facilitate this. In passing, it is worth noting that in fact the Ada objects will be associated with a single instance *per scripting language*, but each language has its own instance. Data is not magically transferred from python to shell!

You should store the list of associated instances with your object. The type *GNATCOLL.Scripts.Instance_List_Access* is meant for that purpose, and provides two *Set* and *Get* primitives to retrieve existing instances.

The final aspect to consider here is how to return existing instances. This cannot be done from the constructor method, since when it is called it has already received the created instance (this is forced by python, and was done the same for other languages for compatibility reasons). There are two ways to work around that limitation:

- Static *get* methods

With each of your classes, you can export a static method generally called *get* that takes in parameter a way to identify an existing instance, and either return it or create a new one. It is also recommended to disable the constructor, ie force it to raise an error. Let's examine the python code as it would be used:

```
ed = Editor ("file.adb") # constructor
=> Error, cannot construct instances
ed = Editor.get ("file.adb")
=> Create a new instance
ed2 = Editor.get ("file.adb")
=> Return existing instance
ed == ed2
=> True
```

The corresponding Ada code would be something like:

```
type MyType is record
  Val : Integer;
  Inst : Instance_List_Access;
end record;
type MyTypeAccess is access all MyType;
procedure Handler
  (Data : in out Callback_Data'Class; Cmd : String)
is
```

(continues on next page)

(continued from previous page)

```

Inst : Class_Instance;
Tmp  : MyTypeAccess;
begin
  if Cmd = Constructor_Method then
    Set_Error_Msg (Data, "cannot construct instances");
  elsif Cmd = "get" then
    Tmp := check_if_exists (Nth_Arg (Data, 1));
    if Tmp = null then
      Tmp := create_new_mytype (Nth_Arg (Data, 1));
      Tmp.Inst := new Instance_List;
    end if;
    Inst := Get (Tmp.Inst.all, Get_Script (Data));
    if Inst = No_Class_Instance then
      Inst := New_Instance (Get_Script (Data), MyClass);
      Set (Tmp.Inst.all, Get_Script (Data), Inst);
      Set_Data (Inst, Tmp);
    end if;
    Set_Return_Value (Data, Inst);
  end if;
end Handler;

```

- Factory classes

The standard way to do this in python, which applies to other languages as well, is to use the Factory design pattern. For this, we need to create one class (*MyClassImpl*) and one factory function (*MyClass*).

The python code now looks like:

```

ed = MyClass ("file.adb") # Create new instance
=> ed is of type MyClassImpl
ed = MyClass ("file.adb") # return same instance
ed.do_something()

```

It is important to realize that in the call above, we are not calling the constructor of a class, but a function. At the Ada level, the function has basically the same implementation as the one we gave for *get* above. But the python code looks nicer because we do not have these additional *.get()* calls. The name of the class *MyClassImpl* doesn't appear anywhere in the python code, so this is mostly transparent.

However, if you have more than one scripting language, in particular for the shell, the code looks less nice in this case:

```

MyClass "file.adb"
=> <MyClassImpl_Instance_0x12345>
MyClassImpl.do_something %1

```

and the new name of the class is visible in the method call.

3.2.4 Executing startup scripts

The final step in starting up your application is to load extensions or plug-ins written in one of the scripting languages. There is not much to be said here, except that you should use the *GNATCOLL.Scripts.Execute_File* procedure to do so.

3.2.5 Multithreading applications and scripts

Python itself is not thread-safe. So a single thread can call the python C API at a time. To enforce this, the python interpreter provides a global interpreter lock, which you must acquire before calling the C API, and release when you are done. To simulate multitasking, the python interpreter will in fact release and reacquire the lock every 100 micro-instructions (opcodes in the python virtual machine), to give a chance to run to other tasks. So this is preemptive multitasking.

The threads that are created in Ada that do not need access to python do not need any special handling. However, those that need access to python must make a special function call before they first call the python C API, so that python can create a thread-specific data for them.

GNATCOLL.Scripts.Python contains a number of subprograms to interact with the global interpreter lock of the python engine. The initialization of your application needs to do two extra calls:

```
Register_Python_Scripting (...);
Initialize_Threads_Support;    -- Also acquires the lock
Begin_Allow_Threads;          -- Releases the lock
```

Whenever a task needs to execute python commands (or basically use any subprogram from *GNATCOLL.Scripts*, it needs to do the following:

```
Ensure_Thread_State;    -- Block all python threads
... access to python C API as usual
Begin_Allow_Threads;    -- Let other python threads run
```

In some cases, the simplest is to get the lock at the beginning of the task, and release it when done. This assumes the task executes fast enough. In other cases, you will need finer grain control over the lock.

3.2.6 Debugging scripts

GNATColl provides a convenient hook to debug your script. By default, a script (python for instance) will call your Ada callback, which might raise errors. Most of the time, the error should indeed be reported to the user, and you can thus raise a standard exception, or call *Set_Error_Msg*.

But if you wish to know which script was executing the command, it is generally not doable. You can however activate a trace (*Traces: Logging information*) called “*PYTHON.TB*” (for “traceback”), which will output the name of the command that is being executed, as well as the full traceback within the python scripts. This will help you locate which script is raising an exception.

TRACES: LOGGING INFORMATION

Most applications need to log various kinds of information: error messages, information messages or debug messages among others. These logs can be displayed and stored in a number of places: standard output, a file, the system logger, an application-specific database table,...

The package `GNATCOLL.Traces` addresses the various needs, except for the application-specific database, which of course is specific to your business and needs various custom fields in any case, which cannot be easily provided through a general interface.

This module is organized around two tagged types (used through access types, in fact, so the latter are mentioned below as a shortcut):

Trace_Handle

This type defines a handle (similar to a file descriptor in other contexts) which is latter used to output messages. An application will generally define several handles, which can be enabled or disabled separately, therefore limiting the amount of logging.

Trace_Stream

Streams are the ultimate types responsible for the output of the messages. One or more handles are associated with each stream. The latter can be a file, the standard output, a graphical window, a socket,... New types of streams can easily be defined in your application.

4.1 Configuring traces

As mentioned above, an application will generally create several *Trace_Handle* (typically one per module in the application). When new features are added to the application, the developers will generally need to add lots of traces to help investigate problems once the application is installed at a customer's site. The problem here is that each module might output a lot of information, thus confusing the logs; this also does not help debugging.

The `GNATCOLL.Traces` package allows the user to configure which handles should actually generate logs, and which should just be silent and not generate anything. Depending on the part of the application that needs to be investigated, one can therefore enable a set of handles or another, to be able to concentrate on that part of the application.

This configuration is done at two levels:

- either in the source code itself, where some *trace_handle* might be disabled or enabled by default. This will be described in more details in later sections.
- or in a configuration file which is read at runtime, and overrides the defaults set in the source code.

The configuration file is found in one of three places, in the following order:

- The file name is specified in the source code in the call to *Parse_Config_File*.
- If no file name was specified in that call, the environment variable `ADA_DEBUG_FILE` might point to a configuration file.

- If the above two attempts did not find a suitable configuration file, the current directory is searched for a file called `.gnatdebug`. Finally, the user's home directory will also be searched for that file.

In all cases, the format of the configuration file is the same. Its goal is to associate the name of a *trace_handle* with the name of a *trace_stream* on which it should be displayed.

Streams are identified by a name. You can provide additional streams by creating a new tagged object (*Defining custom stream types*). Here are the various possibilities to reference a stream:

“name”

where name is a string made of letters, digits and slash (`'/'`) characters. This is the name of a file to which the traces should be redirected. The previous contents of the file is discarded. If the name of the file is a relative path, it is relative to the location of the configuration file, not necessarily to the current directory when the file is parsed. In the file name, `$$` is automatically replaced by the process number. `$D` is automatically replaced by the current date. `$T` is automatically replaced by the current date and time. Other patterns of the form `$name`, `$(name)`, or `$(name)` are substituted with the value of the named environment variable, if it exists. If `>>` is used instead of `>` to redirect to that stream, the file is appended to, instead of truncated.

“&1”

This syntax is similar to the one used on Unix shells, and indicates that the output should be displayed on the standard output for the application. If the application is graphical, and in particular on Windows platforms, it is possible that there is no standard output!

“&2”

Similar to the previous one, but the output is sent to standard error.

“&syslog”

Logging to syslog.

Comments in a configuration file must be on a line of their own, and start with `–`. Empty lines are ignored. The rest of the lines represent configurations, as in:

- If a line contains the single character `+`, it activates all *trace_handle* by default. This means the rest of the configuration file should disable those handles that are not needed. The default is that all handles are disabled by default, and the configuration file should activate the ones it needs. The Ada source code can change the default status of each handles, as well
- If the line starts with the character `>`, followed by a stream name (as defined above), this becomes the default stream. All handles will be displayed on that stream, unless otherwise specified. If the stream does not exist, it defaults to standard output.
- Otherwise, the first token on the line is the name of a handle. If that is the only element on the line, the handle is activated, and will be displayed on the default stream.

Otherwise, the next element on the line should be a `=` sign, followed by either `yes` or `no`, depending on whether the handle should resp. be enabled or disabled.

Finally, the rest of the line can optionally contain the `>` character followed by the name of the stream to which the handle should be directed.

There is are two special cases for the names on this line: they can start with either `*.` or `.*` to indicate the settings apply to a whole set of handles. See the example below.

Here is a short example of a configuration file. It activates all handles by default, and defines four handles: two of them are directed to the default stream (standard error), the third one to a file on the disk, and the last one to the system logger syslog (if your system supports it, otherwise to the default stream, ie standard error):

```
+
>&2
MODULE1
```

(continues on next page)

(continued from previous page)

```

MODULE2=yes
SYSLOG=yes >&syslog:local0:info
FILE=yes >/tmp/file

-- decorators (see below)
DEBUG.COLORS=yes

-- Applies to FIRST.EXCEPTIONS, LAST.EXCEPTIONS,...
-- and forces them to be displayed on stdout
*.EXCEPTIONS=yes > stdout

-- Applies to MODULE1, MODULE1.FIRST,... This can be used to
-- disable a whole hierarchy of modules.
-- As always, the latest config overrides earlier ones, so the
-- module MODULE1.EXCEPTIONS would be disabled as well.

MODULE1.*=no

```

4.2 Using the traces module

If you need or want to parse an external configuration file as described in the first section, the code that initializes your application should contain a call to `GNATCOLL.Traces.Parse_Config_File`. As documented, this takes in parameter the name of the configuration file to parse. When none is specified, the algorithm specified in the previous section will be used to find an appropriate configuration:

```
GNATCOLL.Traces.Parse_Config_File;
```

The code, as written, will end up looking for a file `.gnatdebug` in the current directory.

The function `Parse_Config_File` must be called to indicate that you want to activate the traces. It must also end up finding a configuration file. If it does not, then none of the other functions will ever output anything. This is to make sure your application does not start printing extra output just because you happen to use an external library that uses `GNATCOLL.Traces`. It also ensures that your application will not try to write to `stdout` unless you think it is appropriate (since `stdout` might not even exist in fact).

You then need to declare each of the *trace_handle* (or *logger*) that your application will use. The same handle can be declared several times, so the recommended approach is to declare locally in each package body the handles it will need, even if several bodies actually need the same handle. That helps to know which traces to activate when debugging a package, and limits the dependencies of packages on a shared package somewhere that would contain the declaration of all shared handles.

Function `Trace_Handle Create Name Default Stream Factory Finalize`

This function creates (or return an existing) a *trace_handle* with the specified *Name*. Its default activation status can also be specified (through *Default*), although the default behavior is to get it from the configuration file. If a handle is created several times, only the first call that is executed can define the default activation status, the following calls will have no effect.

Stream is the name of the stream to which it should be directed. Here as well, it is generally better to leave things to the configuration file, although in some cases you might want to force a specific behavior.

Factory is used to create your own child types of *trace_handle* (*Log decorators*).

Here is an example with two package bodies that define their own handles, which are later used for output:

```
package body Pkg1 is
  Me : constant Trace_Handle := Create ("PKG1");
  Log : constant Trace_Handle := Create ("LOG", Stream => "@syslog");
end Pkg1;

package body Pkg2 is
  Me : constant Trace_Handle := Create ("PKG2");
  Log : constant Trace_Handle := Create ("LOG", Stream => "@syslog");
end Pkg2;
```

Once the handles have been declared, output is a matter of calling the *GNATCOLL.Traces.Trace* procedure, as in the following sample:

```
Trace (Me, "I am here");
```

An additional subprogram can be used to test for assertions (pre-conditions or post-conditions in your program), and output a message whether the assertion is met or not:

```
Assert (Me, A = B, "A is not equal to B");
```

If the output of the stream is done in color, a failed assertion is displayed with a red background to make it more obvious.

4.2.1 Logging unexpected exceptions

A special version of *Trace* is provided, which takes an *Exception_Occurrence* as argument, and prints its message and backtrace into the corresponding log stream.

This procedure will in general be used for unexcepted exceptions. Since such exceptions should be handled by developers, it is possible to configure *GNATCOLL.TRACES* to use special streams for those.

Trace (Me, E) will therefore not use *Me* itself as the log handle, but will create (on the fly, the first time) a new handle with the same base name and an *.EXCEPTIONS* suffix. Therefore, you could put the following in your configuration file:

```
# Redirect all exceptions to stdout
*.EXCEPTIONS=yes >& stdout
```

and then the following code will output the exception trace to stdout:

```
procedure Proc is
  Me : Create ("MYMODULE");
begin
  ...
exception
  when E : others =>
    Trace (Me, E, Msg => "unexpected exception:");
end Proc;
```

4.2.2 Checking whether the handle is active

As we noted before, handles can be disabled. In that case, your application should not spend time preparing the output string, since that would be wasted time. In particular, using the standard Ada string concatenation operator requires allocating temporary memory. It is therefore recommended, when the string to display is complex, to first test whether the handle is active. This is done with the following code:

```
if Active (Me) then
  Trace (Me, A & B & C & D & E);
end if;
```

4.3 Log decorators

Speaking of color, a number of decorators are defined by *GNATCOLL.Traces*. Their goal is not to be used for outputting information, but to configure what extra information should be output with all log messages. They are activated through the same configuration file as the traces, with the same syntax (i.e either “=yes” or “=no”).

Here is an exhaustive list:

DEBUG.ABSOLUTE_TIME

If this decorator is activated in the configuration file, the absolute time when Trace is called is automatically added to the output, when the streams supports it (in particular, this has no effect for syslog, which already does this on its own).

DEBUG.MICRO_TIME

If active, the time displayed by *DEBUG.ABSOLUTE_TIME* will use a microseconds precision, instead of milliseconds.

DEBUG.ELAPSED_TIME

If this decorator is activated, then the elapsed time since the last call to Trace for the same handle is also displayed.

DEBUG.STACK_TRACE

If this decorator is activated, then the stack trace is also displayed. It can be converted to a symbolic stack trace through the use of the external application *addr2line*, but that would be too costly to do this automatically for each message.

DEBUG.LOCATION

If this decorator is activated, the location of the call to Trace is automatically displayed. This is a *file:line:column* information. This works even when the executable wasn't compiled with debug information

DEBUG.ENCLOSING_ENTITY

Activate this decorator to automatically display the name of the subprogram that contains the call to *Trace*.

DEBUG.COLORS

If this decorator is activated, the messages will use colors for the various fields, if the stream supports it (syslog doesn't).

DEBUG.COUNT

This decorator displays two additional numbers on each line: the first is the number of times this handle was used so far in the application, the second is the total number of traces emitted so far. These numbers can for instance be used to set conditional breakpoints on a specific trace (break on *gnat.traces.log* or *gnat.traces.trace* and check the value of *Handle.Count*. It can also be used to refer to a specific line in some comment file.

DEBUG.MEMORY

Every time a message is output, display the amount of memory currently in use by the application.

DEBUG.SPLIT_LINES

When this is enabled, messages are split at each newline character. Each line then starts with the name of the logger, indentation level and so on. This might result in more readable output, but is slightly slower.

DEBUG.FINALIZE_TRACES

This handle is activated by default, and indicates whether *GNATCOLL.Traces.Finalize* should have any effect. This can be set to False when debugging, to ensure that traces are available during the finalization of your application.

Here is an example of output where several decorators were activated. In this example, the output is folded on several lines, but in reality everything is output on a single line:

```
[MODULE] 6/247 User Message (2007-07-03 13:12:53.46)
(elapsed: 2ms)(loc: gnatcoll-traces.adb:224)
(entity:GNATCOLL.Traces.Log)
(callstack: 40FD9902 082FCFDD 082FE8DF )
```

Depending on your application, there are lots of other possible decorators that could be useful (for instance the current thread, or the name of the executable when you have several of them,...). Since *GNATCOLL.Traces* cannot provide all possible decorators, it provides support, through tagged types, so that you can create your own decorators.

This needs you to override the *Trace_Handle_Record* tagged type. Since this type is created through calls to *GNATCOLL.Traces.Create*. This is done by providing an additional *Factory* parameter to *Create*; this is a function that allocates and returns the new handle.

Then you can override either (or both) of the primitive operations *Pre_Decorator* and *Post_Decorator*. The following example creates a new type of handles, and prints a constant string just after the module name:

```
type My_Handle is new Trace_Handle_Record with null record;
procedure Pre_Decorator
(Handle : in out My_Handle;
 Stream : in out Trace_Stream_Record'Class;
 Message : String) is
begin
  Put (Stream, "TEST");
  Pre_Decorator (Trace_Handle_Record (Handle), Stream, Message);
end**;

function Factory return Trace_Handle is
begin
  return new My_Handle;
end;

Me : Trace_Handle := Create ("MODULE", Factory => Factory'Access);
```

As we will see below (*Dynamically disabling features*), you can also make all or part of your decorators conditional and configurable through the same configuration file as the trace handles themselves.

4.4 Defining custom stream types

We noted above that several predefined types of streams exist, to output to a file, to standard output or to standard error. Depending on your specific needs, you might want to output to other media. For instance, in a graphical application, you could have a window that shows the traces (perhaps in addition to filing them in a file, since otherwise the window would disappear along with its contents if the application crashes); or you could write to a socket (or even a CORBA ORB) to communicate with another application which is charge of monitoring your application.

You do not need the code below if you simply want to have a new stream in your application (for instance using one for logging Info messages, one for Error messages, and so on). In this case, the function *Create* is all you need.

GNATCOLL.Traces provides the type *Trace_Stream_Record*, which can be overridden to redirect the traces to your own streams.

Let's assume for now that you have defined a new type of stream (called "mystream"). To keep the example simple, we will assume this stream also redirects to a file. For flexibility, however, you want to let the user configure the file name from the traces configuration file. Here is an example of a configuration file that sets the default stream to a file called *foo*, and redirects a specific handle to another file called *bar*. Note how the same syntax that was used for standard output and standard error is also reused (ie the stream name starts with the "&" symbol, to avoid confusion with standard file names):

```
>&mystream:foo
MODULE=yes >&mystream:bar
```

You need of course to do a bit of coding in Ada to create the stream. This is done by creating a new child of *Trace_Stream_Record*, and override the primitive operation *Put*.

The whole output message is given as a single parameter to *Put*:

```
type My_Stream is new Trace_Stream_Record with record
  File : access File_Type;
end record;

procedure Put
  (Stream : in out My_Stream; Str : Msg_Strings.XString)
is
  S : Msg_Strings.Unconstrained_String_Access;
  L : Natural;
begin
  Str.Get_String (S, L);
  Put (Stream.File.all, String (S (1 .. L)));
end Put;
```

The above code did not open the file itself, as you might have noticed, nor did it register the name "mystream" so that it can be used in the configuration file. All this is done by creating a factory, ie a function in charge of creating the new stream.

A factory is also a tagged object (so that you can store custom information in it), with a single primitive operation, *New_Stream*, in charge of creating and initializing a new stream. This operation receives in parameter the argument specified by the user in the configuration file (after the ":" character, if any), and must return a newly allocated stream. This function is also never called twice with the same argument, since *GNATCOLL.Traces* automatically reuses an existing stream when one with the same name and arguments already exists:

```
type My_Stream_Factory is new Stream_Factory with null record;

overriding function New_Stream
```

(continues on next page)

(continued from previous page)

```
(Self : My_Stream_Factory; Args : String) return Trace_Stream
is
  Str : access My_Stream := new My_Stream;
begin
  Str.File := new File_Type;
  Open (Str.File, Out_File, Args);
  return Str;
end Factory;

Fact : access My_Stream_Factory := new My_Stream_Factory;
Register_Stream_Factory ("mystream", Fact);
```

4.5 Logging to syslog

Among the predefined streams, GNATColl gives access to the system logger *syslog*. This is a standard utility on all Unix systems, but is not available on other systems. When you compile GNATColl, you should specify the switch *-enable-syslog* to configure to activate the support. If either this switch wasn't specified, or configure could not find the relevant header files anyway, then support for *syslog* will not be available. In this case, the package *GNATCOLL.Traces.Syslog* is still available, but contains a single function that does nothing. If your configuration files redirect some trace handles to "*syslog*", they will instead be redirect to the default stream or to standard output.

Activating support for syslog requires the following call in your application:

```
GNATCOLL.Traces.Syslog.Register_Syslog_Stream;
```

This procedure is always available, whether your system supports or not syslog, and will simply do nothing if it doesn't support syslog. This means that you do not need to have conditional code in your application to handle that, and you can let GNATColl take care of this.

After the above call, trace handles can be redirected to a stream named "*syslog*".

The package *GNATCOLL.Traces.Syslog* also contains a low-level interface to syslog, which, although fully functional, you should probably not use, since that would make your code system-dependent.

Syslog itself dispatches its output based on two criteria: the *facility*, which indicates what application emitted the message, and where it should be filed, and the *level* which indicates the urgency level of the message. Both of these criteria can be specified in the *GNATCOLL.Traces* configuration file, as follows:

```
MODULE=yes >&syslog:user:error
```

The above configuration will redirect to a facility called *user*, with an urgency level *error*. See the enumeration types in `gnatcoll-traces-syslog.ads` for more information on valid facilities and levels.

4.6 Dynamically disabling features

Although the trace handles are primarily meant for outputting messages, they can be used in another context. The goal is to take advantage of the external configuration file, without reimplementing a similar feature in your application. Since the configuration file can be used to activate or deactivate a handle dynamically, you can then have conditional sections in your application that depend on that handle, as in the following example:

```
CONDITIONAL=yes
```

and in the Ada code:

```
package Pkg is
  Me : constant Trace_Handle := Create ("CONDITIONAL");
begin
  if Active (Me) then
    ... conditional code
  end if;
end Pkg;
```

In particular, this can be used if you write your own decorators, as explained above.

STRINGS: HIGH-PERFORMANCE STRINGS

The generic package `GNATCOLL.Strings_Impl` (and its default instantiation in `GNATCOLL.Strings`) provides a high-performance strings implementation.

It comes in addition to Ada's own *String* and *Unbounded_String* types, although it attempts to find a middle ground in between (flexibility vs performance).

`GNATCOLL.Strings` therefore provides strings (named *XString*, as in extended-strings) that can grow as needed (up to *Natural'Last*, like standard strings), yet are faster than unbounded strings. They also come with an extended API, which includes all primitive operations from unbounded strings, in addition to some subprograms inspired from `GNATCOLL.Utils` and the python and C++ programming languages.

5.1 Small string optimization

`GNATCOLL.Strings` uses a number of tricks to improve on the efficiency. The most important one is to limit the number of memory allocations. For this, we use a trick similar to what all C++ implementations do nowadays, namely the small string optimization.

The idea is that when a string is short, we can avoid all memory allocations altogether, while still keeping the string type itself small. We therefore use an `Unchecked_Union`, where a string can be viewed in two ways:

Small **string**

```
[f][s][ characters of the string 23 bytes ]
  f = 1 bit for a flag, set to 0 for a small string
  s = 7 bits for the size of the string (i.e. number of significant
      characters in the array)
```

Big **string**

```
[f][c ][size ][data ][first ][pad ]
  f = 1 bit for a flag, set to 1 for a big string
  c = 31 bits for half the capacity. This is the size of the buffer
      pointed to by data, and which contains the actual characters of
      the string.
  size = 32 bits for the size of the string, i.e. the number of
      significant characters in the buffer.
  data = a pointer (32 or 64 bits depending on architecture)
  first = 32 bits, see the handling of substrings below
  pad = 32 bits on a 64 bits system, 0 otherwise.
      This is because of alignment issues.
```

So in the same amount of memory (24 bytes), we can either store a small string of 23 characters or less with no memory allocations, or a big string that requires allocation. In a typical application, most strings are smaller than 23 bytes, so we are saving very significant time here.

This representation has to work on both 32 bits systems and 64 bits systems, so we have careful representation clauses to take this into account. It also needs to work on both big-endian and little-endian systems. Thanks to Ada's representation clauses, this one in fact relatively easy to achieve (well, okay, after trying a few different approaches to emulate what's done in C++, and that did not work elegantly). In fact, emulating via bit-shift operations ended up with code that was less efficient than letting the compiler do it automatically because of our representation clauses.

5.2 Character types

Applications should be able to handle the whole set of Unicode characters. In Ada, these are represented as the `Wide_Character` type, rather than `Character`, and stored on 2 bytes rather than 1. Of course, for a lot of applications it would be wasting memory to always store 2 bytes per character, so we want to give flexibility to users here.

So the package `GNATCOLL.Strings_Impl` is a generic. It has several formal parameters, among which:

- `Character_Type` is the type used to represent each character. Typically, it will be `Character`, `Wide_Character`, or even possibly `Wide_Wide_Character`. It could really be any scalar type, so for instance we could use this package to represent DNA with its 4-valued nucleobases.
- `Character_String` is an array of these characters, as would be represented in Ada. It will typically be a `String` or a `Wide_String`. This type is used to make this package work with the rest of the Ada world.

Note about Unicode: we could also always use a `Character`, and use UTF-8 encoding internally. But this makes all operations (from taking the length to moving the next character) slower, and more fragile. We must make sure not to cut a string in the middle of a multi-byte sequence. Instead, we manipulate a string of code points (in terms of Unicode). A similar choice is made in Ada (`String` vs `Wide_String`), Python and C++.

5.3 Configuring the size of small strings

The above is what is done for most C++ implementations nowadays. The maximum 23 characters we mentioned for a small string depends in fact on several criteria, which impact the actual maximum size of a small string:

- on 32 bits system, the size of the big string is 16 bytes, so the maximum size of a small string is 15 bytes.
- on 64 bits system, the size of the big string is 24 bytes, so the maximum size of a small string is 23 bytes.
- If using a `Character` as the character type, the above are the actual number of characters in the string. But if you are using a `Wide_Character`, this is double the maximum length of the string, so a small string is either 7 characters or 11 characters long.

This is often a reasonable number, and given that applications mostly use small strings, we are already saving a lot of allocations. However, in some cases we know that the typical length of strings in a particular context is different. For instance, `GNATCOLL.Traces` builds messages to output in the log file. Such messages will typically be at most 100 characters, although they can of course be much larger sometimes.

We have added one more formal parameter to `GNATCOLL.Strings_Impl` to control the maximum size of small strings. If for instance we decide that a “small” string is anywhere from 1 to 100 characters long (i.e. we do not want to allocate memory for those strings), it can be done via this parameter.

Of course, in such cases the size of the string itself becomes much larger. In this example it would be 101 bytes long, rather than the 24 bytes. Although we are saving on memory allocations, we are also spending more time copying data when the string is passed around, so you'll need to measure the performance here.

The maximum size for the small string is 127 bytes however, because this size and the 1-bit flag need to fit in 1 bytes in the representation clauses we showed above. We tried to make this more configurable, but this makes things significantly more complex between little-endian and big-endian systems, and having large “small” strings would not make much sense in terms of performance anyway.

Typical C++ implementations do not make this small size configurable.

5.4 Task safety

Just like unbounded strings, the strings in this package are not thread safe. This means that you cannot access the same string (read or write) from two different threads without somehow protecting the access via a protected type, locks,...

In practice, sharing strings would rarely be done, so if the package itself was doing its own locking we would end up with very bad performance in all cases, for a few cases where it might prove useful.

As we’ll discuss below, it is possible to use two different strings that actually share the same internal buffer, from two different threads. Since this is an implementation detail, this package takes care of guaranteeing the integrity of the shared data in such a case.

5.5 Copy on write

There is one more formal parameter, to configure whether this package should use copy-on-write or not. When copy on write is enabled, you can have multiple strings that internally share the same buffer of characters. This means that assigning a string to another one becomes a reasonably fast operation (copy a pointer and increment a refcount). Whenever the string is modified, a copy of the buffer is done so that other copies of the same string are not impacted.

But in fact, there is one drawback with this scheme: we need reference counting to know when we can free the shared data, or when we need to make a copy of it. This reference counting must be thread safe, since users might be using two different strings from two different threads, but they share data internally.

Thus the reference counting is done via atomic operations, which have some impact on performance. Since multiple threads try to access the same memory addresses, this is also a source of contention in multi-threaded applications.

For this reason, the current C++ standard prevents the use of copy-on-write for strings.

In our case, we chose to make this configurable in the generic, so that users can decide whether to pay the cost of the atomic operations, but save on the number of memory allocations and copy of the characters. Sometimes it is better to share the data, sometimes to systematically copy it. Again, actual measurements of the performance are needed for your specific application.

5.6 Growth strategy

When the current size of the string becomes bigger than the available allocated memory (for instance because you are appending characters), this package needs to reallocate memory. There are plenty of strategies here, from allocating only the exact amount of memory needed (which saves on memory usage, but is very bad in terms of performance), to doubling the current size of the string until we have enough space, as currently done in the GNAT unbounded strings implementation.

The latter approach would therefore allocate space for two characters, then for 4, then 8 and so on.

This package has a slightly different strategy. Remember that we only start allocating memory past the size of small strings, so we will for instance first allocate 24 bytes. When more memory is needed, we multiply this size by 1.5, which some researchers have found to be a good compromise between waste of memory and number of allocations. For

very large strings, we always allocate multiples of the memory page size (4096 bytes), since this is what the system will make available anyway. So we will basically allocate the following: 24, 36, 54, 82, 122,...

An additional constraint is that we only ever allocate even number of bytes. This is called the capacity of the string. In the layout of the big string, as shown above, we store half that capacity, which saves one bit that we use for the flag.

5.7 Substrings

One other optimization performed by this package (which is not done for unbounded strings or various C++ implementations) is to optimize substrings when also using copy-on-write.

We simply store the index of the first character of the string within the shared buffer, instead of always starting at the first.

From the user's point of view, this is an implementation detail. Strings are always indexed from 1, and internally we convert to an actual position in the buffer. This means that if we need to reallocate the buffer, for instance when the string is modified, we transparently change the index of the first character, but the indexes the user was using are still valid.

This results in very significant savings, as shown below in the timings for Trim for instance. Also, we can do an operation like splitting a string very efficiently.

For instance, the following code doesn't allocate any memory, beside setting the initial value of the string. It parses a file containing some "key=value" lines, with optional spaces, and possibly empty lines:

```
declare
  S, Key, Value : XString;
  L              : XString_Array (1 .. 2);
  Last           : Natural;
begin
  S.Set (".....");

  -- Get each line
  for Line in S.Split (ASCII.LF) loop

    -- Split into at most two substrings
    Line.Split ('=', Into => L, Last => Last);

    if Last = 2 then
      Key := L (1);
      Key.Trim;    -- Removing leading and trailing spaces

      Value := L (2);
      Value.Trim;

    end if;
  end loop;
end;
```

5.8 API

This package provides a very extensive set of API that apply to *XString*, please check the spec in `gnatcoll-strings_impl.ads` for a fully documented list.

MEMORY: MONITORING MEMORY USAGE

The GNAT compiler allocates and deallocates all memory either through type-specific debug pools that you have defined yourself, or defaults to the standard *malloc* and *free* system calls. However, it calls those through an Ada proxy, in the package *System.Memory* that you can also replace in your own application if need be.

Like this:

```
procedure Ada
```

gnatcoll provides such a possible replacement. Its implementation is also based on *malloc* and *free*, but if you so chose you can activate extra monitoring capabilities to help you find out which parts of your program is allocating the most memory, or where memory is allocated at any moment in the life of your application.

This package is called *GNATCOLL.Memory*. To use it requires a bit of preparation in your application:

- You need to create your own version of *s-memory.adb* with the template below, and put it somewhere in your source path. This file should contain the following bit of code:

```
with GNATCOLL.Memory;

package body System.Memory is
  package M renames GNATCOLL.Memory;

  function Alloc (Size : size_t) return System.Address is
  begin
    return M.Alloc (M.size_t (Size));
  end Alloc;

  procedure Free (Ptr : System.Address) renames M.Free;

  function Realloc
    (Ptr : System.Address;
     Size : size_t)
    return System.Address is
  begin
    return M.Realloc (Ptr, M.size_t (Size));
  end Realloc;
end;
```

- You then need to compile your application with the extra switch *-a* passed to *gnatmake* or *gprbuild*, so that this file is appropriately compiled and linked with your application
- If you only do this, the monitor is disabled by default. This basically has zero overhead for your application (apart from the initial small allocation of some internal data). When you call the procedure *GNAT-*

COLL.Memory.Configure to activate the monitor, each memory allocation or deallocation will result in extra overhead that will slow down your application a bit. But at that point you can then get access to the information stored in the monitor

We actually recommend that the activation of the monitor be based on an environment variable or command line switch of your application, so that you can decide at any time to rerun your application with the monitor activated, rather than have to go through an extra recompilation.

All allocations and deallocations are monitor automatically when this module is activated. However, you can also manually call *GNATCOLL.Memory.Mark_Traceback* to add a dummy entry in the internal tables that matches the current stack trace. This is helpful for instance if you want to monitor the calls to a specific subprogram, and know both the number of calls, and which callers executed it how many times. This can help find hotspots in your application to optimize the code.

The information that is available through the monitor is the list of all chunks of memory that were allocated in Ada (this does not include allocations done in other languages like C). These chunks are grouped based on the stack trace at the time of their invocation, and this package knows how many times each stack trace executed each allocation.

As a result, you can call the function *GNATCOLL.Memory.Dump* to dump on the standard output various types of data, sorted. To limit the output to a somewhat usable format, *Dump* asks you to specify how many blocks it should output.

Debugging dangling pointer Using a dangling pointer can lead (and usually it does) to no crash or no side effects. Frequently, freed buffers still contains valid data and are still part of pages owned by your process. Probably, this occurs more often on linux compare to windows.

Writing 0 or 0xDD pattern when a memory is freed will be (because of the exception that will be thrown) detected at the first usage of a freed buffer. The crash occurrence will be higher and less random. This makes solid reproducer more easy to build.

For dangling pointer usage debugging, use `Memory_Free_Pattern` parameter when calling *GNATCOLL.Memory.Configure* procedure.

Memory usage

Blocks are sorted based on the amount of memory they have allocated and is still allocated. This helps you find which part of your application is currently using the most memory.

Allocations count

Blocks are sorted based on the number of allocation that are still allocated. This helps you find which part of your application has done the most number of allocations (since malloc is a rather slow system call, it is in general a good idea to try and reduce the number of allocations in an application).

Total number of allocations

This is similar to the above, but includes all allocations ever done in this block, even if memory has been deallocated since then.

Marked blocks

These are the blocks that were created through your calls to *GNATCOLL.Memory.Mark_Traceback*. They are sorted by the number of allocation for that stacktrace, and also shows you the total number of such allocations in marked blocks. This is useful to monitor and analyze calls to specific places in your code

MMAP: READING AND WRITING FILES

Most applications need to efficiently read files from the disk. Some also need in addition to modify them and write them back. The Ada run-time provides several high-level functions to do so, most notably in the `Ada.Text_IO` package. However, these subprograms require a lot of additional housekeeping in the run-time, and therefore tend to be slow.

GNAT provides a number of low-level functions in its `GNAT.OS_Lib` package. These are direct import of the usual C system calls `read()`, `write()` and `open()`. These are much faster, and suitable for most applications.

However, if you happen to manipulate big files (several megabytes and much more), these functions are still slow. The reason is that to use `read` you basically need a few other system calls: allocate some memory to temporarily store the contents of the file, then read the whole contents of the file (even if you are only going to read a small part of it, although presumably you would use `lseek` in such a case).

On most Unix systems, there exists an additional system call `mmap()` which basically replaces `open`, and makes the contents of the file immediately accessible, in the order of a few micro-seconds. You do not need to allocate memory specifically for that purpose. When you access part of the file, the actual contents is temporarily mapped in memory by the system. To modify the file, you just modify the contents of the memory, and do not worry about writing the file back to the disk.

When your application does not need to read the whole contents of the file, the speed up can be several orders of magnitude faster than `read()`. Even when you need to read the whole contents, using `mmap()` is still two or three times faster, which is especially interesting on big files.

GNATColl's `GNATCOLL.Mmap` package provides a high-level abstraction on top of the `mmap` system call. As for most other packages in GNATColl, it also nicely handles the case where your system does not actually support `mmap`, and will in that case fallback on using `read` and `write` transparently. In such a case, your application will perform a little slower, but you do not have to modify your code to adapt it to the new system.

Due to the low-level C API that is needed underneath, the various subprograms in this package do not directly manipulate Ada strings with valid bounds. Instead, a new type `Str_Access` was defined. It does not contain the bounds of the string, and therefore you cannot use the usual `'First` and `'Last` attributes on that string. But there are other subprograms that provide those values.

Here is how to read a whole file at once. This is what your code will use in most cases, unless you expect to read files bigger than `Integer'Last` bytes long. In such cases you need to read chunks of the file separately. The `mmap` system call is such that its performance does not depend on the size of the file you are mapping. Of course, this could be a problem if `GNATCOLL.Mmap` falls back on calling `read`, since in that case it needs to allocate as much memory as your file. Therefore in some cases you will also want to only read chunks of the file at once:

```
declare
  File : Mapped_File;
  Reg  : Mapped_Region;
  Str  : Long.Str_Access;
begin
  File := Open_Read ("/tmp/file_on_disk");
```

(continues on next page)

(continued from previous page)

```

Reg := Read (File);  -- map the whole file*
Close (File);

Str := Long.Data (File);
for S in 1 .. Long.Last (File) loop
  Put (Str (S));
end loop;
Free (Reg);
end;

```

The above example works for files larger than 2Gb, on 64 bits system (up to a petabyte in fact), on systems that support the *mmap* system call.

To read only a chunk of the file, your code would look like the following. At the low-level, the system call will always read chunks multiple of a size called the *page_size*. Although *GNATCOLL.Mmap* takes care of rounding the numbers appropriately, it is recommended that you pass parameters that are multiples of that size. That optimizes the number of system calls you will need to do, and therefore speeds up your application somewhat:

```

declare
  File   : Mapped_File;
  Reg    : Mapped_Region;
  Str    : Str_Access;
  Offs   : Long_Integer := 0;
  Page   : constant Integer := Get_Page_Size;
begin
  File := Open_Read ("/tmp/file_on_disk");
  while Offs < Length (File) loop
    Read (File, Reg, Offs, Length => Long_Integer (Page) * 4);
    Str := Data (File);

    -- Print characters for this chunk.*
    for S in Integer (Offs - Offset (File)) + 1 .. Last (File) loop
      Put (Str (S));
    end loop;

    Offs := Offs + Long_Integer (Last (File));
  end loop;
  Free (Reg);
  Close (File);
end;

```

There are a number of subtle details in the code above. Since the system call only manipulates chunk of the file on boundaries multiple of the code size, there is no guarantee that the part of the file we actually read really starts exactly at *Offs*. It could in fact start before, for rounding issues. Therefore when we loop over the contents of the buffer, we make sure to actually start at the *Offs*-th character in the file.

In the particular case of this code, we make sure we only manipulate multiples of the *page_size*, so we could in fact replace the loop with the simpler:

```

for S in 1 .. Last (File) loop

```

If you intend to modify the contents of the file, not that *GNATCOLL.Mmap* currently gives you no way to change the size of the file. The only difference compared to the code used for reading the file is the call to open the file, which should be:

```
File := Open_Write ("/tmp/file_on_disk");
```

Modifications to *Str* are automatically reflected in the file. However, there is no guarantee this saving is done immediately. It could be done only when you call *Close*. This is in particular always the case when your system does not support *mmap* and *GNATCOLL.Mmap* had to fallback on calls to *read*.

BOYER-MOORE: SEARCHING STRINGS

Although the Ada standard provides a number of string-searching subprograms (most notably in the *Ada.Strings.Fixed*, *Ada.Strings.Unbounded* and *Ada.Strings.Bounded* packages through the *Index* functions), these subprograms do not in general provide the most efficient algorithms for searching strings.

The package **GNATCOLL.Boyer_Moore** provides one such optimize algorithm, although there exists several others which might be more efficient depending on the pattern.

It deals with string searching, and does not handle regular expressions for instance.

This algorithm needs to preprocess its key (the searched string), but does not need to perform any specific analysis of the string to be searched. Its execution time can be sub-linear: it doesn't need to actually check every character of the string to be searched, and will skip over some of them. The worst case for this algorithm has been proved to need approximately $3 * N$ comparisons, hence the algorithm has a complexity of $O(n)$.

The longer the key, the faster the algorithm in general, since that provides more context as to how many characters can be skipped when a non-matching character is found..

We will not go into the details of the algorithm, although a general description follows: when the pattern is being preprocessed, Boyer-Moore computes how many characters can be skipped if an incorrect match is found at that point, depending on which character was read. In addition, this algorithm tries to match the key starting from its end, which in general provides a greater number of characters to skip.

For instance, if you are looking for "ABC" in the string "ABDEFG" at the first position, the algorithm will compare "C" and "D". Since "D" does not appear in the key "ABC", it knows that it can immediately skip 3 characters and start the search after "D".

Using this package is extremely easy, and it has only a limited API:

```
declare
  Str : constant String := "ABDEABCFGABC";
  Key : Pattern;
  Index : Integer;
begin
  Compile (Key, "ABC");
  Index := Search (Key, Str);
end
```

Search will either return -1 when the pattern did not match, or the index of the first match in the string. In the example above, it will return 5.

If you want to find the next match, you have to pass a substring to search, as in:

```
Index := Search (Key, Str (6 .. Str'Last));
```


PARAGRAPH FILLING: FORMATTING TEXT

The package *GNATCOLL.Paragraph_Filling* provides several algorithms for filling paragraphs—formatting them to take up the minimal number of lines and to look better. *Knuth_Fill* is based on an algorithm invented by Donald Knuth, and used in TeX. *Pretty_Fill* uses a different algorithm, which was judged by some to produce more aesthetically pleasing output.

More detailed documentation may be found in the comments in the package spec.

TEMPLATES: GENERATING TEXT

This module provides convenient subprograms for replacing specific substrings with other values. It is typically used to replace substrings like “%{version}” in a longer string with the actual version, at run time.

This module is not the same as the templates parser provided in the context of AWS, the Ada web server, where external files are parsed and processed to generate other files. The latter provides advanced features like filters, loops,...

The substrings to be replaced always start with a specific delimiter, which is set to % by default, but can be overridden in your code. The name of the substring to be replaced is then the identifier following that delimiter, with the following rules:

- If the character following the delimiter is the delimiter itself, then the final string will contain a single instance of that delimiter, and no further substitution is done for that delimiter. An example of this is “%%”.
- If the character immediately after the delimiter is a curly brace (/), then the name of the identifier is the text until the next closing curly brace. It can then contain any character except a closing curly brace. An example of this is “%{long name}”
- If the first character after the delimiter is a digit, then the name of the identifier is the number after the delimiter. An example of this is “%12”. As a special case, if the first non-digit character is the symbol -, it is added as part of the name of the identifier, as in “%1-”. One use for this feature is to indicate you want to replace it with all the positional parameters %1%2%3%4. For instance, if you are writing the command line to spawn an external tool, to which the user can pass any number of parameter, you could specify that command line as “tool -o %1 %2-” to indicate that all parameters should be concatenated on the command line.
- If the first character after the delimiter is a letter, the identifier follows the same rules as for Ada identifiers, and can contain any letter, digit, or underscore character. An example of this is “%ab_12”. For readability, it is recommended to use the curly brace notation when the name is complex, but that is not mandatory.
- Otherwise the name of the identifier is the single character following the delimiter

For each substring matching the rules above, the *Substitute* subprogram will look for possible replacement text in the following order:

- If the *Substrings* parameter contains an entry for that name, the corresponding value is used.
- Otherwise, if a *callback* was specified, it is called with the name of the identifier, and should return the appropriate substitution (or raise an exception if no such substitution makes sense).
- A default value provided in the substring itself
- When no replacement string was found, the substring is kept unmodified

EMAIL: PROCESSING EMAIL MESSAGES

GNATColl provides a set of packages for managing and processing email messages. Through this packages, you can extract the various messages contained in an existing mailbox, extract the various components of a message, editing previously parsed messages, or create new messages from scratch.

This module fully supports MIME-encoded messages, with attachments.

This module currently does not provide a way to send the message through the SMTP protocol. Rather, it is used to create an in-memory representation of the message, which you can then convert to a string, and pass this to a socket. See for instance the [AWS library](#)) which contains the necessary subprograms to connect with an SMTP server.

11.1 Message formats

The format of mail messages is defined through numerous RFC documents. GNATColl tries to conform to these as best as possible. Basically, a message is made of two parts:

The headers

These are various fields that indicate who sent the message, when, to whom, and so on

The payload (aka body)

This is the actual contents of the message. It can either be a simple text, or made of one or more attachments in various formats. These attachments can be HTML text, images, or any binary file. Since email transfer is done through various servers, the set of bytes that can be sent is generally limited to 7 bit characters. Therefore, the attachments are generally encoded through one of the encoding defined in the various MIME RFCs, and they need to be decoded before the original file can be manipulated again.

GNATColl gives you access to these various components, as will be seen in the section [Parsing messages](#).

The package `GNATCOLL.Email.Utils` contains various subprograms to decode MIME-encoded streams, which you can use independently from the rest of the packages in the email module.

The headers part of the message contains various pieces of information about the message. Most of the headers have a well-defined semantics and format. However, a user is free to add new headers, which will generally start with *X-* prefix. For those fields where the format is well-defined, they contain various pieces of information:

Email addresses

The *From*, *TO* or *CC* fields, among others, contain list of recipients. These recipients are the usual email addresses. However, the format is quite complex, because the full name of the recipient can also be specified, along with comments. The package `GNATCOLL.Email.Utils` provides various subprograms for parsing email addresses and list of recipients.

Dates

The *Date* header indicates when the message was sent. The format of the date is also precisely defined in the RFC, and the package `GNATCOLL.Email.Utils` provides subprograms for parsing this date (or, on the contrary, to create a string from an existing time).

Text

The *Subject* header provides a brief overview of the message. It is a simple text header. However, one complication comes from the fact that the user might want to use extended characters not in the ASCII subset. In such cases, the Subject (or part of it) will be MIME-encoded. The package `GNATCOLL.Email.Utils` provides subprograms to decode MIME-encoded strings, with the various charsets.

11.2 Parsing messages

There are two ways a message is represented in memory: initially, it is a free-form *String*. The usual Ada operations can be used on the string, of course, but there is no way to extract the various components of the message. For this, the message must first be parsed into an instance of the *Message* type.

This type is controlled, which means that the memory will be freed automatically when the message is no longer needed.

The package `GNATCOLL.Email.Parser` provides various subprograms that parse a message (passed as a string), and create a *Message* out of it. Parsing a message might be costly in some cases, for instance if a big attachment needs to be decoded first. In some cases, your application will not need that information (for instance you might only be looking for a few of the headers of the message, and not need any information from the body). This efficiency concern is why there are multiple parsers. Some of them will ignore parts of the message, and thus be more efficient if you can use them.

Once a *Message* has been created, the subprograms in *GNATCOLL.Email* can be used to access its various parts. The documentation for these subprograms is found in the file *gnatcoll-email.ads* directly, and is not duplicated here.

11.3 Parsing mailboxes

Most often, a message is not found on its own (unless you are for instance writing a filter for incoming messages). Instead, the messages are stored in what is called a mailbox. The latter can contain thousands of such messages.

There are traditionally multiple formats that have been used for mailboxes. At this stage, GNATColl only supports one of them, the *mbx* format. In this format, the messages are concatenated in a single file, and separated by a newline.

The package *GNATCOLL.Email.Mailboxes* provides all the types and subprograms to manipulate mailboxes. Tagged types are used, so that new formats of mailboxes can relatively easily be added later on, or in your own application.

Here is a small code example that opens an mbox on the disk, and parses each message it contains:

```
declare
  Box : Mbox;
  Curs : Cursor;
  Msg : Message;
begin
  Open (Box, Filename => "my_mbox");
  Curs := Mbox_Cursor (First (Box));
  while Has_Element (Curs) loop
    Get_Message (Curs, Box, Msg);
    if Msg /= Null_Message then
      ...
    end if;
    Next (Curs, Box);
  end loop;
end;
```

As you can see, the mailbox needs to be opened first. Then we get an iterator (called a cursor, to match the Ada2005 containers naming scheme), and we then parse each message. The *if* test is optional, but recommended: the message that is returned might be null if the mailbox was corrupted and the message could not be parsed. There are still chances that the next message will be readable, so only the current message should be ignored.

11.4 Creating messages

The subprograms in *GNATCOLL.Email* can also be used to create a message from scratch. Alternatively, if you have already parsed a message, you can alter it, or easily generate a reply to it (using the *Reply_To* subprogram. The latter will preset some headers, so that message threading is preserved in the user's mailers.

RAVENSCAR: PATTERNS FOR MULTITASKING

GNATColl provides a set of patterns for concurrent programming using Ravenscar-compliant semantics only. The core goal of the GNATCOLL.Ravenscar (sub) packages is to ease the development of high-integrity multitasking applications by factorizing common behavior into instantiable, Ravenscar-compliant, generic packages. Instances of such generic packages guarantee predictable timing behavior and thus permit the application of most common timing analysis techniques.

12.1 Tasks

The *GNATCOLL.Ravenscar.Simple_Cyclic_Task* generic package lets instantiate a cyclic tasks executing the same operation at regular time intervals; on the other side, the *GNATCOLL.Ravenscar.Simple_Sporadic_Task* task lets instantiate sporadic tasks enforcing a minimum inter-release time.

12.2 Servers

Servers present a more sophisticated run-time semantics than tasks: for example, they can fulfill different kind of requests (see multiple queues servers). *Gnat.Ravenscar.Sporadic_Server_With_Callback* and *Gnat.Ravenscar.Timed_Out_Sporadic_Server* are particularly interesting. The former shows how synchronous inter-task communication can be faked in Ravenscar (the only form of communication permitted by the profile is through shared resources): the server receives a request to fulfill, computes the result and returns it by invoking a call-back. The latter enforces both a minimum and a maximum inter-release time: the server automatically releases itself and invokes an appropriate handler if a request is not posted within a given period of time.

12.3 Timers

Gnat.Ravenscar.Timers.One_Shot_Timer is the Ravenscar implementation of time-triggered event through Ada 2005 Timing Events.

STORAGE POOLS: CONTROLLING MEMORY MANAGEMENT

Ada gives full control to the user for memory management. That allows for a number of optimization in your application. For instance, if you need to allocate a lot of small chunks of memory, it is generally more efficient to allocate a single large chunk, which is later divided into smaller chunks. That results in a single system call, which speeds up your application.

This can of course be done in most languages. However, that generally means you have to remember not to use the standard memory allocations like *malloc* or *new*, and instead call one of your subprograms. If you ever decide to change the allocation strategy, or want to experiment with several strategies, that means updating your code in several places.

In Ada, when you declare the type of your data, you also specify through a *'Storage_Pool* attribute how the memory for instances of that type should be allocated. And that's it. You then use the usual *new* keyword to allocate memory.

GNATColl provides a number of examples for such storage pools, with various goals. There is also one advanced such pool in the GNAT run-time itself, called *GNAT.Debug_Pools*, which allows you to control memory leaks and whether all accesses do reference valid memory location (and not memory that has already been deallocated).

In GNATColl, you will find the following storage pools:

'GNATCOLL.Storage_Pools.Alignment'

This pool gives you full control over the alignment of your data. In general, Ada will only allow you to specify alignments up to a limited number of bytes, because the compiler must only accept alignments that can be satisfied in all contexts, in particular on the stack.

This package overcomes that limitation, by allocating larger chunks of memory than needed, and returning an address within that chunk which is properly aligned.

'GNATCOLL.Storage_Pools.Headers'

This pool allows you to allocate memory for the element and reserve extra space before it for a header. This header can be used to store per-element information, like for instance a reference counter, or next and previous links to other elements in the same collection.

In many cases, this can be used to reduce the number of allocations, and thus speed up the overall application.

VFS: MANIPULATING FILES

Ada was meant from the beginning to be a very portable language, across architectures. As a result, most of the code you write on one machine has good chances of working as is on other machines. There remains, however, some areas that are somewhat system specific. The Ada run-time, the GNAT specific run-time and GNATColl all try to abstract some of those operations to help you make your code more portable.

One of these areas is related to the way files are represented and manipulated. Reading or writing to a file is system independent, and taken care of by the standard run-time. Other differences between systems include the way file names are represented (can a given file be accessed through various casing or not, are directories separated with a backslash or a forward slash, or some other mean, and a few others). The GNAT run-time does a good job at providing subprograms that work on most types of filesystems, but the relevant subprograms are split between several packages and not always easy to locate. GNATColl groups all these functions into a single convenient tagged type hierarchy. In addition, it provides the framework for transparently manipulating files on other machines.

Another difference is specific to the application code: sometimes, a subprogram needs to manipulate the base name (no directory information) of a file, whereas sometimes the full file name is needed. It is somewhat hard to document this in the API, and certainly fills the code with lots of conversion from full name to base name, and sometimes reverse (which, of course, might be an expansive computation). To make this easier, GNATColl provides a type that encapsulates the notion of a file, and removes the need for the application to indicate whether it needs a full name, a base name, or any other part of the file name.

14.1 Filesystems abstraction

There exists lots of different filesystems on all machines. These include such things as FAT, VFAT, NTFS, ext2, VMS,... However, all these can be grouped into three families of filesystems:

- windows-based filesystems

On such filesystems, the full name of a file is split into three parts: the name of the drive (c:, d:,...), the directories which are separated by a backslash, and the base name. Such filesystems are sometimes inaccurately said to be case insensitive: by that, one means that the same file can be accessed through various casing. However, a user is generally expecting a specific casing when a file name is displayed, and the application should strive to preserve that casing (as opposed to, for instance, systematically convert the file name to lower cases).

A special case of a windows-based filesystems is that emulated by the cygwin development environment. In this case, the filesystem is seen as if it was unix-based (see below), with one special quirk to indicate the drive letter (the file name starts with “/cygwin/c/”).

- unix-based filesystems

On such filesystems, directories are separated by forward slashed. File names are case sensitive, that is a directory can contain both “foo” and “Foo”, which is not possible on windows-based filesystems.

- vms filesystem

This filesystem represents path differently than the other two, using brackets to indicate parent directories

A given machine can actually have several file systems in parallel, when a remote disk is mounted through NFS or samba for instance. There is generally no easy way to guess that information automatically, and it generally does not matter since the system will convert from the native file system to that of the remote host transparently (for instance, if you mount a windows disk on a unix machine, you access its files through forward slash- separated directory names).

GNATColl abstracts the differences between these filesystems through a set of tagged types in the *GNATCOLL.Filesystem* package and its children. Such a type has primitive operations to manipulate the names of files (retrieving the base name from a full name for instance), to check various attributes of the file (is this a directory, a symbolic link, is the file readable or writable), or to manipulate the file itself (copying, deleting, reading and writing). It provides similar operations for directories (creating or deleting paths, reading the list of files in a directory,...).

It also provides information on the system itself (the list of available drives on a windows machine for instance).

The root type *Filesystem_Record* is abstract, and is specialized in various child types. A convenient factory is provided to return the filesystem appropriate for the local machine (*Get_Local_Filesystem*), but you might chose to create your own factory in your application if you have specialized needs (*Remote_filesystems*).

14.1.1 file names encoding

One delicate part when dealing with filesystems is handling files whose name cannot be described in ASCII. This includes names in asian languages for instance, or names with accented letters.

There is unfortunately no way, in general, to know what the encoding is for a filesystem. In fact, there might not even be such an encoding (on linux, for instance, one can happily create a file with a chinese name and another one with a french name in the same directory). As a result, GNATColl always treats file names as a series of bytes, and does not try to assume any specific encoding for them. This works fine as long as you are interfacing the system (since the same series of bytes that was returned by it is also used to access the file later on).

However, this becomes a problem when the time comes to display the name for the user (for instance in a graphical interface). At that point, you need to convert the file name to a specific encoding, generally UTF-8 but not necessarily (it could be ISO-8859-1 in some cases for instance).

Since GNATColl cannot guess whether the file names have a specific encoding on the file system, or what encoding you might wish in the end, it lets you take care of the conversion. To do so, you can use either of the two subprograms *Locale_To_Display* and *Set_Locale_To_Display_Encoder*

14.2 Remote filesystems

Once the abstract for filesystems exists, it is tempting to use it to access files on remote machines. There are of course lots of differences with filesystems on the local machine: their names are manipulated similarly (although you need to somehow indicate on which host they are to be found), but any operation of the file itself needs to be done on the remote host itself, as it can't be done through calls to the system's standard C library.

Note that when we speak of disks on a remote machine, we indicate disks that are not accessible locally, for instance through NFS mounts or samba. In such cases, the files are accessed transparently as if they were local, and all this is taken care of by the system itself, no special layer is needed at the application level.

GNATColl provides an extensive framework for manipulating such remote files. It knows what commands need to be run on the remote host to perform the operations ("cp" or "copy", "stat" or "dir /a-d",...) and will happily perform these operations when you try to manipulate such files.

There are however two operations that your own application needs to take care of to take full advantage of remote files.

14.2.1 Filesystem factory

GNATColl cannot know in advance what filesystem is running on the remote host, so it does not try to guess it. As a result, your application should have a factory that creates the proper instance of a *Filesystem_Record* depending on the host. Something like:

```

type Filesystem_Type is (Windows, Unix);
function Filesystem_Factory
  (Typ : Filesystem_Type;
   Host : String)
  return Filesystem_Access
is
  FS : Filesystem_Access;
begin
  if Host = "" then
    case Typ is
      when Unix =>
        FS := new Unix_Filesystem_Record;
      when Windows =>
        FS := new Windows_Filesystem_Record;
    end case;
  else
    case Typ is
      when Unix =>
        FS := new Remote_Unix_Filesystem_Record;
        Setup (Remote_Unix_Filesystem_Record (FS.all),
              Host      => Host,
              Transport => ...); -- see below*
      when Windows =>
        FS := new Remote_Windows_Filesystem_Record;
        Setup (Remote_Windows_Filesystem_Record (FS.all),
              Host      => Host,
              Transport => ...);
    end case;
  end if;

  Set_Locale_To_Display_Encoder
    (FS.all, Encode_To_UTF8'Access);
  return FS;
end Filesystem_Factory;

```

14.2.2 Transport layer

There exists lots of protocols to communicate with a remote machine, so as to be able to perform operations on it. These include protocols such as *rsh*, *ssh* or *telnet*. In most of these cases, a user name and password is needed (and will likely be asked to the user). Furthermore, you might not want to use the same protocol to connect to different machines.

GNATColl does not try to second guess your intention here. It performs all its remote operations through a tagged type defined in *GNATCOLL.Filesystem.Transport*. This type is abstract, and must be overridden in your application. For instance, GPS has a full support for choosing which protocol to use on which host, what kind of filesystem is running on that host, to recognize password queries from the transport protocol,... All these can be encapsulated in the transport protocol.

Once you have created one or more children of *Filesystem_Transport_Record*, you associate them with your instance of the filesystem through a call to the *Setup* primitive operation of the filesystem. See the factory example above.

14.3 Virtual files

As we have seen, the filesystem type abstracts all the operations for manipulating files and their names. There is however another aspect when dealing with file names in an application: it is often unclear whether a full name (with directories) is expected, or whether the base name itself is sufficient. There are also some aspects about a file that can be cached to improve the efficiency.

For these reasons, GNATColl provides a new type *GNATCOLL.VFS.Virtual_File* which abstracts the notion of file. It provides lots of primitive operations to manipulate such files (which are of course implemented based on the filesystem abstract, so support files on remote hosts among other advantages), and encapsulate the base name and the full name of a file so that your API becomes clearer (you are not expecting just any string, but really a file).

This type is reference counted: it takes care of memory management on its own, and will free its internal data (file name and cached data) automatically when the file is no longer needed. This has of course a slight efficiency cost, due to controlled types, but we have found in the context of GPS that the added flexibility was well worth it.

TRIBOOLEANS: THREE STATE LOGIC

Through the package *GNATCOLL.Tribooleans*, GNATColl provides a type that extends the classical *Boolean* type with an *Indeterminate* value.

There are various cases where such a type is useful. One example we have is when a user is doing a search (on a database or any set of data), and can specify some optional boolean criteria (“must the contact be french?”). He can choose to only see french people (“True”), to see no french people at all (“False”), or to get all contacts (“Indeterminate”). With a classical boolean, there is no way to cover all these cases.

Of course, there are more advanced use cases for such a type. To support these cases, the *Tribooleans* package overrides the usual logical operations “*and*”, “*or*”, “*xor*”, “*not*” and provides an *Equal* function.

See the specs of the package to see the truth tables associated with those operators.

GEOMETRY: PRIMITIVE GEOMETRIC OPERATIONS

GNATColl provides the package *GNATCOLL.Geometry*. This package includes a number of primitive operations on geometric figures like points, segments, lines, circles, rectangles and polygons. In particular, you can compute their intersections, the distances,...

This package is generic, so that you can specify the type of coordinates you wish to handle:

```
declare
  package Float_Geometry is new GNATCOLL.Geometry (Float);
  use Float_Geometry;

  P1 : constant Point := (1.0, 1.0);
  P2 : constant Point := (2.0, 3.0);
begin
  Put_Line ("Distance P1-P2 is" & Distance (P1, P2)'Img);
  -- Will print 2.23607
end;
```

Or some operations involving a polygon:

```
declare
  P3 : constant Point := (3.7, 2.0);
  P  : constant Polygon :=
    ((2.0, 1.3), (4.1, 3.0), (5.3, 2.6), (2.9, 0.7), (2.0, 1.3));
begin
  Put_Line ("Area of polygon:" & Area (P)); -- 3.015
  Put_Line ("P3 inside polygon ? " & Inside (P3, P)'Img); -- True
end;
```


PROJECTS: MANIPULATING GPR FILES

The package *GNATCOLL.Projects* provides an extensive interface to parse, manipulate and edit project files (.gpr files).

Although the interface is best used using the Ada05 notation, it is fully compatible with Ada95.

Here is a quick example on how to use the interface, although the spec file itself contains much more detailed information on all the subprograms related to the manipulation of project files:

```
pragma Ada_05;
with GNATCOLL.Projects; use GNATCOLL.Projects;
with GNATCOLL.VFS;      use GNATCOLL.VFS;

procedure Test_Project is
  Tree : Project_Tree;
  Files : File_Array_Access;
begin
  Tree.Load (GNATCOLL.VFS.Create ("path_to_project.gpr"));

  -- List the source files for project and all imported projects

  Files := Tree.Root_Project.Source_Files (Recursive => True);
  for F in Files'Range loop
    Put_Line ("File is: " & Files (F).Display_Full_Name);
  end loop;

  Tree.Unload;
end Test_Project;
```


REFCOUNT: REFERENCE COUNTING

Memory management is often a difficulty in defining an API. Should we let the user be responsible for freeing the types when they are no longer needed, or can we do it automatically on his behalf ?

The latter approach is somewhat more costly in terms of efficiency (since we need extra house keeping to know when the type is no longer needed), but provides an easier to use API.

Typically, such an approach is implemented using reference counting: all references to an object increment a counter. When a reference disappears, the counter is decremented, and when it finally reaches 0, the object is destroyed.

This approach is made convenient in Ada using controlled types. However, there are a number of issues to take care of to get things exactly right. In particular, the Ada Reference Manual specifies that *Finalize* should be idempotent: it could be called several times for a given object, in particular when exceptions occur.

An additional difficulty is task-safety: incrementing and decrementing the counter should be task safe, since the controlled object might be referenced from several task (the fact that other methods on the object are task safe or not is given by the user application, and cannot be ensured through the reference counting mechanism).

To make things easier, GNATColl provides the package *GNATCOLL.Refcount*. This package contains a generic child package.

To use it, you need to create a new tagged type that extends *GNATCOLL.Refcount.Refcounted*, so that it has a counter. Here is an example:

```
with GNATCOLL.Refcount; use GNATCOLL.Refcount;

package My_Pkg is
  type My_Type is new Refcounted with record
    Field1 : ...; -- Anything
  end record;

  package My_Type_Ptr is new Smart_Pointers (My_Type);
end My_Pkg;
```

The code above makes a *Ref* available. This is similar in semantics to an access type, although it really is a controlled type. Every time you assign the *Ref*, the counter is incremented. When the *Ref* goes out of scope, the counter is decremented, and the object is potentially freed.

Here an example of use of the package:

```
declare
  R : Ref;
  Tmp : My_Type := ...;
begin
  Set (R, Tmp); -- Increment counter
```

(continues on next page)

(continued from previous page)

```

    Get (R).Field1 := ...;  -- Access referenced object
end;

-- R out of scope, so decrement counter, and free Tmp

```

Although reference counting solves most of the issues with memory management, it can get tricky: when there is a cycle between two reference counted objects (one includes a reference to the other, and the other a reference to the first), their counter can never become 0, and thus they are never freed.

There are, however, common design patterns where this can severely interfere: imagine you want to have a *Map*, associating a name with a reference counted object. Typically, the map would be a cache of some sort. While the object exists, it should be referenced in the map. So we would like the Map to store a reference to the object. But that means the object will then never be freed while the map exists either, and memory usage will only increase.

The solution to this issue is to use *weak references*. These hold a pointer to an object, but do not increase its counter. As a result, the object can eventually be freed. At that point, the internal data in the weak reference is reset to *null*, although the weak reference object itself is still valid.

Here is an example:

```

with GNATCOLL.Refcount.Weakref;
use GNATCOLL.Refcount.Weakref;

type My_Type is new Weak_Refcounted with...;

package Pointers is new Weakref_Pointers (My_Type);

```

The above code can be used instead of the code in the first example, and provides the same capability (smart pointers, reference counted types,...). However, the type *My_Type* is slightly bigger, but can be used to create weak references:

```

WR : Weak_Ref;

declare
  R   : Ref;
  Tmp : My_Type := ...;
begin
  Set (R, Tmp);           -- Increment counter
  WR := Get_Weak_Ref (R); -- Get a weak reference

  Get (R).Field1 := ...;  -- Access referenced object
  Get (Get (WR)).Field1 := ...; -- Access through weak ref
end;

-- R out of scope, so decrement counter, and free Tmp

if Get (WR) /= Null_Ref then -- access to WR still valid
  -- Always true, since Tmp was freed
end if;

```

The example above is very simplified. Imagine, however, that you store *WR* in a map. Even when *R* is deallocated, the contents of the map remains accessible without a *Storage_Error* (although using *Get* will return *Null_Ref*, as above).

For task-safety issues, *Get* on a weak-reference returns a smart pointer. Therefore, this ensures that the object is never freed while that smart pointer object lives. As a result, we recommend the following construct in your code:

```
declare
  R : constant Ref := Get (WR);
begin
  if R /= Null_Ref then
    -- Get (R) never becomes null while in this block
  end if;
end;
```


READLINE: INTERACTIVE COMMAND LINE

gnatcoll provides an interface to the *readline* library.

License



The GNU *readline* library is licensed under the terms of the GNU General Public License, version 3. This means that if you want to use Readline in a program that you release or distribute to anyone, the program must be free software and have a GPL-compatible license.

When this library is available on the system, it provides various features to enhance command line support in tools. You also need to pass `-enable-gpl` to configure to indicate you understand the license of readline.

In particular, it provides various keybindings to make editing more comfortable than `Ada.Text_IO.Get_Line`. For instance, it is possible to use backspace to edit what you have just typed. It is also possible to move forward or backward by word, go to the start or end of line,...

readline also provides support for completion: by using the `tab` key, users can get all possible completions for the current word. This behavior is controllable from Ada, where your application can provide the list of completions.

Finally, readline comes with support for history. By using the `up` and `down` keys, the user can navigate the commands that were previously typed. It is also possible to preserve the history across sessions.

See the package *GNATCOLL.Readline* for more documentation on the API.

This package is always usable, whether or not the system supports readline. When readline is not available, the package falls back to a *Text_IO* based implementation (still providing input, although none of the advanced features described above will work).

CONFIG: PARSING CONFIGURATION FILES

gnatcoll provides a general framework for reading and manipulating configuration files. These files are in general static configuration for your application, and might be different from the preferences that a user might change interactively. However, it is possible to use them for both cases.

There are lots of possible formats for such configuration files: you could chose to use an XML file (but these are in general hard to edit manually), a binary file, or any other format. One format that is found very often is the one used by a lot of Windows applications (the `.ini` file format).

GNATCOLL.Config is independent from the actual format you are using, and you can add your own parsers compatible with the *GNATCOLL.Config* API. Out of the box, support is provided for `.ini` files, so let's detail this very simply format:

```
# A single-line comment
[Section1]
key1 = value
key2=value2

[Section2]
key1 = value3
```

Comments are (by default) started with '#' signs, but you can configure that and use any prefix you want. The (*key*, *value*) pairs are then organized into optional sections (if you do not start a section before the first key, that key will be considered as part of the "" section). A section then extends until the start of the next section.

The values associated with the various keys can be strings, integers or booleans. Spaces on the left and right of the values and keys are trimmed, and therefore irrelevant.

Support is providing for interpreting the values as file or directory names. In such a case, if a relative name is specified in the configuration file it will be assumed to be relative to the location of the configuration file (by default, but you can also configure that).

GNATCOLL.Config provides an abstract iterator over a config stream (in general, that stream will be a file, but you could conceptually read it from memory, a socket, or any other location). A specific implementation is provided for file-based streams, which is further specialized to parse `.ini` files.

Reading all the values from a configuration file is done with a loop similar to:

```
declare
  C : INI_Parser;
begin
  Open (C, "settings.txt");
  while not At_End (C) loop
    Put_Line ("Found key " & Key (C) & " with value " & Value (C));
    Next (C);
```

(continues on next page)

(continued from previous page)

```
end loop;  
end;
```

This can be made slightly lighter by using the Ada05 dotted notation.

You would only use such a loop in your application if you intend to store the values in various typed constants in your application. But *GNATCOLL.Config* provides a slightly easier interface for this, in the form of a *Config_Pool*. Such a pool is filled by reading a configuration file, and then the values associated with each key can be read at any point during the lifetime of your application. You can also explicitly override the values when needed:

```
Config : Config_Pool;  -- A global variable  
  
declare  
  C : INI_Parser;  
begin  
  Open (C, "settings.txt");  
  Fill (Config, C);  
end;  
  
Put_Line (Config.Get ("section.key"));  -- Ada05 dotted notation
```

Again, the values are by default read as strings, but you can interpret them as integers, booleans or files.

A third layer is provided in *GNATCOLL.Config*. This solves the issue of possible typos in code: in the above example, we could have made a typo when writing “*section.key*”. That would only be detected at run time. Another issue is that we might decide to rename the key in the configuration file. We would then have to go through all the application code to find all the places where this key is references (and that can’t be based on cross-references generated by the compiler, since that’s inside a string).

To solve this issue, it is possible to declare a set of constants that represent the keys, and then use these to access the values, solving the two problems above:

```
Section_Key1 : constant Config_Key := Create ("Key1", "Section");  
Section_Key2 : constant Config_Key := Create ("Key2", "Section");  
  
Put_Line (Section_Key1.Get);
```

You then access the value of the keys using the Ada05 dotted notation, providing a very natural syntax. When and if the key is renamed, you then have a single place to change.

POOLS: CONTROLLING ACCESS TO RESOURCES

The package **GNATCOLL.Pools** provides resource pools.

A pool contains a maximum number of resources, which are created on demand. However, once a resource is no longer needed by the client, it is not freed, but instead it is released to the pool, which will then return it again the next time a client requests a resource.

The typical resource is when the creation of the resources is expensive, for instance a connection to a database or a remote server. The lazy creation then provides a faster startup time (as well as more flexibility, since there is no need to allocate dozens of resources if only one will be needed in the end), and more efficient retrieval through the reuse of resources.

The pool in this package is task safe, and is intended as a global variable (or field of a global variable) somewhere in your application.

The resources are implemented as reference-counted types (through *GNATCOLL.Refcount*). As a result, as soon as the client no longer has a handle on them, they are automatically released to the pool and there is no risk that the client forgets to do so.

GNATCOLL.Pools is a generic package where the formal parameters describe the type of resources, how to create them on demand, what should happen when a resource is released, and finally how to free a resource when the pool itself is freed. See `gnatcoll-pools.ads` for a full and up-to-date description of these parameters.

JSON: HANDLING JSON DATA

JSON is a format often used on the web to communicate between a server and a browser, or between servers. It plays a similar role to XML, but is much lighter in terms of size. On the other hand, it doesn't provide advanced features like validation which XML provides.

The package **GNATCOLL.JSON** provides an Ada for creating JSON data, or parse such data that your application receives.

Most JSON data will generally start with an object, on which attributes can be set. The value for the attributes are also JSON data.

Here is an example of use:

```
pragma Ada_05;
with GNATCOLL.JSON;   use GNATCOLL.JSON;
with Ada.Text_IO;     use Ada.Text_IO;

procedure JSON_Test is
  MyObj : JSON_Value := Create_Object;
begin
  MyObj.Set_Field ("field1", Create (1));
  MyObj.Set_Field ("name", "theName");

  -- Now print the value
  Put_Line (MyObj.Write);
end JSON_Test;
```

This example used the Ada05 dot notation to call the primitive operations, but would also work using the more traditional prefix notation.

It is also possible to create JSON arrays. These are not tagged types, so the prefix notation has to be used. Here is a further example that sets another field in the object we had before:

```
declare
  MyArr : JSON_Array := Empty_Array;
begin
  Append (MyArr, Create (1));
  Append (MyArr, Create ("aString"));

  MyObj.Set_Field ("vals", MyArr);
end;
```

GNATColl automatically takes care of memory management, and all allocated memory is automatically freed when the object is no longer needed.

ICONV: CONVERTING BETWEEN CHARACTER ENCODINGS

This package provides a binding to the libiconv library. This library is standard on most Unix systems. When it is not provided by the system, the GNU libiconv package can be installed instead.

23.1 Compiling

Since this is a binding to an external library, it is optional. By default, GNATCOLL will automatically detect whether the package is available on the system.

If the library is not found, a dummy version of `gnatcoll-iconv.adb` is provided, where the input is always returned unmodified. This ensures that an application can use the API whether or not the library is available, although of course the behavior will be different in both cases.

If your application depends on having a working libiconv, you can specify `-with-iconv` or `-with-iconv=PATH` to GNATCOLL's configure. This will ensure that configure fails if iconv is not found on the system. Specifying an explicit path is recommended in general. If you specifying an explicit path, and also include the `-disable-shared` switch, GNATCOLL will always use a static libiconv found in the given path.

An alternative is to use `-with-iconv=static` to force the use of a static library for iconv. This might be needed on some systems if you have link errors in your application: some versions of libiconv define symbols such as `__iconv_open` whereas others define `__libiconv_open`, and it is mandatory to use the include file corresponding to the linked version of libiconv. For instance, on OSX, the system provides `/usr/include/iconv.h` and `/usr/lib/libiconv.dylib`, but you might also have the GNU version of libiconv in `/opt/local/include/iconv.h` and `/opt/local/lib/libiconv.dylib`. Unfortunately, these directories might also contain other libraries you depend and it sometimes happens that `/usr/include/iconv.h` is used when building, and `/opt/local/lib/libiconv.dylib` when linking. These are incompatible. Using shared libraries solves this issue.

On the other hand, you can avoid compiling GNATCOLL.Iconv altogether by specifying `-without-iconv`. This however does not save much.

23.2 Using GNATCOLL.Iconv

Since this is a binding to an external library, GNATCOLL installs a separate project file for it so that applications must explicitly declare they will use that package (and so that applications that do not use this package do not inherit extra external dependencies).

You need to start your project file with:

```
with "gnatcoll_iconv";  
project Default is
```

(continues on next page)

(continued from previous page)

```
...  
end Default;
```

23.3 API

The whole API is documented in `gnatcoll-iconv.ads`. Here is a simple code sample that converts from iso-8859-1 encoding to UTF8:

```
with GNATCOLL.Iconv;   use GNATCOLL.Iconv;  
procedure Main is  
  EAcute : constant Character := Character'Val (16#E9#);  
  -- in iso-8859-1  
  
  Result : constant String := Iconv  
    ("Some string " & EAcute,  
     To_Code => UTF8,  
     From_Code => ISO_8859_1);  
begin  
  null;  
end Main;
```

A more advanced (and somewhat more efficient) API is available via the *Iconv* procedure. In that procedure, you control the input and output buffers, so you will need less overall memory when you are converting big buffers.

SQL: DATABASE INTERFACE

A lot of applications need to provide **persistence** for their data (or a part of it). This means the data needs to be somehow saved on the disk, to be read and manipulated later, possibly after the application has been terminated and restarted. Although Ada provides various solutions for this (including the use of the streams as declared in the Ada Reference Manual), the common technique is through the use of relational database management systems (**RDBMS**). The term database is in fact overloaded in this context, and has come to mean different things:

- The software system that implements file and query management. This is generally provided by a third-party. The common abbreviation for these is **DBMS**. Queries are generally written in a language called **SQL**. One of the issues is that each DBMS tends to make minor changes to this language. Another issue is that the way to send these SQL commands to the DBMS is vendor-specific. GNATColl tries to abstract this communication through its own API. Optional components can instantiate this framework to specific DBMS, for example **PostgreSQL** and **sqlite**. Common API makes it relatively easy to change between two systems. For instance, development could be done using a local sqlite DBMS, and then deployed (after testing, of course!) on a PostgreSQL system.

The code in GNATColl is such that adding support for a new DBMS should be relatively easy.

- A place where an application stores its data. The term **database** in this document refers to this meaning. In a relational database, this place is organized into tables, each of which contains a number of fields. A row in a table represents one object. The set of tables and their fields is called the **schema** of the database.

Traditionally, writing the SQL queries is done inline: special markers are inserted into your code to delimit sections that contain SQL code (as opposed to Ada code), and these are then preprocessed to generate actual code. This isn't the approach chosen in GNATColl: there are several drawbacks, in particular your code is no longer Ada and various tools will choke on it.

The other usual approach is to write the queries as strings, which are passed, via a DBMS-specific API, to the DBMS server. This approach is very fragile:

- The string might not contain **well-formed SQL**. This will unfortunately only be detected at run time when the DBMS complains.
- This is not **type safe**. You might be comparing a text field with an integer, for instance. In some cases, the DBMS will accept that (sqlite for instance), but in some other cases it won't (PostgreSQL). The result might then either raise an error, or return an empty list.
- There is a risk of **SQL injection**. Assuming the string is constructed dynamically (using Ada's & operator), it might be easy for a user to pass a string that breaks the query, and even destroys things in the database.
- As discussed previously, the SQL code might not be **portable** across DBMS. For instance, creating an automatically increment integer primary key in a table is DBMS specific.
- The string is fragile if the database **schema changes**. Finding whether a schema change impacts any of the queries requires looking at all the strings in your application.
- **performance** might be an issue. Whenever you execute a query, the DBMS will analyze it, decide how to execute it (for instance, whether it should traverse all the rows of a table, or whether it can do a faster lookup), and then

retrieve the results. The analysis pass is typically slow (relatively the overall execution time), and queries can in fact be **prepared** on the server: they are then analyzed only once, and it is possible to run them several times without paying the price of the analysis every time. Such a query can also be **parameterized**, in that some values can be deferred until the query is actually executed. All the above is made easy and portable in GNATColl, instead of requiring DBMS-specific techniques.

- This might require **large amount of code** to setup the query, bind the parameters, execute it, and traverse the list of results.

GNATColl attempts to solve all these issues. It also provides further performance improvements, for instance by keeping connections to the DBMS open and reusing them when possible. A paper was published at the Ada-Europe conference in 2008 which describes the various steps we went through in the design of this library.

24.1 Database Abstraction Layers

GNATColl organizes the API into several layers, each written on top of the previous one and providing more abstraction. All layers are compatible with each other and can be mixed inside a given application.

- **low-level binding.**

This API is DBMS-specific, and is basically a mapping of the C API provided by the DBMS vendors into Ada. If you are porting C code, or working with an existing application, as a way to start using GNATColl before moving to higher levels of abstraction.

The code is found in `gnatcoll-sql-sqlite-gnade.ads` and `gnatcoll-sql-postgres-gnade.ads`. The *gnade* part in the file names indicate that this code was initially inspired by the **GNADE** library that used to be available on the internet. Part of the code might in fact come from that library.

Using this API requires writing the SQL queries as strings, with all the disadvantages that were highlighted at the beginning of this chapter.

- **GNATCOLL.SQL and GNATCOLL.SQL.Exec**

The first of these packages makes it possible to write type-safe queries strongly linked to the database schema (thus with a guarantee that the query is up-to-date with regards to the schema). To accomplish this, it also relies on code that is generated automatically from a description of your database schema, using the tool *gnatcoll_db2ada*. To simplify memory management, the queries are automatically referenced counted and freed when they are no longer needed.

The second of these packages provides communication with the DBMS. It provides a vendor-neutral API. You can send your queries either as strings, or preferably as written with *GNATCOLL.SQL*. It also provides a simple way to prepare parameterized statements on the server for maximum efficiency, as well as the reuse of existing DBMS connections. It provides a simple API to retrieve and manipulate the results from a query.

- **GNATCOLL.SQL.ORM and GNATCOLL.SQL.Sessions**

This is an Object-Relational Mapping (ORM).

The first of these packages makes it possible to manipulate a database without writing SQL. Instead, you manipulate Ada objects (tagged types), whose primitive operations might transparently execute SQL queries. This API provides caching for maximum efficiency. It relies on code automatically generated by *gnatcoll_db2ada* from the schema of your database. The generated objects can then be extended in your own code if needed.

The second package encapsulates DBMS connections into higher-level objects which provide their own caching and work best with the ORM objects. A session is automatically released to a pool when no longer needed and will be reused later on.

The following sections will ignore the lower layer, and concentrate on the other layers. They share a number of types and, again, are fully compatible with each other. You could connect to the database, and then write some queries using `GNATCOLL.SQL` and some using `GNATCOLL.SQL.ORM`.

24.2 Database example

This section describes an example that will be extended throughout this chapter. We will build an application that represents a library. Such a library contains various media (books and DVDs for instance), and customers. A customer can borrow multiple media at the same time, but a media is either at a customer's, or still in the library.

The GNATColl distribution includes an example directory which contains all the code and data for this example.

24.3 Database schema

As was mentioned earlier (*Database Abstraction Layers*), GNATColl relies on automatic code generation to provide a type safe interface to your database. This code is generated by an external tool called *gnatcoll_db2ada* provided as an optional component. In some cases, this tool requires an installation of python (www.python.org) on your machine, since part of the code is written in that language.

This tool is able to output various kind of information, and is fully described in the corresponding component. However, the input is always the same: this is the schema of your database, that is the list of tables and fields that make up your database. There exist two ways to provide that information:

- From a running database

If you pass the DBMS vendor (postgresql, sqlite,...) and the connection parameters to *gnatcoll_db2ada*, it is able to query the schema on its own. However, this should not be the preferred method: this is similar to reverse engineering assembly code into the original high-level code, and some semantic information will be missing. For instance, in SQL we have to create tables just to represent the many-to-many relationships. These extra tables are part of the implementation of the schema, but are just noise when it comes to the semantics of the schema. For this reason, it is better to use the second solution below:

- From a textual description

Using the *-dbmodel* switch to *gnatcoll_db2ada*, you can pass a file that describes the schema. We do not use SQL as the syntax in this, because as explained above this is too low-level. This text file also provides additional capabilities that do not exist when reverse-engineering an existing database, for instance the ability to use name to represent reverse relationships for foreign keys (see below and the ORM).

The most convenient editor for this file is Emacs, using the *org-mode* which provides convenient key shortcuts for editing the contents of ASCII tables. But any text editor will do, and you do not need to align the columns in this file.

All lines starting with a hash sign (`#`) will be ignored.

This file is a collection of ASCII tables, each of which relates to one table or one SQL view in your database. The paragraphs start with a line containing:

```
table ::=
'|' ('ABSTRACT')? ('TABLE'|'VIEW') ['(' supertable ')']
'|' <name> '|' <name_row>
```

“name” is the name of the table. The third pipe and third column are optional, and should be used to specify the name for the element represented by a single row. For instance, if the table is called “books”, the third column could contain “book”. This is used when generating objects for use with *GNATCOLL.SQL.ORM*.

If the first line starts with the keyword *ABSTRACT*, then no instance of that table actually exists in the database. This is used in the context of table inheritance, so define shared fields only once among multiple tables.

The keyword *TABLE* can be followed by the name of a table from which it inherits the fields. Currently, that supertable must be abstract, and the fields declared in that table are simply duplicated in the new table.

Following the declaration of the table, the file then describe their fields, each on a separate line. Each of these lines must start with a pipe character (“|”), and contain a number of pipe-separated fields. The order of the fields is always given by the following grammar:

```
fields ::=
  '|' <name> '|' <type>
  '|' ('PK'|'|' 'NULL'|'NOT NULL'|'INDEX'|'UNIQUE'|'NOCASE')
  '|' [default] '|' [doc] '|'
```

The type of the field is the SQL type (“INTEGER”, “TEXT”, “TIMESTAMP”, “DATE”, “DOUBLE PRECISION”, “MONEY”, “BOOLEAN”, “TIME”, “CHARACTER(1)”). Any maximal length can be specified for strings, not just 1 as in this example. The tool will automatically convert these to Ada when generating Ada code. A special type (“AUTOINCREMENT”) is an integer that is automatically incremented according to available ids in the table. The exact type used will depend on the specific DBMS.

The property ‘NOCASE’ indicates that comparison should be case insensitive for this field.

If the field is a foreign key (that is a value that must correspond to a row in another table), you can use the special syntax for its type:

```
fk_type ::= 'FK' <table_name> [ '(' <reverse_name> ')' ]
```

As you can see, the type of the field is not specified explicitly, but will always be that of the foreign table’s primary key. With this syntax, the foreign table must have a single field for its primary key. GNATColl does not force a specific order for the declaration of tables: it is valid to have a foreign key to a table that hasn’t been declared yet. There is however a restriction if you use the model to create a sqlite database (through the *-createdb* switch of *gnatcoll_db2ada*): in this case, a reference to a table that hasn’t been defined yet may not be through a field marked as NOT NULL. This is a limitation of the sqlite backend itself. The solution in this case is to reorder the declaration of tables, or drop the NOT NULL constraint.

Another restriction is that a foreign key that is also a primary key must reference a table that has already been defined. You need to reorder the declaration of your tables to ensure this is the case.

“reverse_name” is the optional name that will be generated in the Ada code for the reverse relationship, in the context of *GNATCOLL.SQL.ORM*. If the “reverse_name” is empty (the parenthesis are shown), no reverse relationship is generated. If the parenthesis and the reverse_name are both omitted, a default name is generated based on the name of the field.

The third column in the fields definition indicates the constraints of the type. Multiple keywords can be used if they are separated by commas. Thus, “NOT NULL, INDEX” indicates a column that must be set by the user, and for which an index is created to speed up look ups.

- A primary key (“PK”)
- The value must be defined (“NOT NULL”)
- The value can be left undefined (“NULL”)
- A unique constraint and index (“UNIQUE”)
- An index should be created for that column (“INDEX”) to speed up the lookups.
- The automatic index created for a Foreign Key should not be created (“NOINDEX”). Every time a field references another table, GNATColl will by default create an index for it, so that the ORM can more efficiently do a reverse query (from the target table’s row find all the rows in the current table that reference

that target row). This will in general provide more efficiency, but in some cases you never intend to do the reverse query and thus can spare the extra index.

The fourth column gives the default value for the field, and is given in SQL syntax. Strings must be quoted with single quotes.

The fifth column contains documentation for the field (if any). This documentation will be included in the generated code, so that IDEs can provide useful tooltips when navigating your application's code.

After all the fields have been defined, you can specify extract constraints on the table. In particular, if you have a foreign key to a table that uses a tuple as its primary key, you can define that foreign key on a new line, as:

```
FK ::= '|' "FK:" '|' <table> '|' <field_names>*
      '|' <field_names>* '|'
```

For instance:

```
| TABLE | tableA |
| FK: | tableB | fieldA1, fieldA2 | fieldB1, fieldB2 |
```

It is also possible to create multi-column indexes, as in the following example. In this case, the third column contains the name of the index to create. If left blank, a default name will be computed by GNATColl:

```
| TABLE | tableA |
| INDEX: | field1,field2,field3 | name |
```

The same way the unique multi-column constraint and index can be created. The name is optional.

```
TABLE | tableA |
UNIQUE: | field1,field2,field3 | name |
```

Going back to the example we described earlier (*Database example*), let's describe the tables that are involved.

The first table contains the customers. Here is its definition:

TABLE	customers	customer	The customer for the library
id	AUTOINCREMENT	PK	Auto-generated id
first	TEXT	NOT NULL	Customers' first name
last	TEXT	NOT NULL , INDEX	Customers' last name

We highly recommend to set a primary key on all tables. This is a field whose value is unique in the table, and thus that can act as an identifier for a specific row in the table (in this case for a specific customer). We recommend using integers for these ids for efficiency reasons. It is possible that the primary key will be made of several fields, in which case they should all have the "PK" constraint in the third column.

A table with no primary key is still usable. The difference is in the code generated for the ORM (*The Object-Relational Mapping layer (ORM)*), since the *Delete* operation for this table will raise a *Program_Error* instead of doing the actual deletion (that's because there is no guaranteed unique identifier for the element, so the ORM does not know which one to delete – we do not depend on having unique internal ids on the table, like some DBMS have). Likewise, the elements extracted from such a primary key-less table will not be cached locally in the session, and cannot be updated (only new elements can be created in the table).

As we mentioned, the library contains two types of media, books and DVDs. Each of those has a title, an author. However, a book also has a number of pages and a DVD has a region where it can be viewed. There are various ways to represent this in a database. For illustration purposes, we will use table inheritance here: we will declare one abstract table (media) which contains the common fields, and two tables to represent the types of media.

As we mentioned, a media can be borrowed by at most one customer, but a customer can have multiple media at any point in time. This is called a **one-to-many** relationship. In SQL, this is in general described through the

use of a foreign key that goes from the table on the “many” side. In this example, we therefore have a foreign key from media to customers. We also provide a name for the reverse relationship, which will become clearer when we describe the ORM interface.

Here are the declarations:

```
| ABSTRACT TABLE | media | media || The contents of the library |
| id | AUTOINCREMENT | PK || Auto-generated id |
| title | TEXT | || The title of the media |
| author | TEXT | || The author |
| published | DATE | || Publication date |
| borrowed_by | FK customers(items) | NULL || Who borrowed the media |

| TABLE (media) | books | book | | The books in the library |
| pages | INTEGER | | 100 |

| TABLE (media) | dvds | dvd | | The dvds in the library |
| region | INTEGER | | 1 |
```

For this example, all this description is put in a file called `dbschema.txt`.

24.4 Connecting to the database

This library abstracts the specifics of the various database engines it supports. Ideally, code written for one database could be ported almost transparently to another engine. This is not completely doable in practice, since each system has its own SQL specifics, and unless you are writing things very carefully, the interpretation of your queries might be different from one system to the next.

However, the Ada code should remain untouched if you change the engine. Various engines are supported out of the box (PostgreSQL and SQLite), although new ones can be added by overriding the appropriate SQL type (*Database_Connection*). When you compile GNATColl, the build scripts will try and detect what systems are installed on your machine, and only build support for those. It is possible, if no database was installed on your machine at that time, that the database interface API is available (and your application compiles), but no connection can be done to database at run time.

To connect to a DBMS, you need to specify the various connection parameters. This is done via a *GNATCOLL.SQL.Exec.Database_Description* object. The creation of this object depends on the specific DBMS you are connecting to (and this is the only part of your code that needs to know about the specific system). The packages *GNATCOLL.SQL.Postgres* and *GNATCOLL.SQL.SQLite* contain a *Setup* function, whose parameters depend on the DBMS. They provide full documentation for their parameters. Let’s take a simple example from sqlite:

```
with GNATCOLL.SQL.SQLite;  -- or Postgres
declare
    DB_Descr : GNATCOLL.SQL.Exec.Database_Description;
begin
    DB_Descr := GNATCOLL.SQL.SQLite.Setup ("dbname.db");
end
```

At this point, no connection to the DBMS has been done, and no information was exchanged.

To communicate with the database, however, we need to create another object, a **GNATCOLL.SQL.Exec.Database_Connection**. Your application can create any number of these. Typically, one would create one such connection per task in the application, although other strategies are possible (like a pool of reusable connections, where a task might be using two connections and another task none at any point in time).

If you do not plan on using the ORM interface from **GNATCOLL.SQL.ORM**, GNATColl provides a simple way to create a task-specific connection. While in this task, the same connection will always be returned (thus you do not have to pass it around in parameter, although the latter might be more efficient):

```
declare
  DB : GNATCOLL.SQL.Exec.Database_Connection;
begin
  DB := GNATCOLL.SQL.Exec.Get_Task_Connection
    (Description => DB_Descr);
end;
```

If your application is not multi-tasking, or you wish to implement your own strategy for a connection pool, you can also use the following code (using Ada 2005 dotted notation when calling the primitive operation). This code will always create a new connection, not reuse an existing one, as opposed to the code above:

```
declare
  DB : GNATCOLL.SQL.Exec.Database_Connection;
begin
  DB := DB_Descr.Build_Connection;
end;
```

A note on concurrency: if you implement your own pool, you might sometimes end up with dead locks when using sqlite. If a task uses two or more connections to sqlite, and you setup GNATCOLL to create SQL transactions even for *SELECT* statements (see *GNATCOLL.SQL.Sqlite.Always_Use_Transactions*), the following scenario will result in a deadlock:

```
DB1 := ... new connection to sqlite
... execute a SELECT through DB1. The latter then holds a shared
... lock, preventing other connections from writing (but not from
... reading).
DB2 := ... another connection in the same thread
... execute an INSERT through DB2. This tries to get a lock, which
... will fail while DB1 holds the shared lock. Since these are in
... the same thread, this will deadlock.
```

By default, GNATCOLL will not create SQL transactions for select statements to avoid this case, which occurs frequently in code.

If you wish to reuse an existing connection later on, you must reset it. This terminates any on-going SQL transaction, and resets various internal fields that describe the state of the connection:

```
Reset_Connection (DB);
```

In all three cases, the resulting database connection needs to be freed when you no longer needed (which might be when your program terminates if you are using pools) to avoid memory leaks. Nothing critical will appear if you do not close, though, because the transactions to the DBMS server are saved every time you call *Commit* in any case. So the code would end with:

```
Free (DB); -- for all connections you have opened
Free (DB_Descr);
```

At this point, there still hasn't been any connection to the DBMS. This will be done the first time a query is executed. If for some reason the connection to the DBMS server is lost, GNATColl will automatically attempt to reconnect a number of times before it gives up. This might break if there was an ongoing SQL transaction, but simplifies your code since you do not have to handle reconnection when there was a network failure, for instance.

As we saw before, the database interface can be used in multi-tasking applications. In such a case, it is recommended that each thread has its own connection to the database, since that is more efficient and you do not have to handle locking. However, this assumes that the database server itself is thread safe, which most often is the case, but not for *sqlite* for instance. In such a case, you can only connect one per application to the database, and you will have to manage a queue of queries somehow.

If you want to use `GNATCOLL.SQL.Sessions` along with the Object-Relational Mapping API, you will need to initialize the connection pool with the `Database_Description`, but the session will then take care automatically of creating the `Database_Connection`. See later sections for more details.

24.5 Loading initial data in the database

We have now created an empty database. To make the queries we will write later more interesting, we are going to load initial data.

There are various ways to do it:

- Manually or with an external tool

One can connect to the database with an external tool (a web interface when the DBMS provides one for instance), or via a command line tool (*psql* for PostgreSQL or *sqlite3* for SQLite), and start inserting data manually. This shows one of the nice aspects of using a standard DBMS for your application: you can alter the database (for instance to do minor fixes in the data) with a lot of external tools that were developed specifically for that purpose and that provide a nice interface. However, this is also tedious and error prone, and can't be repeat easily every time we recreate the database (for instance before running automatic tests).

- Using `GNATCOLL.SQL.EXEC`

As we will describe later, GNATColl contains all the required machinery for altering the contents of the database and creating new objects. Using `GNATCOLL.SQL.ORM` this can also be done at a high-level and completely hide SQL.

- Loading a data file

A lot of frameworks call such a file that contains initial data a “fixture”. We will use this techniques as an example. At the Ada level, this is a simple call to `GNATCOLL.SQL.Inspect.Load_Data`. The package contains a lot more than just this subprogram (The_gnatcoll_db2ada_tool):

```
declare
  File : GNATCOLL.VFS.Virtual_File := Create ("fixture.txt");
  DB : Database_Connection; -- created earlier
begin
  GNATCOLL.SQL.Inspect.Load_Data (DB, File);
  DB.Commit;
end;
```

The format of this file is described just below.

As we mentioned, GNATColl can load data from a file. The format of this file is similar to the one that describes the database schema. It is a set of ASCII tables, each of which describes the data that should go in a table (it is valid to duplicate tables). Each block starts with two lines: The first one has two mandatory columns, the first of which contains the text “TABLE”, and the second contains the name of the table you want to fill. The second line should contain as many columns as there are fields you want to set. Not all the fields of the table need to have a corresponding column if you want to set their contents to NULL (provided, of course, that your schema allows it). For instance, we could add data for our library example as such:

TABLE	customers			
id	first	last		
-----/				
1	John	Smith		
2	Alain	Dupont		

TABLE	books			
title	author	pages	published	borrowed_by
-----/				
Art of War	Sun Tzu	90	01-01-2000	1
Ada RM	WRG	250	01-07-2005	

A few comments on the above: the *id* for *books* is not specified, although the column is the primary key and therefore cannot be NULL. In fact, since the type of the *id* was set to AUTOINCREMENT, GNATColl will automatically assign valid values. We did not use this approach for the *id* of *customers*, because we need to know this *id* to set the *borrowed_by* field in the *books* table.

There is another approach to setting the *borrowed_by* field, which is to give the value of another field of the *customers* table. This of course only work if you know this value is unique, but that will often be the case in your initial fixtures. Here is an example:

TABLE	dvds			
title	author	region	borrowed_by(&last)	
-----/				
The Birds	Hitchcock	1	&Smith	
The Dictator	Chaplin	3	&Dupont	

Here, the title of the column indicates that any value in this column might be a reference to the *customers.last* value. Values which start with an ampersand ("&") will therefore be looked up in *customers.last*, and the *id* of the corresponding customer will be inserted in the *dvds* table. It would still be valid to use directly customer ids instead of references, this is just an extra flexibility that the references give you to make your fixtures more readable.

However, if we are using such references we need to provide the database schema to *Load_Data* so that it can write the proper queries. This is done by using other services of the *GNATCOLL.SQL.Inspect* package.

The code for our example would be:

```
Load_Data
(DB, Create ("fixture.txt"),
 New_Schema_IO (Create ("dbschema.txt")).Read_Schema);
```

24.6 Writing queries

The second part of the database support in GNATColl is a set of Ada subprograms which help write SQL queries. Traditional ways to write such queries have been through embedded SQL (which requires a preprocessing phase and complicates the editing of source files in Ada-aware editors), or through simple strings that are passed as is to the server. In the latter case, the compiler can not do any verification on the string, and errors such a missing parenthesis or misspelled table or field names will not be detected until the code executes the query.

GNATColl tries to make sure that code that compiles contains syntactically correct SQL queries and only reference existing tables and fields. This of course does not ensure that the query is semantically correct, but helps detect trivial errors as early as possible.

Such queries are thus written via calls to Ada subprograms, as in the following example:

```
with GNATCOLL.SQL; use GNATCOLL.SQL;
with Database; use Database;
declare
  Q : SQL_Query;
begin
  Q := SQL_Select
    (Fields => Max (Ticket_Priorities.Priority)
      & Ticket_Priorities.Category,
    From   => Ticket_Priorities,
    Where  => Ticket_Priorities.Name /= "low",
    Group_By => Ticket_Priorities.Category);
end;
```

The above example will return, for each type of priority (internal or customer) the highest possible value. The interest of this query is left to the user...

This is very similar to an actual SQL query. Field and table names come from the package that was automatically generated by the *gnatcoll_db2ada* tool, and therefore we know that our query is only referencing existing fields. The syntactic correctness is ensured by standard Ada rules. The *SQL_Select* accepts several parameters corresponding to the usual SQL attributes like *GROUP BY*, *HAVING*, *ORDER BY* and *LIMIT*.

The *From* parameter could be a list of tables if we need to join them in some ways. Such a list is created with the overridden “&” operator, just as for fields which you can see in the above example. GNATColl also provides a *Left_Join* function to join two tables when the second might have no matching field (see the SQL documentation).

Similar functions exist for *SQL_Insert*, *SQL_Update* and *SQL_Delete*. Each of those is extensively documented in the *gnatcoll-sql.ads* file.

It is worth noting that we do not have to write the query all at once. In fact, we could build it depending on some other criteria. For instance, imagine we have a procedure that does the query above, and omits the priority specified as a parameter, or shows all priorities if the empty string is passed. Such a procedure could be written as:

```
procedure List_Priorities (Omit : String := "") is
  Q : SQL_Query;
  C : SQL_Criteria := No_Criteria;
begin
  if Omit /= "" then
    C := Ticket_Priorities.Name /= Omit;
  end if;
  Q := SQL_Select
    (Fields => ..., -- as before
    Where  => C);
end;
```

With such a code, it becomes easier to create queries on the fly than it would be with directly writing strings.

The above call has not sent anything to the database yet, only created a data structure in memory (more precisely a tree). In fact, we could be somewhat lazy when writing the query and rely on auto-completion, as in the following example:

```
Q := SQL_Select
  (Fields => Max (Ticket_Priorities.Priority)
    & Ticket_Priorities.Category,
  Where  => Ticket_Priorities.Name /= "low");

Auto_Complete (Q);
```

This query is exactly the same as before. However, we did not have to specify the list of tables (which GNATColl can compute on its own by looking at all the fields referenced in the query), nor the list of fields in the *GROUP BY* clause, which once again can be computed automatically by looking at those fields that are not used in a SQL aggregate function. This auto-completion helps the maintenance of those queries.

There is another case where GNATColl makes it somewhat easier to write the queries, and that is to handle joins between tables. If your schema was build with foreign keys, GNATColl can take advantage of those.

Going back to our library example, let's assume we want to find out all the books that were borrowed by the user "Smith". We need to involve two tables (*Books* and *Customers*), and provide a join between them so that the DBMS knows how to associate the rows from one with the rows from the other. Here is a first example for such a query:

```
Q := SQL_Select
  (Fields => Books.Title & Books.Pages,
   From   => Books & Customers,
   Where  => Books.Borrowed_By = Customers.Id
           and Customers.Last = "Smith");
```

In fact, we could also use auto-completion, and let GNATColl find out the involved tables on its own. We thus write the simpler:

```
Q := SQL_Select
  (Fields => Books.Title & Books.Pages,
   Where  => Books.Borrowed_By = Customers.Id
           and Customers.Last = "Smith");
```

There is one more things we can do to simplify the query and make it more solid if the schema of the database changes. For instance, when a table has a primary key made up of several fields, we need to make sure we always have an "=" statement in the WHERE clause for all these fields between the two tables. In our example above, we could at some point modify the schema so that the primary key for *customers* is multiple (this is unlikely in this example of course). To avoid this potential problems and make the query somewhat easier to read, we can take advantage of the FK subprograms generated by *gnatcoll_db2ada*. Using the Ada05 dotted notation for the call, we can thus write:

```
Q := SQL_Select
  (Fields => Books.Title & Books.Pages,
   Where  => Books.FK (Customers)
           and Customers.Last = "Smith");
```

Regarding memory management, there is no need for explicitly freeing memory in the above code. GNATColl will automatically do this when the query is no longer needed.

24.7 Executing queries

Once we have our query in memory, we need to pass it on to the database server itself, and retrieve the results.

Executing is done through the *GNATCOLL.SQL.Exec* package, as in the following example:

```
declare
  R : Forward_Cursor;
begin
  R.Fetch (Connection => DB, Query => Q);
end;
```

This reuses the connection we have established previously (*DB*) (although now we are indeed connecting to the DBMS for the first time) and sends it the query. The result of that query is then stored in *R*, to be used later.

Some SQL commands execute code on the DBMS, but do not return a result. In this case, you can use *Execute* instead of *Fetch*. This is the case when you execute an *INSERT* or *UPDATE* statement for instance. Using *Execute* avoids the need to declare the local variable *R*.

If for some reason the connection to the database is no longer valid (a transient network problem for instance), GNATColl will attempt to reconnect and re-execute your query transparently, so that your application does not need to handle this case.

We'll describe later (*Getting results*) how to analyze the result of the query.

Some versions of *Fetch* have an extra parameter *Use_Cache*, set to *False* by default. If this parameter is true, and the exact same query has already been executed before, its result will be reused without even contacting the database server. The cache is automatically invalidated every hour in any case. This cache is mostly useful for tables that act like enumeration types. In this case, the contents of the table changes very rarely, and the cache can provide important speedups, whether the server is local or distant. However, we recommend that you do actual measurements to know whether this is indeed beneficial for you. You can always invalidate the current cache with a call to *Invalidate_Cache* to force the query to be done on the database server.

If your query produces an error (whether it is invalid, or any other reason), a flag is toggled in the *Connection* parameter, which you can query through the *Success* subprogram. As a result, a possible continuation of the above code is:

```
if Success (DB) then
  ...
else
  ... -- an error occurred
end if
```

GNATColl also tries to be helpful in the way it handles SQL transactions. Such transactions are a way to execute your query in a sandbox, i.e. without affecting the database itself until you decide to *COMMIT* the query. Should you decide to abort it (or *ROLLBACK* as they say for SQL), then it is just as if nothing happened. As a result, it is in general recommended to do all your changes to the database from within a transaction. If one of the queries fail because of invalid parameters, you just rollback and report the error to the user. The database is still left in a consistent state. As an additional benefit, executing within a transaction is sometimes faster, as is the case for PostgreSQL for instance.

To help with this, GNATColl will automatically start a transaction the first time you edit the database. It is then your responsibility to either commit or rollback the transaction when you are done modifying. A lot of database engines (among which PostgreSQL) will not accept any further change to the database if one command in the transaction has failed. To take advantage of this, GNATColl will therefore not even send the command to the server if it is in a failure state.

Here is code sample that modifies the database:

```
Execute (DB, SQL_Insert (...));
-- Executed in the same transaction

Commit_Or_Rollback (DB);
-- Commit if both insertion succeeded, rollback otherwise
-- You can still check Success(DB) afterward if needed
```

24.8 Prepared queries

The previous section showed how to execute queries and statements. But these were in fact relatively inefficient.

With most DBMS servers, it is possible to compile the query once on the server, and then reuse that prepared query to significantly speed up later searches when you reuse that prepared statement.

It is of course pretty rare to run exactly the same query or statement multiple times with the same values. For instance, the following query would not give much benefit if it was prepared, since you are unlikely to reuse it exactly as is later on:

```
SELECT * FROM data WHERE id=1
```

SQL (and GNATColl) provide a way to parameterize queries. Instead of hard-coding the value *1* in the example above, you would in fact use a special character (unfortunately specific to the DBMS you are interfacing to) to indicate that the value will be provided when the query is actually executed. For instance, *sqlite* would use:

```
SELECT * FROM data WHERE id=?
```

You can write such a query in a DBMS-agnostic way by using GNATColl. Assuming you have automatically generated `database.ads` by using *gnatcoll_db2ada*, here is the corresponding Ada code:

```
with Database; use Database;

Q : constant SQL_Query :=
  SQL_Select
    (Fields => Data.Id & Data.Name
     From   => Data,
     Where  => Data.Id = Integer_Param (1));
```

GNATColl provides a number of functions (one per type of field) to indicate that the value is currently unbound. *Integer_Param*, *Text_Param*, *Boolean_Param*,... All take a single argument, which is the index of the corresponding parameter. A query might need several parameters, and each should have a different index. On the other hand, the same parameter could be used in several places in the query.

Although the query above could be executed as is by providing the values for the parameters, it is more efficient, as we mentioned at the beginning, to compile it on the server. In theory, this preparation is done within the context of a database connection (thus cannot be done for a global variable, where we do not have connections yet, and where the query might be executed by any connection later on).

GNATColl will let you indicate that the query should be prepared. This basically sets up some internal data, but does not immediately compile it on the server. The first time the query is executed in a given connection, though, it will first be compiled. The result of this compilation will be reused for that connection from then on. If you are using a second connection, it will do its own compilation of the query.

So in our example we would add the following global variable:

```
P : constant Prepared_Statement :=
  Prepare (Q, On_Server => True);
```

Two comments about this code:

- You do not have to use global variables. You can prepare the statement locally in a subprogram. A *Prepared_Statement* is a reference counted type, that will automatically free the memory on the server when it goes out of scope.
- Here, we prepared the statement on the server. If we had specified *On_Server => False*, we would still have sped things up, since *Q* would be converted to a string that can be sent to the DBMS, and from then on reused

that string (note that this conversion is specific to each DBMS, since they don't always represent things the same way, in particular parameters, as we have seen above). Thus every time you use P you save the time of converting from the GNATColl tree representation of the query to a string for the DBMS.

Now that we have a prepared statement, we can simply execute it. If the statement does not require parameters, the usual *Fetch* and *Execute* subprograms have versions that work exactly the same with prepared statements. They also accept a *Params* parameter that contains the parameter to pass to the server. A number of “+” operators are provided to create those parameters:

```
declare
  F : Forward_Cursor;
begin
  F.Fetch (DB, P, Params => (1 => +2));
  F.Fetch (DB, P, Params => (1 => +3));
end;
```

Note that for string parameters, the “+” operator takes an access to a string. This is for efficiency, to avoid allocating memory and copying the string, and is safe because the parameters are only needed while *Fetch* executes (even for a *Forward_Cursor*).

Back to our library example. We showed earlier how to write a query that retrieves the books borrowed by customer “Smith”. We will now make this query more general: given a customer name, return all the books he has borrowed. Since we expect to use this often, we will prepare it on the server (in real life, this query is of little interest since the customer name is not unique, we would instead use a query that takes the id of the customer). In general we would create a global variable with:

```
Borrowed : constant Prepared_Statement := Prepare
(SQL_Select
  (Fields => Books.Title & Books.Pages,
   Where => Books.FK (Customers)
   and Customers.Last = Text_Param (1));
Auto_Complete => True,
On_Server => True);
```

Then when we need to execute this query, we would do:

```
declare
  Name : aliased String := "Smith";
begin
  R.Fetch (DB, Borrowed, Params => (1 => +Smith'Access));
end;
```

There is one last property on *Prepared_Statement*'s: when you prepare them, you can pass a *Use_Cache* => *True* parameter. When this is used, the result of the query will be cached by GNATColl, and reuse when the query is executed again later. This is the fastest way to get the query, but should be used with care, since it will not detect changes in the database. The local cache is automatically invalidated every hour, so the query will be performed again at most one hour later. Local caching is disabled when you execute a query with parameters. In this case, prepare the query on the server which will still be reasonably fast.

Finally, here are some examples of timings. The exact timing are irrelevant, but it is interesting to look at the different between the various scenarios. Each of them performs 100_000 simple queries similar to the one used in this section:

```
Not preparing the query, using `Direct_Cursor`:
4.05s
```

```
Not preparing the query, using `Forward_Cursor`, and only
```

(continues on next page)

(continued from previous page)

```
retrieving the first row:
```

```
3.69s
```

```
Preparing the query on the client (`On_Server => False`),  
with a `Direct_Cursor`. This saves the whole `GNATCOLL.SQL`  
manipulations and allocations:
```

```
2.50s
```

```
Preparing the query on the server, using `Direct_Cursor`:
```

```
0.55s
```

```
Caching the query locally (`Use_Cache => True`):
```

```
0.13s
```

24.9 Getting results

Once you have executed a *SELECT* query, you generally need to examine the rows that were returned by the database server. This is done in a loop, as in:

```
while Has_Row (R) loop
  Put_Line ("Max priority=" & Integer_Value (R, 0)'Img
    & " for category=" & Value (R, 1));
  Next (R);
end loop;
```

You can only read one row at a time, and as soon as you have moved to the next row, there is no way to access a previously fetched row. This is the greatest common denominator between the various database systems. In particular, it proves efficient, since only one row needs to be kept in memory at any point in time.

For each row, we then call one of the *Value* or **Value* functions which return the value in a specific row and a specific column.

We mentioned earlier there was no way to go back to a row you fetched previously except by executing the query again. This is in fact only true if you use a *Forward_Cursor* to fetch the results.

But GNATColl provides another notion, a *Direct_Cursor*. In this case, it fetches all the rows in memory when the query executes (thus it needs to allocate more memory to save every thing, which can be costly if the query is big). This behavior is supported natively by *PostgreSQL*, but doesn't exist with *sqlite*, so GNATColl will simulate it as efficiently as possible. But it will almost always be faster to use a *Forward_Cursor*.

In exchange for this extra memory overhead, you can now traverse the list of results in both directions, as well as access a specific row directly. It is also possible to know the number of rows that matched (something hard to do with a *Forward_Cursor* since you would need to traverse the list once to count, and then execute the query again if you need the rows themselves).

Direct_Cursor, produced from prepared statements, could be indexed by the specified field value and routine *Find* could set the cursor position to the row with specified field value.:

```
-- Prepared statement should be declared on package level.

Stmt : Prepared_Statement :=
  Prepare ("select Id, Name, Address from Contact order by Name"
    Use_Cache => True, Index_By => Field_Index'First);
```

(continues on next page)

(continued from previous page)

```

procedure Show_Contact (Id : Integer) is
  CI : Direct_Cursor;
begin
  CI.Fetch (DB, Stmt);
  CI.Find (Id); -- Find record by Id

  if CI.Has_Row then
    Put_Line ("Name " & CI.Value (1) & " Address " & CI.Value (2));
  else
    Put_Line ("Contact id not found.");
  end if;
end Show_Contact;

```

In general, the low-level DBMS C API use totally different approaches for the two types of cursors (when they even provide them). By contrast, GNATColl makes it very easy to change from one to the other just by changing the type of a the result variable. So you would in general start with a *Forward_Cursor*, and if you discover you in fact need more advanced behavior you can pay the extra memory cost and use a *Direct_Cursor*.

For both types of cursors, GNATColl automatically manages memory (both on the client and on the DBMS), thus providing major simplification of the code compared to using the low-level APIs.

24.10 Creating your own SQL types

GNATColl comes with a number of predefined types that you can use in your queries. `gnatcoll_db2ada` will generate a file using any of these predefined types, based on what is defined in your actual database.

But sometimes, it is convenient to define your own SQL types to better represent the logic of your application. For instance, you might want to define a type that would be for a *Character* field, rather than use the general *SQL_Field_Text*, just so that you can write statements like:

```

declare
  C : Character := 'A';
  Q : SQL_Query;
begin
  Q := SQL_Select (.., Where => Table.Field = C);
end

```

This is fortunately easily achieved by instantiating one generic package, as such:

```

with GNATCOLL.SQL_Impl; use GNATCOLL.SQL_Impl;

function To_SQL (C : Character) return String is
begin
  return "'" & C & "'";
end To_SQL;

package Character_Fields is new Field_Types (Character, To_SQL);
type SQL_Field_Character is new Character_Fields.Field
  with null record;

```

This automatically makes available both the field type (which you can use in your database description, as `gnatcoll_db2ada` would do, but also all comparison operators like `<`, `>`, `=`, and so on, both to compare with an-

other character field, or with *Character* Ada variable. Likewise, this makes available the assignment operator = so that you can create *INSERT* statements in the database.

Finally, the package *Character_Fields* contain other generic packages which you can instantiate to bind SQL operators and functions that are either predefined in SQL and have no equivalent in GNATColl yet, or that are functions that you have created yourself on your DBMS server.

See the specs of *GNATCOLL.SQL_Impl* for more details. This package is only really useful when writing your own types, since otherwise you just have to use *GNATCOLL.SQL* to write the actual queries.

See also *GNATCOLL.SQL_Fields* for an example on how to have a full integration with other parts of *GNATCOLL.SQL*.

24.11 Query logs

In *Traces: Logging information* we discovered the logging module of GNATColl. The database interface uses this module to log the queries that are sent to the server.

If you activate traces in your application, the user can then activate one of the following trace handles to get more information on the exchange that exists between the database and the application. As we saw before, the output of these traces can be sent to the standard output, a file, the system logs,...

The following handles are provided:

- **SQL.ERROR** This stream is activated by default. Any error returned by the database (connection issues, failed transactions,...) will be logged on this stream
- **SQL** This stream logs all queries that are not **SELECT** queries, ie mostly all queries that actually modify the database
- **SQL.SELECT** This stream logs all select queries. It is separated from **SQL** because very often you will be mostly interested in the queries that impact the database, and logging all selects can generate a lot of output.

In our library example, we would add the following code to see all SQL statements executed on the server:

```
with GNATCOLL.Traces; use GNATCOLL.Traces;
procedure Main is
begin
  GNATCOLL.Traces.Parse_Config_File (".gnatdebug");
  ... -- code as before
  GNATCOLL.Traces.Finalize; -- reclaim memory
```

and then create a *.gnatdebug* in the directory from which we launch our executable. This file would contain a single line containing "+" to activate all log streams, or the following to activate only the subset of fields related to SQL:

```
SQL=yes
SQL.SELECT=yes
SQL.LITE=yes
```

24.12 Writing your own cursors

The cursor interface we just saw is low-level, in that you get access to each of the fields one by one. Often, when you design your own application, it is better to abstract the database interface layer as much as possible. As a result, it is often better to create record or other Ada types to represent the contents of a row.

Fortunately, this can be done very easily based on the API provided by *GNATCOLL.SQL*. Note that *GNATCOLL.SQL.ORM* provides a similar approach based on automatically generated code, so might be even better. But it is still useful to understand the basics of providing your own objects.

Here is a code example that shows how this can be done:

```
type Customer is record
  Id   : Integer;
  First, Last : Unbounded_String;
end record;

type My_Cursor is new Forward_Cursor with null record;
function Element (Self : My_Cursor) return My_Row;
function Do_Query (DB, ...) return My_Cursor;
```

The idea is that you create a function that does the query for you (based on some parameters that are not shown here), and then returns a cursor over the resulting set of rows. For each row, you can use the *Element* function to get an Ada record for easier manipulation.

Let's first see how these types would be used in practice:

```
declare
  C : My_Cursor := Do_Query (DB, ...);
begin
  while Has_Row (C) loop
    Put_Line ("Id = " & Element (C).Id);
    Next (C);
  end loop;
end;
```

So the loop itself is the same as before, except we no longer access each of the individual fields directly. This means that if the query changes to return more fields (or the same fields in a different order for instance), the code in your application does not need to change.

The specific implementation of the subprograms could be similar to the following subprograms (we do not detail the writing of the SQL query itself, which of course is specific to your application):

```
function Do_Query return My_Cursor is
  Q : constant SQL_Query := ....;
  R : My_Cursor;
begin
  R.Fetch (DB, Q);
  return R;
end Do_Query;

function Element (Self : My_Cursor) return My_Row is
begin
  return Customer'
    (Id   => Integer_Value (Self, 0),
```

(continues on next page)

(continued from previous page)

```

    First => To_Unbounded_String (Value (Self, 1)),
    Last  => To_Unbounded_String (Value (Self, 2));
end Element;
```

There is one more complex case though. It might happen that an element needs access to several rows to fill the Ada record. For instance, if we are writing a CRM application and query the contacts and the companies they work for, it is possible that a contact works for several companies. The result of the SQL query would then look like this:

contact_id	company_id
1	100
1	101
2	100

The sample code shown above will not work in this case, since `Element` is not allowed to modify the cursor. In such a case, we need to take a slightly different approach:

```

type My_Cursor is new Forward_Cursor with null record;
function Do_Query return My_Cursor; -- as before
procedure Element_And_Next
  (Self : in out My_Cursor; Value : out My_Row);
```

where `Element_And_Next` will fill `Value` and call `Next` as many times as needed. On exit, the cursor is left on the next row to be processed. The usage then becomes:

```

while Has_Row (R) loop
  Element_And_Next (R, Value);
end loop;
```

To prevent the user from using `Next` incorrectly, you should probably override `Next` with a procedure that does nothing (or raises a `Program_Error` maybe). Make sure that in `Element_And_Next` you are calling the inherited function, not the one you have overridden, though.

There is still one more catch. The user might depend on the two subprograms `Rows_Count` and `Processed_Rows` to find out how many rows there were in the query. In practice, he will likely be interested in the number of distinct contacts in the tables (2 in our example) rather than the number of rows in the result (3 in the example). You thus need to also override those two subprograms to return correct values.

24.13 The Object-Relational Mapping layer (ORM)

GNATColl provides a high-level interface to manipulate persistent objects stored in a database, using a common paradigm called an object-relational mapping. Such mappings exist for most programming languages. In the design of GNATColl, we were especially inspired by the python interface in *django* and *sqlalchemy*, although the last two rely on dynamic run time introspection and GNATColl relies on code generation instead.

This API is still compatible with *GNATCOLL.SQL*. In fact, we'll show below cases where the two are mixed. It can also be mixed with *GNATCOLL.SQL.Exec*, although this might be more risky. Communication with the DBMS is mostly transparent in the ORM, and it uses various caches to optimize things and make sure that if you modify an element the next query(ies) will also return it. If you use *GNATCOLL.SQL.Exec* directly you are bypassing this cache so you risk getting inconsistent results in some cases.

In ORM, a table is not manipulated directly. Instead, you manipulate objects that are read or written to a table. When we defined our database schema (*Database schema*), we gave two names on the first line of a table definition. There was the name of the table in the database, and the name of the object that each row represent. So for our library example

we have defined *Customer*, *Book* and *Dvd* objects. These objects are declared in a package generated automatically by *gnatcoll_db2ada*.

There is first one minor change we need to do to our library example. The ORM currently does not handle properly cases where an abstract class has foreign keys to other tables. So we remove the *borrowed_by* field from the *Media* table, and change the *books* table to be:

TABLE (media)	books	book	The books in the library
pages	INTEGER	100	
borrowed_by	FK customers(borrowed_books)	NULL	Who borrowed the media

Let's thus start by generating this code. We can replace the command we ran earlier (with the *-api* switch) with one that will also generate the ORM API:

```
gnatcoll_db2ada -dbmode dbschema.txt -api Database -orm ORM
```

The ORM provides a pool of database connections through the package *GNATCOLL.SQL.Sessions*. A session therefore acts as a wrapper around a connection, and provides a lot more advanced features that will be described later. The first thing to do in the code is to configure the session pool. The *Setup* procedure takes a lot of parameters to make sessions highly configurable. Some of these parameters will be described and used in this documentation, others are for special usage and are only documented in *gnatcoll-sql-sessions.ads*. Here will we use only specify the mandatory parameters and leave the default value for the other parameters:

```
GNATCOLL.SQL.Sessions.Setup
(Descr => GNATCOLL.SQL.SQLite.Setup ("library.db"),
 Max_Sessions => 2);
```

The first parameter is the same *Database_Description* we saw earlier (*Connecting to the database*), but it will be freed automatically by the sessions package, so you should not free it yourself.

Once configure, we can now request a session. Through a session, we can perform queries on the database, make objects persistent, write the changes back to the database,... We configured the session pool to have at most 2 sessions. The first time we call *Get_New_Session*, a new session will be created in the pool and marked as busy. While you have a reference to it in your code (generally as a local variable), the session belongs to this part of the code. When the session is no longer in scope, it is automatically released to the pool to be reused for the next call to *Get_New_Session*. If you call *Get_New_Session* a second time while some part of your code holds a session (for instance in a different task), a new session will be created. But if you do that a third time while the other two are busy, the call to *Get_New_Session* is blocking until one of the two sessions is released to the pool.

This technics ensures optimal use of the resources: we avoid creating a new session every time (with the performance cost of connecting to the database), but also avoid creating an unlimited number of sessions which could saturate the server. Since the sessions are created lazily the first time they are needed, you can also configure the package with a large number of sessions with a limited cost.

Let's then take a new session in our code:

```
Session : constant Session_Type := Get_New_Session;
```

and let's immediately write our first simple query. A customer comes at the library, handles his card and we see his id (1). We need to look up in the database to find out who he is. Fortunately, there is no SQL to write for this:

```
C : ORM.Detached_Customer'Class := Get_Customer (Session, Id => 1);
```

The call to *Get_Customer* performs a SQL query transparently, using prepared statements for maximum efficiency. This results in a *Customer* object.

ORM is the package that was generated automatically by *gnatcoll_db2ada*. For each table in the database, it generates a number of types:

- *Customer*

This type represents a row of the *Customers* table. It comes with a number of primitive operations, in particular one for each of the fields in the table. Such an object is returned by a cursor, similarly to what was described in the previous section (*Writing your own cursors*). This object is no longer valid as soon as the cursor moves to the next row (in the currently implementation, the object will describe the next row, but it is best not to rely on this). As a benefit, this object is light weight and does not make a copy of the value of the fields, only reference the memory that is already allocated for the cursor.

This object redefines the equality operator (“=”) to compare the primary key fields to get expected results.

- *Detached_Customer*

A detached object is very similar to the *Customer* object, but it will remain valid even if the cursor moves or is destroyed. In fact, the object has made a copy of the value for all of its fields. This object is heavier than a *Customer*, but sometimes easier to manager. If you want to store an object in a data structure, you must always store a detached object.

A detached object also embeds a cache for its foreign keys. In the context of our demo for instance, a *Book* object was borrowed by a customer. When returning from a query, the book knows the id of that customer. But if call *B.Borrowed_By* this returns a *Detached_Customer* object which is cached (the first time, a query is made to the DBMS to find the customer given his id, but the second time this value is already cached).

One cache create a *Detached_Customer* from a *Customer* by calling the *Detach* primitive operation.

- *Customer_List*

This type extends a *Forward_Cursor* (*Getting results*). In addition to the usual *Has_Row* and *Next* operations, it also provides an *Element* operation that returns a *Customer* for easy manipulation of the results.

- *Direct_Customer_List*

This type extends a *Direct_Cursor*. It also adds a *Element* operation that returns a *Customer* element.

- *Customers_Managers*

This type is the base type to perform queries on the DBMS. A manager provides a number of primitive operations which end up creating a SQL query operation in the background, without making that explicit.

Let’s first write a query that returns all books in the database:

```
declare
  M : Books_Managers := All_Books;
  BL : Book_List := M.Get (Session);
  B : Book;
begin
  while BL.Has_Row loop
    B := BL.Element;
    Put_Line ("Book: " & B.Title);
    Put_Line ("  Borrowed by: " & B.Borrowed_By.Last);
    BL.Next;
  end loop;
end;
```

The manager *M* corresponds to a query that returns all the books in the database. The second line then executes the query on the database, and returns a list of books. We then traverse the list. Note how we access the book’s title by calling a function, rather than by the index of a field as we did with *GNATCOLL.SQL.Exec* with *Value(B, 0)*. The code is much less fragile this way.

The line that calls *Borrowed_By* will execute an additional SQL query for each book. This might be inefficient if there is a large number of books. We will show later how this can be optimized.

The manager however has a lot more primitive operations that can be used to alter the result. Each of these primitive operations returns a modified copy of the manager, so that you can easily chain calls to those primitive operations. Those operations are all declared in the package *GNATCOLL.SQL.ORM.Impl* if you want to look at the documentation. Here are those operations:

- *Get* and *Get_Direct*

As seen in the example above, these are the two functions that execute the query on the database, and returns a list of objects (respectively a *Customer_List* and a *Direct_Customer_List*).

- *Distinct*

Returns a copy of the manager that does not return twice a row with the same data (in SQL, this is the “DISTINCT” operator)

- *Limit* (Count : Natural; From : Natural := 0)

Returns a copy of the manager that returns a subset of the results, for instance the first *Count* ones.

- *Order_By* (By : SQL_Field_List)

Returns a copy of the manager that sorts the results according to a criteria. The criteria is a list of field as was defined in *GNATCOLL.SQL*. We can for instance returns the list of books sorted by title, and only the first 5 books, by replacing *M* with the following:

```
M : Books_Managers := All_Books.Limit (5).Order_By (Books.Title);
```

- *Filter*

Returns a subset of the result matching a criteria. There are currently two versions of Filter: one is specialized for the table, and has one parameter for each field in the table. We can for instance return all the books by Alexandre Dumas by using:

```
M : Books_Managers := All_Books.Filter (Author => "Dumas");
```

This version only provides the equality operator for the fields of the table itself. If for instance we wanted all books with less than 50 pages, we would use the second version of filter. This version takes a *GNATCOLL.SQL.SQL_Criteria* similar to what was explained in previous sections, and we would write:

```
M : Books_Managers := All_Books.Filter (Condition => Books.Pages < 50);
```

More complex conditions are possible, involving other tables. Currently, the ORM does not have a very user-friendly interface for those, but you can always do this by falling back partially to SQL. For instance, if we want to retrieve all the books borrowed by user “Smith”, we need to involve the *Customers* table, and thus make a join with the *Books* table. In the future, we intend to make this join automatic, but for now you will need to write:

```
M : Books_Managers := All_Books.Filter
(Books.FK (Customers)
 and Customers.Last = "Smith");

-- SQL query: SELECT books.pages, books.borrowed_by, books.id,
--             books.title, books.author, books.published
--             FROM books, customers
--             WHERE books.borrowed_by=customers.id AND customers.last='Smith'
```

This is still simpler code than we were writing with *GNATCOLL.SQL* because we do not have to specify the fields or tables, and the results are objects rather than fields with specific indexes.

- *Select_Related* (Depth : Integer; Follow_Left_Join : Boolean)

This function returns a new manager that will retrieve all related objects. In the example we gave above, we mentioned that every time *B.Borrowed_By* was called, this resulted in a call to the DBMS. We can optimize this by making sure the manager will retrieve that information. As a result, there will be a single query rather than lots. Be careful however, since the query will return more data, so it might sometimes be more efficient to perform multiple smaller queries.

Depth indicates on how many levels the objects should be retrieved. For instance, assume we change the schema such that a Book references a Customer which references an Address. If we pass 1 for *Depth*, the data for the book and the customer will be retrieved. If however you then call *B.Borrowed_By.Address* this will result in a query. So if you pass 2 for *Depth* the data for book, customers and addresses will be retrieved.

The second parameter related to efficiency. When a foreign key was mentioned as *NOT NULL* in the schema, we know it is always pointing to an existing object in another table. *Select_Related* will always retrieve such objects. If, however, the foreign key can be null, i.e. there isn't necessarily a corresponding object in the other table, the SQL query needs to use a *LEFT JOIN*, which is less efficient. By default, GNATColl will not retrieve such fields unless *Follow_Left_Join* was set to True.

In our example, a book is not necessarily borrowed by a customer, so we need to follow the left joins:

```
M : Books_Managers := All_Books.Filter
  (Books.FK (Customers)
   and Customers.Last = "Smith")
.Select_Related (1, Follow_Left_Join => True);

-- SQL query: SELECT books.pages, books.borrowed_by, books.id,
--             books.title, books.author, books.published,
--             customers.id, customers.first, customers.last
--             FROM (books LEFT JOIN customers ON books.borrowed_by=customers.id)
--             WHERE books.borrowed_by=customers.id AND customers.last='Smith'
```

24.13.1 reverse relationships

In fact, the query we wrote above could be written differently. Remember we have already queried the *Customer* object for id 1 through a call to *Get_Customer*. Since our schema specified a *reverse_name* for the foreign key *borrowed_by* in the table *books*, we can in fact simply use:

```
BL := C.Borrowed_Books.Get (Session);

-- SQL: SELECT books.pages, books.borrowed_by, books.id, books.title,
--       books.author, books.published FROM books
--       WHERE books.borrowed_by=1
```

Borrowed_Books is a function that was generated because there was a *reverse_name*. It returns a *Books_Managers*, so we could in fact further filter the list of borrowed books with the same primitive operations we just saw. As you can see, the resulting SQL is optimal.

Let's optimize further the initial query. We have hard-coded the customer name, but in fact we could be using the same subprograms we were using for prepared statements (*Prepared queries*), and even prepare the query on the server for maximum efficiency. Since our application is likely to use this query a lot, let's create a global variable:

```
M : constant Books_Managers := All_Books.Filter
  (Books.FK (Customers)
```

(continues on next page)

(continued from previous page)

```

    and Customers.Id = Integer_Param (1))
    .Select_Related (1, Follow_Left_Join => True);

MP : constant ORM_Prepared_Statement :=
    M.Prepare (On_Server => True);

... later in the code

Smith_Id : constant Natural := 1;
BL : Book_List := MP.Get (Session, Params => (1 => Smith_Id));

```

The last call to *Get* is very efficient, with timing improvements similar to the ones we discussed on the session about prepared statements (*Prepared queries*).

24.14 Modifying objects in the ORM

The ORM is much more than writing queries. Once the objects are persistent, they can also be simply modified, and they will be saved in the database transparently.

Let's start with a simple example. In the previous section, we retrieve an object *C* representing a customer. Let's change his name, and make sure the change is in the database:

```

C := Get_Customer (Session, 1);
C.Set_Last ("Smit");
C.Set_First ("Andrew");
Session.Commit;

```

A reasonable way to modify the database. However, this opens a can of complex issues that need to be dealt with.

When we called *Set_Last*, this modifies the objects in memory. At this point, printing the value of *C.Last* would indeed print the new value as expected. The object was also marked as modified. But no change was made in the database.

Such a change in the database might in fact be rejected, depending on whether there are constraints on the field. For instance, say there existed a constraint that *Last* must be the same *First* (bear with me, this is just an example). If we call *Set_Last*, the constraint is not satisfied until we also call *Set_First*. But if the former resulted in an immediate change in the database, it would be rejected and we would not even get a change to call *Set_First*.

Instead, the session keeps a pointer to all the objects that have been modified. When it is committed, it traverses this list of objects, and commits their changes into the database. In the example we gave above, the call to *Commit* will thus commit the changes to *C* in the database. For efficiency, it uses a single SQL statement for that, which also ensures the constraint remains valid:

```

UPDATE customers SET first='Andrew', last='Smit' WHERE customers.id=1;

```

We can create a new customer by using similar code:

```

C := New_Customer;
C.Set_First ("John");
C.Set_Last ("Lee");
Session.Persist (C);

Session.Commit;

```

New_Customer allocates a new object in memory. However, this object is not persistent. You can call all the *Set_** subprograms, but the object will not be saved in the database until you add it explicitly to a session with a call to *Persist*, and then *Commit* the session as usual.

Another issue can occur when objects can be modified in memory. Imagine we retrieve a customer, modify it in memory but do not commit to the database yet because there are other changes we want to do in the same SQL transaction. We then retrieve the list of all customers. Of course, the customer we just modified is part of this list, but the DBMS does not know about the change which currently only exists in memory.

Thankfully, GNATColl takes care of this issue automatically: as we mentioned before, all modified objects are stored in the session. When traversing the list of results, the cursors will check whether the session already contains an element with the same id that it sees in the result, and if yes will return the existing (i.e. modified) element. For instance:

```
C := Get_Customer (Session, Id => 1);
C.Set_Last ("Lee");

CL : Customer_List := All_Customers.Get (Session);
while CL.Has_Row loop
  Put_Line (CL.Element.Last);
  CL.Next;
end loop;
```

The above example uses *CL.Element*, which is a light-weight *Customer* object. Such objects will only see the in-memory changes if you have set *Flush_Before_Query* to true when you configured the sessions in the call to *GNATCOLL.SQL.Sessions.Setup*. Otherwise, it will always return what's really in the database.

If the example was using *Detached_Customer* object (by calling *CL.Element.Detach* for instance) then GNATColl looks up in its internal cache and returns the cached element when possible. This is a subtlety, but this is because an *Customer* only exists as long as its cursor, and therefore cannot be cached in the session. In practice, the *Flush_Before_Query* should almost always be true and there will be not surprising results.

24.15 Object factories in ORM

Often, a database table is used to contain objects that are semantically of a different kind. In this section, we will take a slightly different example from the library. We no longer store the books and the DVDs in separate tables. Instead, we have one single *media* table which contains the title and the author, as well as a new field *kind* which is either 0 for a book or 1 for a DVD.

Let's now look at all the media borrowed by a customer:

```
C : constant Customer'Class := Get_Customer (Session, Id => 1);
ML : Media_List := C.Borrowed_Media.Get (Session);

while ML.Has_Row loop
  case ML.Element.Kind is
    when 0 =>
      Put_Line ("A book " & ML.Element.Title);
    when 1 =>
      Put_Line ("A DVD " & ML.Element.Title);
  end case;
  ML.Next;
end loop;
```

This code works, but requires a case statement. Now, let's imagine the check out procedure is different for a book and a DVD (for the latter we need to check that the disk is indeed in the box). We would have two subprograms

Checkout_Book and *Checkout_DVD* and call them from the case. This isn't object-oriented programming.

Instead, we will declare two new types:

```
type My_Media is abstract new ORM.Detached_Media with private;
procedure Checkout (Self : My_Media) is abstract;

type Detached_Book is new My_Media with private;
overriding Checkout (Self : Detached_Book);

type Detached_DVD is new My_Media with private;
overriding Checkout (Self : Detached_DVD);
```

We could manually declare a new *Media_List* and override *Element* so that it returns either of the two types instead of a *Media*. But then we would also need to override *Get* so that it returns our new list. This is tedious.

We will instead use an element factory in the session. This is a function that gets a row of a table (in the form of a *Customer*), and returns the appropriate type to use when the element is detached (by default, the detached type corresponding to a *Customer* is a *Detached_Customer*, and that's what we want to change).

So let's create such a factory:

```
function Media_Factory
  (From      : Base_Element'Class;
   Default   : Detached_Element'Class) return Detached_Element'Class
is
begin
  if From in Media'Class then
    case Media (From).Kind is
      when 0 =>
        return R : Detached_Book do null; end return;
      when 1 =>
        return R : Detached_DVD do null; end return;
      when others =>
        return Default;
    end case;
  end if;
  return Default;
end Media_Factory;

Session.Set_Factory (Media_Factory'Access);
```

This function is a bit tricky. It is associated with a given session (although we can also register a default factory that will be associated with all sessions by default). For all queries done through this session (and for all tables) it will be called. So we must first check whether we are dealing with a row from the *Media* table. If not, we simply return the suggested *Default* value (which has the right *Detached_** kind corresponding to the type of *From*).

If we have a row from the *Media* table, we then retrieve its kind (through the usual automatically generated function) to return an instance of *Detached_Book* or *Detached_DVD*. We use the Ada05 notation for extended return statements, but we could also use a declare block with a local variable and return that variable. The returned value does not need to be further initialized (the session will take care of the rest of the initialization).

We can now write our code as such:

```
C : constant Customer'Class := Get_Customer (Session, Id => 1);
ML : Media_List := C.Borrowed_Media.Get (Session);
```

(continues on next page)

(continued from previous page)

```
while ML.Has_Row loop
  Checkout (ML.Element.Detach);  -- Dispatching
  ML.Next;
end loop;
```

The loop is cleaner. Of course, we still have the case statement, but it now only exists in the factory, no matter how many loops we have or how many primitive operations of the media we want to define.

XREF: CROSS-REFERENCING SOURCE CODE

When manipulating source code, programmers need to know where the various symbols are defined, where they are used, and so on. This is generally available directly from their IDE and editors. But computing this information in the first place is tricky, especially for languages that support overloading of subprograms.

Some compilers like GNAT and gcc can generate this information for Ada code bases. For instance, GNAT will generate `.ali` files, which contain the navigation information, when compiling Ada or SPARK code.

GNATCOLL.Xref can then be used to parse and aggregate all those files into a single sqlite database, which can be conveniently used to answer queries such as “give me the declaration for this entity”, “list all places where this entity is used”, “show all subprograms that could be called in practice at this dispatching call”, “what files does this file depend on”, “show me the call graph for this application”,...

To use this package, some initialization needs to be performed first:

```
with GNATCOLL.Xref;      use GNATCOLL.Xref;
with GNATCOLL.SQL.Sqlite;
with GNATCOLL.Projects; use GNATCOLL.Projects;  -- 1
with GNATCOLL.VFS;      use GNATCOLL.VFS;
with GNAT.Strings;

procedure Support is
  DB : Xref_Database;
  Tree : Project_Tree_Access := new Project_Tree;
  Error : GNAT.Strings.String_Access;

begin
  Tree.Load (Create ("prj.gpr"));  -- 2
  Setup_DB
    (DB, Tree,
     GNATCOLL.SQL.Sqlite.Setup (Database => "testdb.db"),
     Error);  -- 3
  Free (Error);
  Parse_All_LI_Files (DB, Tree.Root_Project);  -- 4
end Support;
```

GNATCOLL needs to be able to find the `*li` files. For this, it depends on project files (as supported by GNATCOLL.Projects). So the first thing to do is to parse the project (step 2).

We then need to tell GNATCOLL where the cross-reference information needs to be aggregated. In this example, it will be stored in a sqlite database on the disk. By using a name “:memory:” instead, we would create a temporary in-memory database. This is in general faster, but uses more memory and needs to be recreated every time the program is restarted. We could also decide to store the information in any other database supported by GNATCOLL.SQL.Exec, for instance PostgreSQL.

Finally, in step 4 we let GNATCOLL parse all the *.i files that are relevant for this project. This operation can take a while, depending on the size of the project. However, if the database already exists on the disk, it will simply be updated by parsing the files that are not already up-to-date. When all files are up-to-date, this operation is almost immediate.

At this point, we now have a database that we can start querying. Here are a few examples, but see the documentation `gnatcoll-xref.ads` for more types of queries. All these queries have a similar API: they return a **cursor** which iterates over the result returned by a SQL query. There are various kinds of cursors, depending on whether they return files, entities, or references to entities. But they all support the *Has_Element*, *Element* and *Next* operations, so all loops will look similar:

```
pragma Ada_05;    -- use object-dotted-notation
with GNATCOLL.VFS; use GNATCOLL.VFS;

declare
  Entity : Entity_Information;
  Ref     : Entity_Reference;
  File    : Virtual_File;
  Refs    : References_Cursor;
begin
  File := Tree.Create ("source.ads");    -- 5
  Ref := DB.Get_Entity ("Method", File, Line => 2);    -- 6
  Entity := Ref.Entity;

  DB.References (Entity, Refs);    -- 7
  while Refs.Has_Element loop
    Ref := Refs.Element;
    Put_Line (" at " & Ref.File.Display_Full_Name & ':'
              & Ref.Line'Img & ':' & Ref.Column'Img);
    Refs.Next;
  end loop;
end;
```

This example will print all the references to the entity that is referenced in file `source.ads` at line 2 (the column is unspecified).

Step 5 gets a handle on the source file. Here, we depend on the project to find the precise directory in which the source file is located. We can of course use an absolute file name instead.

Step 6 gets handle on the entity referenced on line 2 in this file. Such an entity is the starting point for most queries defined in *GNATCOLL.Xref*.

Finally, on step 7 and the loop below we iterate over all references, and print their location on the standard output.

Let's do a much more complex query: we want to see all references to that entity, but also places where the entity might be called through a *renames* statement, or called through a dispatching call via an overriding method defined on a child tagged type (assuming this is a primitive operation of a tagged type in the first place). We also want to see all locations where a method that overrides "Method" is called:

```
declare
  Refs : Recursive_References_Cursor;
begin
  DB.Recursive (Entity, GNATCOLL.Xref.References'Access,
                From_Overriding => True, From_Overridden => True,
                From_Renames => True);
  while Refs.Has_Element loop
    ... same as before
```

(continues on next page)

(continued from previous page)

```
    Refs.Next;  
  end loop;  
end;
```

As shown above, the programing pattern is always the same.

GNATCOLL.Xref provides many more subprogram to get information like the list of fields for a record type (or a C structure), the list of primitive operations or methods for a tagged object or a class, the call graph for a subprogram,...

It is also able to extract documentation for an entity from the source code, by looking at the lines of code just before or just after the declaration or the body of the entity.

XREF: GNATINSPECT

As discussed in the previous section, GNATCOLL provides an Ada API to perform cross-references queries.

There exist a few alternatives when you want to reuse that cross-reference information from other tools, or command line scripts.

You can of course access the sqlite database directly. Most programming languages have an interface to sqlite. For instance python does.

But GNATCOLL provides a command line tool dedicated to that purpose, named **gnatinspect**.

When it is first started on a project, this tool will refresh the xref database by parsing all the ALI files from the project. This might take a while (up to several minutes) the first time, unless of course the xref were already up-to-date because you had loaded the project in GPS first, or already run gnatinspect.

gnatinspect then displays an interactive prompt that lets you perform various queries on the database. The full list of queries is available by typing “help” at the prompt, but this documentation will demonstrate some of them.

Let’s first look at a number of command line switches that might be useful:

- **-db=ARG**: this switch can be used to specify the name of the database.

By default, gnatinspect checks in the project whether there exists an attribute `IDE'Xref_Database`, which should specify a file name (relative to the project’s `object_dir`) for the database.

If this attribute does not exist, it defaults to “gnatinspect.db” in the project’s object directory.

If there is no object directory defined in the project, the file is created in the project’s directory itself. You can however specify any name, including an absolute path, or a path relative to the project’s object directory.

An alternative is to specify `:memory:`, which creates the database in memory. This is of course a temporary database which will disappear when gnatinspect exits, and cannot be shared with other tools.

- **-nightlydb=ARG**: this switch can help speed up the initial startup of gnatinspect. The idea is that in a lot of cases, the software on which a team works is build nightly in a common setup. Running gnatinspect in that setup will create or update an xref database. Individual developers can then create their own copy of the database by starting from the contents of the nightly database (which is pointed to by the `-nightlydb` switch), and then gnatinspect will parse the ALI files in the user’s setup that are different from the nightly ones.
- **-runtime**: by default, gnatinspect will only parse the ALI files from your project (and of course the ones from imported projects). It will not however parse the ALI files found in predefined directories, like for instance the GNAT runtime. This saves time in general. If you click on a call to one of the runtime subprograms in your own code, gnatinspect will be able to point you to its declaration. However, you will not have access to the body, because the link from declaration to body is found in the ALI files of the runtime.
- **-command=ARG**: gnatinspect will update the xref database as usual, then execute a command, display its result, and exit. This can be convenient when calling gnatinspect from another tool, like Emacs or vi.

- `-file=ARG`: similar to `-command`, but reads the commands to execute from a file. The file can contain comments (starting with `'-'`). See also the `-lead` switch.
- `-lang=LANG:SPEC:BODY:OBJ`: specifies a naming scheme for a language. The preferred approach is to use a configuration project file (such as those generated by `gprconfig` for instance), that would define attributes such as *Naming'Spec_Suffix*, *Naming'Body_Suffix* and *Compiler'Object_File_Suffix*. However, this switch provides an alternative whereby you can specify the same values directly on the command line. For instance, the equivalent of:

```
configuration project Autoconf is
  package Naming is
    for Spec_Suffix ("MyLang") use ".myl";
  end Naming;
  package Compiler is
    for Object_File_Suffix ("MyLang") use ".ali";
  end Compiler;
end Autoconf;
```

is to use:

```
--lang=MyLang:.myl:.ali
```

A third alternative is to have the same contents as the configuration project file above, directly in your own project file. This has the same effect, but needs to be duplicated in each of your project file.

Given one of the above, and assuming your project file includes:

```
for Languages use ("Ada", "MyLang");
```

then any file with the `.myl` extension will be correctly detected by the project manager, and any `.ali` file with the same base name will be parsed by `gnatinspect` to find cross-reference information. Remember that the switch `-config=autoconf.cgpr` must be passed to `gnatinspect` if the information is provided via a config project file.

- `-lead=ARG` should be used in coordination with `-file`, and specify lines to ignore from the file. All lines starting with the given prefix will be ignored.
- `-basenames`: controls the display of file names in the output. By default, `gnatinspect` outputs full path information.
- `-exit`: if this switch is specified, `gnatinspect` updates the xref database and exits immediately.
- `-project=ARG` or `-P ARG` specifies the name of the project to load. This switch is mandatory.
- `-X VAR=VALUE` is used to specify the value of scenario variables used in your project. This is similar to the homonym switch in `gprbuild`.
- `-symlinks` should be specified if your project uses symbolic links for files. This will ensure that the links are fully resolved as stored in the database, and thus that when a file is visible through different links, the information is appropriately coalesced in the database for that file.
- `-subdirs=ARG` is similar to the homonym switch in `gprbuild`
- `-tracefile=ARG` is used to point to a file compatible with GNATCOLL.Traces that controls the debug information generated by `gnatinspect`. By default, `gnatinspect` parses a file called `'gnatdebug'` in the current directory.
- `-encoding=ARG` is the character encoding used for source and ALI files. By default, `gnatinspect` assumes they are encoded in UTF-8.

Once it has finished parsing the xref information, `gnatinspect` displays an interactive prompt, where a number of commands can be used to perform queries. In a lot of cases, these commands take some file information as argument (either just the file, or an entity name and the file in which it is defined).

The file names can be given as either a base name, or relative to the current directory, or even a full name. But file names are ambiguous (even when a full path is specified) when aggregate projects are used. It is valid for a given file to be part of multiple aggregate projects, and depending on the project we are considering the result of the xref queries might vary).

To remove the ambiguity, it is possible to specify the project to which the file belongs. The project is specified either as a project name (which itself could be ambiguous with aggregate projects), or as a full path.

In all commands below, whenever the parameter specifies “:file”, you can use instead “:file:project” if there are ambiguities. It is also possible not to specify the file, in which case the entity will be looked for in all sources of the project.

Here is the full list of commands supported by gnatinspect:

- *decl name:file:line:column* is probably the most useful command. Given a reference to an entity, it will indicate where the entity is declared. The line and column informations are optional:

```
>>> decl Func:file.adb:12
Func:/some/path/file2.adb:20:9
```

- *body name:file:line:column* is similar to *decl*, but will return the location of the body of the entity. When the entity is an Ada private type, its body is in fact the location of the full declaration for that type.
- *refs name:file:line:column* displays all known references to the entity.
- *refs_overriding name:file:line:column* displays all known references to the entity or one of its overriding entities
- *doc name:file:line:column* will display documentation for the entity. The exact format for the entity might change in future versions of gnatinspect, but will in general include the type of entity, the location of its declaration, and any comment associated with it in the source code:

```
>>> doc Func:file.adb
procedure declared at /some/path/file2.adb:20:9

And the comments written below Func in file2.adb
```

- *fields name:file:line:column* displays the fields of an Ada record type or a C struct:

```
>>> fields Rec:file.ads:20
A:/some/path/file.ads:21
B:/some/path/file.ads:22
```

- *child_types name:file:line:column* lists all child types for this entity, for instance classes that inherit from the entity. This is the opposite of *parent_types*.
- *child_types_recursive name:file:line:column* is similar to *child_types* but will also list the child types of the children. This query can be used to find a whole tagged type hierarchy (or class hierarchy in C++).
- *parent_types name:file:line:column* returns the parent types for the entity, for instance the classes or interfaces from which it derives. See also *child_types*.
- *methods name:file:line:column* returns the list of methods (or primitive operations) for the entity.
- *method_of name:file:line:column* returns the class or tagged type for which the entity is a method.
- *calls name:file:line:column* lists all entities called by the entity. This includes all entities defined within the scope of the entity (so for a subprogram this will be the list of local variables, but for a package this includes all subprograms and nested packages defined within that package).
- *callers name:file:line:column* lists all entities that call the entity. This information is also available from a call to ‘refs’, but ‘callers’ return the callers directly, instead of references to the original entity.

- *overrides name:file:line:column* returns the entity that is overridden by the entity (generally a method from a parent class).
- *overridden name:file:line:column* returns the list of entities that override the parameter (generally methods from children classes).
- *overridden_recursive name:file:line:column* returns the list of entities that override the parameter (generally methods from children classes). This is recursive.
- *type name:file:line:column* returns the type of the entity (variable or constant). For an enumeration literal, this returns the corresponding enumeration.
- *component name:file:line:column* returns the component type of the entity (for arrays for instance).
- *literals name:file:line:column* returns the valid literal values for an enumeration.
- *pointed name:file:line:column* returns the type pointed to by the entity.
- *qname name:file:line:column* returns the fully qualified name for the entity.
- *params name:file:line:column* returns the list of parameters for the subprogram.

A number of queries are related to the source files of the project:

- *importing filename* lists the files that import the file (via with statements in Ada or #include in C for instance)
- *imports filename* lists the files that the file imports (via with statements in Ada or #include in C for instance). See also *depends_on*.
- *depends filename* lists the files that the file depends on (recursively calling *imports*)
- *entities file* lists all entities referenced or declared in the file.

Finally, some commands are not related to entities or source files:

- *refresh* refreshes the contents of the xref database, by parsing all ALI files that have been changed.
- *shell* Execute a shell command (an alternative is to use ‘!’ as the command).
- *scenario VARIABLE VALUE* changes the value of a scenario variable, and reparses the project.
- *time command arguments* executes the command as usual, and report the time it took to execute it.

TERMINAL: CONTROLLING THE CONSOLE

Applications generally provide user feedback either via full-fledge graphical interfaces, or via a simpler, console-based output.

The basic support for console-based output is provided directly via *Ada.Text_IO*. But more advanced features are highly system-dependent, and somewhat tricky to develop.

The package *GNATCOLL.Terminal* provide cross-platform support for manipulating colors in terminals, as well as a few basic cursor manipulation subprograms.

27.1 Colors

Most modern terminals support color output, generally with a limit set of colors. On Unix systems, these colors are set by using escape sequences in the output; on Windows systems, these are manipulated by calling functions on a file handle.

GNATCOLL will automatically try to guess whether its output is sent to a color enabled terminal. In general, this will be true when outputting to standard output or standard error, and false when outputting to files or to pipes. You can override this default value to force either color support or black-and-white support.

Here is an example:

```
with Ada.Text_IO;           use Ada.Text_IO;
with GNATCOLL.Terminal;     use GNATCOLL.Terminal;

procedure Test_Colors is
  Info : Terminal_Info;
begin
  Info.Init_For_Stdout (Auto);

  Info.Set_Color (Standard_Output, Blue, Yellow);
  Put_Line ("A blue on yellow line");

  Info.Set_Color (Standard_Output, Style => Reset_All);
  Put_Line ("Back to standard colors -- much better");
end Test_Colors;
```

27.2 Cursors

It is often useful for an application to display some progress indicator during long operations. *GNATCOLL.Terminal* provides a limit set of subprograms to do so, as in:

```
with Ada.Text_IO;           use Ada.Text_IO;
with GNATCOLL.Terminal;     use GNATCOLL.Terminal;

procedure Test_Colors is
  Info : Terminal_Info;
begin
  Info.Init_For_Stdout (Auto);
  for J in 1 .. 1_000 loop
    if J mod 10 = 0 then
      Put ("Processing file" & J'Img & " with long name");
    else
      Put ("Processing file" & J'Img);
    end if;
    delay 0.1;
    Info.Beginning_Of_Line;
    Info.Clear_To_End_Of_Line;
  end loop;
end Test_Colors;
```

PROMISES: DEFERRING WORK

This package provides a way to synchronize work between some asynchronous workers (or threads).

Promises are a way to encapsulate a yet unknown value, immediately return to the caller, and work in the background to actually execute the work.

For instance, you could have a function that reads some data from a socket. This takes time, and we do not want to block the application while retrieving the data (and if there is an error retrieving it, we certainly want to properly handle it).

The main thread (for instance a graphical user interface) needs to keep processing events and refresh itself. As soon as the data becomes available from the socket, we should let this main thread know so that it can take further action, like post-processing the data and then displaying it.

A general scheme to do that is to have a callback function that is called whenever the work is finished. Promises build on that simple idea so that you can easily chain multiple callbacks to build more complex actions.

See the extensive documentation in `gnatcoll-promises.ads`

INDICES AND TABLES

- `genindex`

This document may be copied, in whole or in part, in any form or by any means, as is or with alterations, provided that (1) alterations are clearly marked as alterations and (2) this copyright notice is included unmodified in any copy.

Symbols

.gnatdebug, 25

A

ADA_DEBUG_FILE, 25

B

Boyer-Moore, 47

C

charset, 83
 Class Callback_Data, 17
 Class Class_Instance, 17
 Class Console, 12
 class diagram, 16
 Class Instance_Property, 18
 Class Scripting_Language, 17
 Class Scripts_Repository, 17
 Class Subprogram_Record, 18
 Console.clear, 12
 Console.flush, 12
 Console.isatty, 12
 Console.read, 12
 Console.readline, 12
 Console.write, 12

D

decorator, 29

E

email, 53
 encoding, 53
 Exception Exception, 11
 Exception Invalid_Argument, 11
 Exception Missing_Arguments, 11
 Exception Unexpected_Exception, 11

F

filling, 49
 Flush_Before_Query, 109
 Function clear_cache, 11

Function echo, 11
 Function exec_in_console, 11
 Function load, 10

G

gnat sources, 5
 gnat.traces.syslog, 32
 gnat_util, 5
 GNATCOLL.Email, 54
 GNATCOLL.Email.Mailboxes, 54
 GNATCOLL.Email.Parser, 54
 GNATCOLL.Email.Utils, 53
 GNATCOLL.Projects, 5
 GNATCOLL.Python, 11
 gnatcoll-python.ads, 11

I

iconv, 83

J

json, 81

K

Knuth, 49

L

log, 29
 Logger, 27

M

MIME, 53
 mmap, 43

P

paragraph filling, 49
 Procedure Register_Python_Scripting, 14
 Procedure Register_Shell_Scripting, 14
 Procedure Register_Standard_Classes, 14
 projects, 5
 aggregate projects, 118
 Python, 11

R

ravenscar, [57](#)
Record Class_Type, [17](#)
reference, [72](#)
reference counting, [71](#)
Register_Command, [18](#)

S

script module, [16](#)
search, [47](#)
syslog, [32](#)

T

templates, [51](#)
test driver, [9](#)
testing your application, [9](#)
Trace_Handle, [27](#)

V

Virtual_Console.GetInstance, [15](#)
Virtual_Console.Insert_Error, [15](#)
Virtual_Console.Insert_Log, [15](#)
Virtual_Console.Insert_Prompt, [15](#)
Virtual_Console.Insert_Text, [15](#)
Virtual_Console.Read, [15](#)
Virtual_Console.Set_As_Default_Console, [15](#)
Virtual_Console.Set_Data_Primitive, [15](#)

W

weak, [72](#)