

GNAT-AJIS User's Guide

GNAT Ada-Java Interfacing Suite

A Toolkit for GNAT, the GNU Ada Compiler
GNAT GPL Edition, Version 2009

AdaCore

Copyright © 1995-2008, Free Software Foundation

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with the Invariant Sections being “GNU Free Documentation License”, with the Front-Cover Texts being “GNAT-AJIS User’s Guide” and “GNAT Ada-Java Interfacing Suite”, and with no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

About This Guide

This guide describes the features and the use of GNAT-AJIS, the GNAT Ada-Java Interfacing Suite that can be used with the GNAT Ada development environment.

What This Guide Contains

This guide contains the following chapters:

- [Chapter 1 \[Getting Started with GNAT-AJIS\], page 3](#), describes how to set up your environment and illustrates the use of the GNAT-AJIS toolset
- [Chapter 2 \[Using ada2java to Generate Java Classes\], page 9](#), describes how to use `ada2java` to generate Java classes that can be used as a “thin binding” to an Ada package specification.
- [Chapter 3 \[Mapping Ada to Java\], page 15](#), describes how Ada features are mapped to Java by `ada2java`.
- [Chapter 4 \[Advanced ada2java Topics\], page 27](#), describes some of the advanced aspects of the mapping of Ada to Java.
- [Chapter 5 \[Using javastub to Generate Ada Package Specifications\], page 47](#), describes how to use the `javastub` utility.
- [Chapter 6 \[Using jvm2ada to Interface Ada with Java\], page 49](#), describes to use the `jvm2ada` utility.
- [Appendix A \[Using JNI Directly\], page 51](#), describes how to write native methods in Ada with the same low-level style as in C.

What You Should Know before Reading This Guide

Before reading this manual you should be familiar with the following:

- The Ada programming language, and in particular the Ada 2005 Object-Oriented Programming and separate compilation enhancements (e.g., `limited with`, `interfaces`) that facilitate interfacing between Ada and Java.
- *The GNAT User Guide*
- The Java programming language

It will also be useful if you have a basic knowledge of the following:

- JNI (Java Native Interface), for example as described in *The Java Native Interface Programmer's Guide and Specification*, S. Liang, Addison-Wesley; 1999.

1 Getting Started with GNAT-AJIS

This chapter summarizes GNAT-AJIS's basic capabilities and illustrates how to use the GNAT-AJIS tools for some simple applications.

1.1 Introduction

GNAT-AJIS (GNAT Ada-Java Interfacing Suite) is a collection of GNAT add-on tools for developing mixed-language Ada / Java applications where the Java components run on a JVM and the Ada components are compiled natively. Through GNAT-AJIS you can realize the following scenarios:

1. In a Java application, invoke subprograms from natively-compiled Ada packages (i.e., either interface with an existing Ada API, or implement Java native methods in Ada);
2. In a natively compiled Ada program, access methods and fields from Java classes or objects.

GNAT-AJIS addresses these scenarios through an Ada binding to the JNI services and “binding generator” tools that automate the generation of the necessary “glue code”:

- | | |
|-----------------------|---|
| <code>ada2java</code> | Takes an Ada package specification as input and produces one or more Java classes, with native methods corresponding to the Ada subprograms. This allows you to call Ada from Java. |
| <code>javastub</code> | Takes a Java classfile and produces an Ada package spec for the native methods found in these files. This allows you to implement Java native methods in Ada. |
| <code>jvm2ada</code> | When given the ‘-jni’ option, takes a Java classfile and produces corresponding Ada specs that can be used to invoke the Java methods. This allows you to call Java from Ada. |

An alternative technique for combining Ada and Java is to compile Ada directly to JVM bytecodes. This functionality is provided by the JGNAT tool, which allows direct (i.e., not JNI-related) communication between Ada and Java based on the standard Ada interfacing mechanisms. JGNAT handles a subset of Ada that is compatible with the semantics of the Java Virtual Machine (for example, representation clauses are not supported). It is a separate tool, not part of the GNAT Ada-Java Interfacing Suite, and is fully described in the *JGNAT User's Guide*.

1.2 GNAT-AJIS Installation Structure

Installing the GNAT-AJIS tools results in the following directory structure:¹

```
$GNATAJIS_INSTALL_DIR/
  bin/
    ada2java (Solaris, Linux) or ada2java.exe (Windows) – executable
  include/
    ajis/
      Various ‘.ads’ and ‘.adb’ files
    gnatjni/
      Various ‘.ads’ and ‘.adb’ files
  lib/
    libajis.so (Solaris, Linux) or ajis.dll (Windows)
    libgnatjni.so (Solaris, Linux) or gnatjni.dll (Windows)
    ajis.jar
    ajis/
      Various ‘.ali’ files
    gnat/
      Various files
    gnatjni/
      Various ‘.ali’ files
```

1.3 GNAT-AJIS / GNAT Compatibility

ada2java is based on ASIS and require a compatible version of the GNAT compiler. To check the status of your installation, run ada2java with the ‘-v’ switch; this will indicate the version of the GNAT compiler that was used to build the GNAT-AJIS suite. Your GNAT-AJIS installation is compatible with that GNAT version:

- Running ada2java requires using that specific GNAT version;
- On the other hand, the generated Ada files may be compiled with that version or any later one.

The gnatjni and ajis libraries have been prebuilt for a specific version of GNAT. If you need to compile them for some other version of GNAT, you can rebuild the libraries manually:

```
gprbuild -P ajis.gpr -XExternal_Build=false -XObject_Dir=<some-dir>
```

where <some-dir> is a local directory where the temporary objects will be placed.

1.4 A Simple Example: Calling Ada from Java

This section illustrates how to invoke an Ada subprogram (compiled natively) from Java running on a JVM. In summary, the steps are as follows:

¹ For simplicity, Unix-style notation is used throughout this manual in depicting directories and other host system conventions. For Windows, please make the relevant transformations (e.g. ‘\’ for ‘/’ in path names).

- Make sure that the relevant environment variables are properly defined.
- Write a package specification for the subprogram(s) to be called from Java, and a corresponding package body.
- Invoke the GNAT-AJIS tool `ada2java` on the Ada package spec, to produce the corresponding Java classes (source files) and the necessary JNI “glue” code (additional Ada source files). Providing the ‘`-L libname`’ switch will cause a project file to be generated, which will help to automate some of the processing.
- Invoke the Java compiler `javac` on the Java source files;
- Invoke `gprbuild` on the project file generated by `ada2java`; this will compile the Ada files into a shared library (Solaris, Linux) or `dll` (Windows);
- Invoke the Java interpreter to run a Java main class that invokes methods from the Java classes generated by `ada2java`.

These steps will now be described in detail.

1.4.1 Environment Setup

Since you will be using both the Ada and Java toolsets, you need to ensure that several environment variables are set. You can automate this step by defining these variables in a shell script / batch file. For convenience you will also find it useful to define an environment variable that “points to” the root directory for the GNAT-AJIS tool installation. The description below assumes that `GNATAJIS_INSTALL_DIR` has this role.

`PATH` Must contain the directories for the GNAT tools and for the GNAT-AJIS tools. The latter will be in the `GNATAJIS_INSTALL_DIR/bin` directory. On Windows, it needs to contain the directory where the shared libraries are generated, typically ‘`./lib`’ although you can override this.

`LD_LIBRARY_PATH` On Solaris and Linux, must contain the directories where your native libraries will reside (generally the ‘`./lib`’ subdirectory). This variable is not needed on Windows.

`CLASSPATH` Must contain `GNATAJIS_INSTALL_DIR/lib/ajis.jar`, which is the parent directory of the ‘`com.adacore.ajis`’ Java package.

`ADA_PROJECT_PATH` Must contain ‘`GNATAJIS_INSTALL_DIR/lib/gnat`’, the directory that holds the GNAT project files needed for building applications with GNAT-AJIS.

1.4.2 An Ada Package

Assume that you would like to invoke an Ada procedure that displays the text `Hello from Ada`, followed by an integer value passed to Ada from Java. Declare a procedure `Hello` in a package spec `Hello_Pkg` (file `'hello_pkg.ads'`) and implement the body (file `'hello_pkg.adb'`):

```
package Hello_Pkg is
  procedure Hello (Item : in Integer);
end Hello_Pkg;

with Ada.Text_IO; use Ada.Text_IO;
package body Hello_Pkg is
  procedure Hello (Item : in Integer) is
  begin
    Put_Line("Hello from Ada: " & Integer'Image(Item));
  end Hello;
end Hello_Pkg;
```

1.4.3 Invoking `ada2java`

Change to the directory containing the Ada source files, and invoke the command

```
ada2java hello_pkg.ads -L hello_proj
```

This will generate a number of files and directories, including:

```
Hello_Pkg/
  Hello_Pkg_Package.java
```

```
Ada2Java/
  Library.java
```

```
hello_proj.gpr
```

Specs and bodies for the `JNI_Binding` package hierarchy

These have the following significance:

Directory `'Hello_Pkg'`

In the absence of an option that specifies the output directory for the generated Java file, `ada2java` creates a new directory with the same name as the Ada input unit and places the Java file in this directory.

File `'Hello_Pkg_Package.java'`

`ada2java` generates a Java source file with native method(s) corresponding to the visible subprogram(s) in the Ada package. (In general `ada2java` may generate several Java source files, based on the contents of the Ada package spec. In this example only one Java file is produced.) The name of this file is the same as the Ada unit,

with `_Package` appended (since the input file is a package, rather than a procedure or function). The casing of the file name is the same as that specified on the Ada unit declaration.

Ada parameters are mapped to Java types; here Ada's `Integer` corresponds to the Java type `int`.

In skeletal form, here is the Java class that is generated:

```
package Hello_Pkg;

public final class Hello_Pkg_Package {

    static public void Hello (int Item){...}
    ...
}
```

Directory 'Ada2Java' and file 'Library.java'

`ada2java` generates the boilerplate file 'Library.java' to automate the library load step.

File 'hello_proj.gpr'

This is a GNAT project file that automates building the application and loading the dynamic library.

Specs and bodies for the `JNI_Binding` package hierarchy

These files provide various "boilerplate" packages as well as the package containing the "glue code" procedure whose signature complies with the required JNI protocol and which invokes the `Hello` procedure supplied in the original `Hello_Pkg` package.

1.4.4 Compiling the Java class

Invoke the Java compiler on the generated Java class:

```
$ javac Hello_Pkg/Hello_Pkg_Package.java
```

This will generate the classfile 'Hello_Pkg_Package.class' in the 'Hello_Pkg' directory.

1.4.5 Building the Application

Run `gprbuild`, using the project file generated by `ada2java` at an earlier step:

```
$ gprbuild -p -P hello_proj.gpr
```

This will generate a dynamic library – 'libhello_proj.so' (Solaris, Linux) or 'hello_proj.dll' – in the subdirectory './lib' of the current directory, and will produce the necessary object files in the './obj' subdirectory. The two subdirectories will be created if they do not already exist.

The dynamic library will be loaded automatically at run-time, by one of the generated Java classes.

1.4.6 Running the Program

Write a main Java class, for example a file 'Hello.java':

```
import Hello_Pkg.Hello_Pkg_Package;

public class Hello{
    public static void main(String[] args){
        Hello_Pkg_Package.Hello(100);
    }
}
```

Compile this class:

```
$ javac Hello.java
```

Run the Java program:

```
$ java Hello
```

This will produce the following output:

```
Hello from Ada: 100
```

2 Using `ada2java` to Generate Java Classes

The `ada2java` tool takes one or more Ada package specs and produces as output a Java “binding” to these packages, implemented through JNI. The binding consists of a set of Java classes, with methods that access the Ada package’s visible entities.

More specifically, `ada2java` generates two sets of source files as output:

- The Java classes that make up the binding, and
- The necessary Ada “glue code” that hides the details of how JNI is used for interfacing between Ada and Java.

You will need to compile the Java files to bytecodes for execution on a JVM, and you will need to compile the Ada files to native code in a dynamic library.

This chapter explains how to use the `ada2java` tool and describes the mapping from package spec contents to Java classes.

2.1 Using the Tool

The `ada2java` tool is invoked with at least one input file, and any number of switches, in any order:

```
$ ada2java {switch | input-file} input-file {switch | input-file}
```

Each *input-file* must be the name of a source file for an Ada package spec (including the extension).

The following *switch* values are allowed:

‘-h’ Display help

‘-c *JavaClassOutputDirectory*’

The root directory used as the destination for the output classes. The directory will be created if it does not already exist. In the absence of this switch, the current directory is used. See below for the relationship with the ‘-b’ switch.

‘-b *BaseJavaBindingPackage*’

The base package for the generated Java classes; this will be relative to the directory specified in the ‘-c’ switch, or relative to the current directory if no ‘-c’ switch was supplied.

‘-o *AdaGlueOutputDirectory*’

The destination directory for the “glue” packages (‘ads’ and ‘adb’ files) generated by `ada2java`. The current directory will be used if this switch is not supplied. The generated packages will need to be compiled into a dynamic library.

'-P *ProjectFile*'

The project file that applies to the processing of the *input-files* submitted to `ada2java`. This can specify compiler switches, source directories, etc. *ProjectFile* must be a “flat” project (sources from “with”ed projects are not yet supported).

'-L *LibraryName*'

A mechanism for automating the loading of the native Ada dynamic library in Java. This switch causes the generation of a project file '*LibraryName.gpr*' in the directory specified by the '-o' switch (or in the current directory if the '-o' switch was not supplied). The resulting project file can be submitted to `gprbuild` to build the dynamic library:

```
$ gprbuild -p -P LibraryName.gpr
```

which will generate a 'lib/' subdirectory that contains the file '`libLibraryName.so`' (Solaris, Linux) or '`LibraryName.dll`' (Windows). This library will be loaded automatically whenever one of the Java classes produced by `ada2java` is loaded; there is no need for the user to explicitly include an invocation of `System.loadLibrary`.

'-M *MainName*'

A mechanism for automating the creation of an Ada main subprogram, embedding both the native code and a JVM. See [Section 2.2.3 \[Compiling as an Ada Main Subprogram\], page 12](#) for more details. Implies `-link-method=register_natives`.

'--main-class=*java main class*'

Changes the name of the java main class to use, in case the '-M' switch is used. See [Section 2.2.3 \[Compiling as an Ada Main Subprogram\], page 12](#) for more details.

'--link-method=(*export* | *register_natives*)'

The Java virtual machine has two ways of discovering the functions declared in the native environment. Either it checks the correspondence between the exported symbol and the Java native declaration name (*export* mode), or the JNI code registers manually the symbols using the `Register_Native` JNI function (*register_natives* mode). Note that if the code is not in a shared library but compiled with a main native subprogram, then only *register_natives* mode will work.

'--no-monitor[-finalize]'**'--monitor[-finalize]-(*check* | *protect*)'**

Sets the default monitor for subprograms. See [Section 4.3 \[Thread Safety\], page 34](#).

- '--[no-]attach-(parameter|access|controlling|ada2005)'
Sets the default attachment policy. See [Section 4.7 \[Managing Attachment to Java Proxies\]](#), page 45
- '--[no-]assume-escaped'
Controls whether checks for object ownership are enabled. See [Section 4.2.6 \[Restrictions on Proxy-Owned Objects passed to Subprograms\]](#), page 32
- '--[no-]java-enum'
Controls whether java enumerations should be used to bind ada enumerations, or if static integers should be used instead (java enumeration is default).

Example:

```
$ ada2java -c mydir pack1.ads -b foo.bar
```

This results in the placement of the Java binding classes in the relative directory `'mydir/foo/bar/'`.

Note that the actual directory containing the generated Java classes will need to be on the `CLASSPATH` environment variable in order to successfully run a Java application that uses the binding.

2.2 Compiling and Running the Generated Code

2.2.1 Issues with the Ada Generated Code

Two sets of Ada units need to be compiled – the original packages and the generated "glue" code. The Ada glue depends on the `ajis` project installed in the `'lib/gnat'` directory of the GNAT-AJIS installation.

It is highly recommended that you use the project generation switches `'-L'` (for a shared library) or `'-M'` (for an Ada main subprogram). However, even if these switches handle most cases, you may need to write your own build procedures to address more advanced usage. In such a situation please note that some compiler options may have an impact on the `ajis` library and thus need to be taken into consideration:

- '-O2 -O3' If you compile with a high optimization level, you should deactivate strict aliasing using the compiler switch `'-fno-strict-aliasing'`.
- '-fstack-check'
The stack checking mechanism is based on signals that are deactivated by the GNAT AJIS library, so this switch will have no effect and should not be used.
- '-fPIC' On Linux / Solaris, all the code has to be relocatable, which is specified through the `'-fPIC'` switch. If you are creating a shared library

that integrates components compiled externally, you have to ensure that they have been compiled using the `'-fPIC'` switch.

2.2.2 Compiling as an Ada Shared Library

The most common architecture of an Ada / Java program, and a Java / Native program in general, is to compile the native code into a shared library, and then load that shared library at run time. In this case, the main entry point is a Java main method, written by the developer.

In order to implement this scheme, you will need to create a SAL (Stand-Alone Library) project containing the sources of the input packages plus the “glue”, and use it to compile the library.

A simple standalone library project is generated if you use the `'-L'` switch. The generated project can then be compiled with `gprbuild`, for example:

```
$ ada2java my_package.ads -o ada -c java -P my_project.gpr -b base -L my_lib
$ gprbuild -p -P ada/my_lib.gpr
```

Note that the native library will then be loaded automatically by the generated Java glue code.

2.2.3 Compiling as an Ada Main Subprogram

If compiling the native code into a shared library is not practical, an alternative is to create an Ada main subprogram embedding a Java Virtual Machine.

`ada2java` provides an easy way to generate a project and an Ada main subprogram, through the `'-M'` switch. This switch takes the name of the main as parameter and will generate an Ada main that will automatically create a Java virtual machine, and then call a Java method defined as follows:

```
package <base_package>;

public class <main_name> {

    public static void main (String [] args) {
    }

}
```

This class (and thus the method implementation) has to be provided by the developer. If it is not present, the main subprogram will fail with an error at run time.

The generated main will look into the `CLASSPATH` environment variable to find the Java classes when initializing the Java virtual machine. So for example, if that you provide the following class:

```
package java_code;

import java_code.Test.Test_Package;
```

```
public class Main {  
  
    public static void main () {  
        Test_Package.Call_Something ();  
    }  
  
}
```

using the following Ada API:

```
package Test is  
    procedure Call_Something;  
end Test;
```

with the appropriate `'test.gpr'` project referencing the `Test` code, you will be able to compile and run the code as follows:

```
$ ada2java test.ads -P test -b java_code -o ada -c java -M Main  
$ gnatmake -P ada/main.gpr  
$ CLASSPATH='pwd': 'pwd'/java:$CLASSPATH  
$ export CLASSPATH  
$ javac java_code/Main.java  
$ ada/obj/main
```

You can explicitly specify the name of the Java main class to use, through the `'--main-class'` switch, e.g.:

```
$ ada2java test.ads -P test -b java_code -o ada -c java \  
> -M Main --main-class=some.main.My_Main
```

In this case, the Ada main will look for a main subprogram in `some.main.My_Main`, instead of `java_code.Main`.

Note that you may need to define the `LPATH`, `LD_LIBRARY_PATH` or `PATH` environment variables so that the code can be compiled against `'jvm.lib'` or `'libjvm.a'`, and then run with `'jvm.dll'` or `'libjvm.so'`.

2.2.4 Compiling the Java Generated Classes

The Java application needs to load the library before any of the Ada subprograms are invoked. If you did not supply the `'-L'` switch to `ada2java`, then you will need to do this explicitly; conventional style is to invoke `System.loadLibrary ("library-name")` in a static initializer in the main Java class. This step is automated if you use the `'-L'` switch, as described above.

Before running the Java code, you need to ensure that the `CLASSPATH` environment variable contains both the directory of the generated Java code, and the JAR for the GNAT-AJIS-related predefined classes. The latter archive exists as `$GNATAJIS_INSTALL_DIR/lib/ajis.jar` where `GNATAJIS_INSTALL_DIR` is the root directory for the GNAT-AJIS installation.

2.3 Pragma Annotate and `ada2java`

Pragma `Annotate` (see *GNAT Reference Manual*) has several uses in conjunction with the GNAT-AJIS tools, each with the form:

```
pragma Annotate (AJIS, AJIS_annotation_identifier {, argument});
```

GNAT-AJIS annotation names are defined in the package `AJIS.Annotations`, which is a part of the `ajis.gpr` project installed with GNAT-AJIS. You need to have visibility on this package using a `with` and possibly a `use` clause before being able to use these pragmas.

The following GNAT-AJIS annotation pragmas are supported:

- `Annotation_Renaming` – [Section 4.1 \[Dealing with Name Clashes\], page 27](#)
- `Assume_Escaped` – [Section 4.2.6 \[Restrictions on Proxy-Owned Objects passed to Subprograms\], page 32](#)
- `Attached` – [Section 4.7 \[Managing Attachment to Java Proxies\], page 45](#)
- `Monitor` – [Section 4.3 \[Thread Safety\], page 34](#)
- `Rename` – [Section 4.1 \[Dealing with Name Clashes\], page 27](#)

3 Mapping Ada to Java

To allow an Ada package to be used from Java, `ada2java` generates one or more Java classes (source files that will need to be compiled to bytecodes by a Java compiler) based on the content of the visible part of the Ada package spec. This section explains and illustrates the mapping for each of the various kinds of entities declared in a package that can be used from Java.

In brief:

- Although there are some exceptions to this rule, in general a type and certain of its associated subprograms declared in an Ada package are mapped to a Java class with methods corresponding to the Ada subprograms. Such entities are said to be *attached* to the resulting class.
- Other entities declared in the Ada package map to static members defined in a “default class” generated by `ada2java`. In particular, variables and constants in the Ada package map to private static fields in the default class and are accessed through “getter” (and “setter” for variables) methods. Such entities are said to be *unattached*.

If the default class is generated, its name is that of the original Ada package (with the same casing as the identifier in the package declaration) suffixed with `_Package`.

In the examples, only the portions of the Java classes needed by users of the classes are shown.

3.1 Types

Types used in the Ada package map to Java types in the generated class(es). This section explains the correspondence. As a general rule, note that while most forms of type declarations have a correspondence in Java, subtype declarations are ignored, as there is no equivalent to subtypes in Java. However, subtype constraints imposed on Ada entities, such as variables or formal parameters, must be respected when referenced from Java, and can result in exceptions when constraints are violated.

3.1.1 Scalar Types

The following table shows how Ada scalar types are mapped to Java primitive types:

Ada type	Java type
<i>Integer type</i> <= 32 bits	<code>int</code>
<i>Integer type</i> > 32 bits	<code>long</code>
<code>Boolean</code>	<code>boolean</code>

Character	char
<i>Other enumeration type</i>	int
<i>Fixed-point type</i>	double
<i>Floating-point type</i>	double

Constraint checks generated in the Ada glue code detect errors that may result from the range mismatches between Ada and Java. For example, since a 16-bit Ada integer will be mapped to 32-bit `int` in Java, the Java code might attempt to pass an out-of-range value to Ada. This will raise a `Constraint_Error` exception in Ada, which will be propagated back to Java as an `AdaException` exception.

For an enumeration type, a Java final class is created, with the same name as the enumeration type. This class defines the possible values for the enumeration.

Example:

```
package Pkg is
  type Enum is (A, B, C);
end Pkg;
```

will give:

```
package Pkg;

public final class Enum {
  public static final int A = 0;
  public static final int B = 1;
  public static final int C = 2;
}
```

Representation clauses for enumeration types are not currently supported.

A discussion of subprogram formal parameters of scalar types may be found in [Section 3.3.2 \[Subprogram parameters\]](#), page 22.

3.1.2 Arrays

Mapping Ada arrays to Java arrays would be very expensive, since it would imply a copy of the whole array each time a parameter has to be passed. Thus for efficiency an Ada array type is mapped to a dedicated “proxy” class with methods that serve as accessors to attributes and components. For example:

```
package Ex1 is
  type T1 is array(Integer range <>) of Float;
end Ex1;
```

will yield the following class:

```
public final class T1 extends com.adacore.ajis.internal.ada.AdaProxy {
  ...
  public T1 (int First_1, int Last_1){...}

  final public double Get_Element_At (int Index_1){...}
```

```

    final public void Set_Element_At (int Index_1, double Value){...}

    final public int First (){...}

    final public int Last (){...}

    final public int Length (){...}
}

```

A subprogram that takes a parameter of the Ada array type is mapped to a method taking a parameter of the corresponding Java “proxy” class; note that this method is located in the default class, and not in the proxy class.

3.1.3 Strings

Directly passing `String` data between Ada and Java would require expensive copying, and thus an alternative approach is used. The Ada type `String` is mapped to the Java class `AdaString`, which encapsulates the accesses.

More specifically, an Ada parameter of type `String` of any mode, and an Ada access `String` parameter, are both mapped to a Java parameter of type `AdaString`.

For efficiency, an `AdaString` object caches both its Ada and Java string values after they have been computed. As an example, if the Ada spec is:

```

package Pckg is
  procedure P (V : String);
end Pckg;

```

then the generated Java will be:

```

public final class Pckg_Package {
  public static void P (V : AdaString) {...}
}

```

If we now write:

```

AdaString str = new AdaString ("A string from Java");
Pckg_Package.P (str);
Pckg_Package.P (str);

```

Only the first call will require the expensive string translation from Java to Ada. The second invocation will directly use the cached value.

Please note that Java strings are UTF16-encoded, whereas the corresponding Ada strings will be UTF8-encoded. This may have significant impact when computing character offset on Java strings.

3.1.4 Simple Record Types

Simple (that is, not tagged) record types are mapped to Java final classes. Components are accessed through a set of generated accessors (“getter”/“setter”

methods). As a current limitation, `ada2java` does not yet support accessing discriminant components.

Example:

```
package Pkg is
  type R is
    record
      F1 : Integer;
      F2 : Float;
    end record;
end Pkg;
```

will give:

```
package Pkg;
public final class R {
  public R () {...}

  public final int F1 () {...}
  public final void F1 (int Value) {...}

  public final double F2 () {...}
  public final void F2 (double Value) {...}
}
```

A component that has an access-to-record type is treated as though it were of the record type itself. For example:

```
package Pkg is
  type R is
    record
      F1 : Integer;
      F2 : Float;
    end record;
  type S is
    record
      G1 : R;
      G2 : access R;
    end record;
end Pkg;
```

will result in both a class `R` as above, and the following class `S`:

```
package Pkg;
public final class S {
  public S () {...}

  public R G1 () {...}
  public void G1 (R Value) {...}

  public R G2 () {...}
  public void G2 (R Value) {...}
}
```

Only one level of indirection is implemented; `ada2java` does not support access to access-to-record.

A private (untagged) type is treated like a record type, except that it does not have any component-accessing methods. (A later release of `ada2java` will generate methods for accessing discriminants if the type has any.)

3.1.5 Tagged Types

A tagged type is mapped to a Java class of the same name. If the Ada type is abstract, then the Java type will be abstract as well.

3.1.5.1 General principles

A primitive (i.e., dispatching) subprogram of a tagged type is mapped to a corresponding Java instance method. A current restriction is that the first parameter of the Ada subprogram must be a controlling parameter; otherwise the subprogram is mapped to a method in the default class. (Thus a function that delivers a value of the tagged type, but has no controlling parameter, is mapped to a method in the default class, and not to a method in the class corresponding to the Ada type.) The first Ada parameter is mapped to the Java method's implicit `this` parameter.

A subprogram with a class-wide parameter is mapped to a method of the tagged type's Java class whose corresponding parameter has the Java class type. However, as this is not properly a dispatching primitive of the Ada type, it is declared as a final method.

As an example:

```
package Ex1 is
  type T is tagged null record;
  procedure P1 (X : in out T; F : Float);
  procedure P2 (X : T'Class);
  procedure Q1 (I : Integer; X : T);
  procedure Q2 (I : Integer; X : T'Class);
  function F return T;
end Ex1;
```

is mapped to:

```
public final class Ex1_Package {
  ...
  static public void Q1 (int I, Ex1.T X){...}

  static public void Q2 (int I, Ex1.T X){...}

  static public T F (){...}
} // Ex1_Package
```

```
public class T extends com.adacore.ajis.internal.ada.AdaProxy {
    ...
    public void P1 (double F){...}

    final public void P2 () {...}
} // T
```

3.1.5.2 Ada type hierarchies

Hierarchies of Ada types are preserved in the generated Java classes. Therefore, the following structure:

```
type R is tagged record;

type R_Child is new R with null record;
```

will result in:

```
public class R {...}

public class R_Child extends R {...}
```

Consistency of Java types is guaranteed at run time. For example, the following function:

```
package Pckg is
    function F return R'Class;
end Pckg;
```

will result in:

```
public final class Pckg_Package {
    public R F () {...}
}
```

However, if the actual type of the returned object is `R_Child`, then the value returned by the Java function will be of the Java type corresponding to `R_Child`.

3.1.5.3 Java class hierarchies

It is possible to extend a Java class that was generated by `ada2java` from an Ada tagged type.

For example:

```
package Rec_Pckg
    type Rec is tagged null record;
    procedure P (R : Rec);
end Rec_Pckg;
```

results in a Java class `Rec` with an instance method `P`:

```
class Rec extends com.adacore.ajis.internal.ada.AdaProxy {
    ...
    public void P(){...}
}
```

You can then write:

```

class Rec_Child extends Rec {
    public void P () {
        System.out.println ("Hello from Java");
    }
}
...
Rec ref = new Rec_Child();
ref.P(); // Displays "Hello from Java"

```

3.2 Global Variables and Constants

A package containing global variables (that is, variables declared in the package spec's visible part) is mapped to a default class containing “getter” and “setter” methods for accessing and updating the variables. Global constants are treated analogously, but they have only a “getter” method. Variables of limited types also only have a “getter” method. Note that Ada named numbers, which are really just values intended for use in static compile-time computations, are not mapped to Java.

For example:

```

package Globals is
    V : Integer;
    C : constant Integer := 100;
    N : constant := 3.14159;
end Globals;

```

will result in the default class:

```

public class Globals_Package {

    static public int V () {...}

    static public void V (int Value){...}

    static public int C () {...}
}

```

3.3 Subprograms

Ada procedures and functions are mapped to Java methods. Nondispatching subprograms are marked as `final`. Dispatching subprograms are discussed in [Section 3.1.5 \[Tagged Types\], page 19](#).

3.3.1 Method placement

In general, nondispatching subprograms are mapped to methods defined in the default class. For example, if the input package spec is:

```
package Pkg is
  function F return Integer;
end Pkg;
```

then `ada2java` will generate the following default class:

```
public class Pkg_Package{
  ...
  public static int F(){...}
}
```

However, there are cases where the subprogram can be attached to the class of its first parameter. Attachment can be enabled / disabled depending on user requirement. In this case, the explicit initial Ada parameter is mapped to the implicit `this` parameter in Java. See [Section 4.7 \[Managing Attachment to Java Proxies\]](#), [page 45](#) for further details.

3.3.2 Subprogram parameters

The following rules and restrictions apply to the types of subprogram formal parameters:

- Scalar types
 - Access-to-scalar types are not supported.
 - A scalar type with mode `in` is mapped to the corresponding Java type. For example:

```
procedure P (V : Integer);
```

will result in:

```
public void P (int V) {...}
```

- A formal scalar with mode `out` or `in out` will be mapped to a corresponding “wrapper” class: `BooleanRef`, `CharacterRef`, `DoubleRef`, `IntegerRef`, and `LongRef` respectively encapsulating the primitive type `boolean`, `character`, `double`, `int` and `long`).

Each of these classes defines `setValue` and `getValue` methods for accessing the encapsulated value.

The Java application needs to construct an object of the relevant wrapper class and pass it to the method that corresponds to the Ada subprogram. After the return from the method, the Java application can invoke the `getValue` method to retrieve the new value of the actual parameter.

- Record and private types
 - An Ada `in`, `in out`, or `out` formal parameter of a record or private type (either tagged or untagged), is mapped to a Java formal parameter of the class corresponding to the Ada type. Similarly, an Ada formal parameter of an access-to-record-type or access-to-private-type (either

anonymous or named) is mapped to a Java formal parameter of the class corresponding to the Ada type.

Example:

```
package Example is
  type R is null record;
  type Access_R is access all R;
  procedure P(V1 : R;
             V2 : out R;
             V3 : in out R;
             V4 : access R;
             V5 : Access_R);
end Example;
```

The resulting Java class is:

```
public final class Example_Pkg {
  ...
  public void P (R V1, R V2, R V3, R V4, R V5){...}
}
```

- An Ada `out` or `in out` parameter of an access-to-record (or access-to-private) type is mapped to a nested class. For example:

```
package Example is
  type R is null record;
  type Access_R is access all R;
  procedure P(V : out Access_R);
end Example;
```

will generate the default class and a class for R

```
public class Example_Package {
  ...
  static public void P (R.Ref V){...}
}

public class R extends com.adacore.ajis.internal.ada.AdaProxy {
  ...
  public static class Ref implements com.adacore.ajis.IProxyRef {
    public void setValue (Object r) {...}
    public Object getValue () {...}
  }
}
```

The Java application needs to construct an object of the class `R.Ref` and pass it to `P`. On return, the `getValue` method may be called to retrieve the value in the `out` parameter returned by the Ada procedure.

- Further indirection, such as an access-to-access type for a formal parameter, is not supported.

3.4 Exceptions

Exceptions are bound into classes derived from `com.adacore.ajis.NativeException`. It is then possible to throw or handle them directly in Java code.

Example:

```
package Example is
  An_Exception : exception;
  procedure Raise_An_Exception;
end Example;

package body Example is
  procedure Raise_An_Exception is
  begin
    raise An_Exception;
  end Raise_An_Exception;
end Example;
```

The resulting Java class is:

```
public final class An_Exception extends com.adacore.ajis.NativeException {
  ...
}
```

And can be used in e.g.:

```
try {
  Example_Package.Raise_An_Exception ();
} catch (An_Exception e) {
  // process the exception
}
```

3.5 Renamings

Renamings of objects and subprograms are supported by `ada2java`. Object renamings are mapped in the same way as global objects, by means of “setter” and “getter” methods in the default class for the containing package. A subprogram renaming is represented by a method with the name of the renaming that invokes the renamed subprogram, declared in the appropriate class. In other words, the same rules that apply to other subprograms apply to subprogram renamings.

3.6 Predefined Environment

In order to access descendants of Ada or GNAT from Java, you need to manually invoke `ada2java` on the Ada source files from the GNAT installation directories, to generate the corresponding Java binding classes. This step will be automated in a future release of GNAT-AJIS.

3.7 Current Limitations

The following features are not supported:

- *Discriminants*. Discriminants are not accessible from the Java class generated for a discriminated type.
- *Anonymous arrays*. Objects with an anonymous array type are not supported, but array type declarations which declare a constrained first subtype are supported.
- *Generic units*. Generic declarations and instantiations are ignored.
- *Interfaces*. No mapping is currently provided from Ada interface types to Java interfaces.
- *Tasking features*. Tasks and protected objects/types are ignored.

4 Advanced ada2java Topics

This chapter discusses a number of issues that `ada2java` users should be aware of.

4.1 Dealing with Name Clashes

If Ada subprograms from the same package spec produce the same Java profile, the generated Java class will fail to compile. To prevent this, you can use the `Rename` pragma to define the Java name corresponding to an Ada entity:

```
pragma Annotate (AJIS, Rename, identifier, static_string_expression);
```

The *identifier* argument denotes the Ada entity. The *static_string_expression* is the name that will be used for the corresponding Java entity generated by `ada2java`.

Example:

```
package Example is
  type I1 is new Integer;
  type I2 is new Integer;

  function F return I1;
  pragma Annotate (AJIS, Rename, F, "F_I1");

  function F return I2;
  pragma Annotate (AJIS, Rename, F, "F_I2");
end Example;
```

This will result in two Java functions:

```
int F_I1 ();
int F_I2 ();
```

Adding a pragma to an Ada package specification is not always practical, and indeed may be impossible if the specification is from an external library. With GNAT-AJIS, you can provide the pragma in a separate Ada file, applying it to an Ada entity that is itself a renaming declaration.

This entity must be marked by another `Annotate` pragma:

```
pragma Annotate (AJIS, Annotation_Renaming, identifier).
```

where *identifier* denotes an entity defined by an Ada renaming declaration. The `Annotation_Renaming` pragma applies to all AJIS pragmas that are specified for *identifier*.

```
package Example is
  type I1 is new Integer;
  type I2 is new Integer;

  function F return I1;
  function F return I2;
end Example;
```

```
with Example;
package Renamings is
  function F return Example.I1 renames Example.F;
  pragma Annotate (AJIS, Annotation_Renaming, F);
  pragma Annotate (AJIS, Rename, F, "F_I1");

  function F return Example.I2 renames Example.F;
  pragma Annotate (AJIS, Annotation_Renaming, F);
  pragma Annotate (AJIS, Rename, F, "F_I2");
end Renamings;
```

Entities annotated with `pragma Annotate(AJIS, Annotation_Renaming)` will not be mapped to Java entities; they are assumed to be used only to define annotations.

4.2 Memory Model

An object on the Ada side – a so-called *native object* – is accessed in Java through a *proxy object*: an instance of a class (the *proxy class*) generated from the Ada object's type. The proxy object contains a reference to the native object. Invoking a member function on the proxy object results in a call of a native subprogram on the native object. The native object may be either allocated or declared.

This section explains the implications of this model on the usage of the generated binding.

4.2.1 Requirements for Non-null Parameter Values

If an Ada subprogram's formal parameter is not an access parameter (i.e., it has `in`, `in out` or `out` mode), then invoking the corresponding Java method requires a non-null reference to a proxy object. For example:

```
package P is
  type T is null record;

  procedure Proc (V : out T);
end P;
```

results in the following Java classes:

```
public class T {
  ...
}

public final class P_Package {
  public static Proc (T V) {
    ...
  }
  ...
}
```

You need to ensure that the `v` parameter passed to `Proc` is not `null`. Hence, the following will throw an exception, since `v` is implicitly initialized to `null`:

```
T v;
P_Package.Proc (v);
```

A simple way to provide a non-null reference is to initialize it to an allocated object:

```
T v = new T ();
P_Package.Proc (v);
```

4.2.2 Allocating Ada Objects from Java

If `A` is an Ada type that is mapped by `ada2java` to a Java class `J`, then the execution of a Java constructor `J()` will create two objects:

- A Java object (the “proxy” object) of class `J`, allocated on the Java heap and subject to Garbage Collection by the JVM, and
- An Ada object of type `A`, allocated on the Ada heap, referenced from the proxy object.

4.2.3 Automatic Creation of Native Objects

Under certain circumstances, the generated Java code may construct extra native objects on the Ada heap. To help explain this, here is an example where such allocation is not needed, namely, a function returning an access value:

```
package P is
  type T is record ... end record;

  type T_Acc is access all T;

  function Create_T_Acc return T_Acc;
end P;
```

The generated code will look like:

```
public class T { ...}

public final class P_Package {
  final public T Create_T_Acc () { ...}
  ...
}
```

On the Java side, the method `Create_T_Acc` returns a proxy object that contains the value of the pointer returned by the call of the Ada function. So the user can write:

```
T v = P_Package.Create_T_Acc ();
```

and then access the data of `v`. Note that all the standard precautions that apply in using an Ada pointer have to be taken in this case as well. In particular, after the object has been freed on the Ada side, there should be no further references

to it from Java. The Ada programmer needs to document how the returned value should or should not be used, and the Java programmer needs to adhere to these guidelines.

However, a function returning a record rather than an access value raises additional issues:

```
package P is
  type T is record ... end record;

  function Create_T return T;
end P;
```

results in a default class similar to the one in the previous example:

```
public final class P_Package {
  final public T Create_T () { ... }
  ...
}
```

The Java programmer can write:

```
T v = P_Package.Create_T ();
```

But in this case, the value returned by the function is not a pointer. A new Ada native object is automatically allocated on the heap, initialized to a copy of the value returned by the Ada function, and is referenced through the proxy object constructed by the Java method `Create_T`. This is similar to the Ada situation where calling `Create_T_Acc` would not involve a copy of a `T` object, whereas `Create_T` would.

Because of implementation constraints there may be more than one copy involved in the call of `Create_T` (but only one native object will be created on the heap). This will be discussed further in [Section 4.5 \[Clone and Copy Semantics\]](#), page 39.

4.2.4 Native Ownership

As illustrated above, native objects may be created automatically by Java methods corresponding to Ada subprograms. This raises the issue of when/how such objects are to be deallocated.

In general, the `ada2java` approach is based on the following principles:

- The environment (Java or Ada) that allocates an object is responsible for its deallocation;
- In Java, all deallocation is performed implicitly, by the Garbage Collector;
- In Ada, the programmer is responsible for manually deallocating the objects;
- Dangling references should be prevented; i.e., the Ada object should not be deallocated as long as there are still live references to the object.

A native object created from Java is said to be *owned* by its proxy object. It has been created by the Java program, and it must be freed by the Java environment.

Such an object is tightly linked to its proxy – when the proxy object doesn't exist (i.e. is garbage collected), the native object becomes inaccessible.¹ The native object deallocation will occur automatically, when the Java proxy is garbage collected (more specifically through the implementation of the `finalize` method).

A native object not owned by a proxy – for example, one that has been obtained from the Ada API through an Ada pointer – will not be deleted automatically.

The ownership state of a native object may be queried through the `getOwner` method. This function returns a value from the enumeration `com.adacore.ajis.IProxy.Owner`, either `NATIVE` or `PROXY`, specifying who is responsible for managing the memory. For example after these assignments:

```
T v1 = new T ();
T v2 = P_Package.Create_T_Acc ();
T v3 = P_Package.Create_T ();
```

the following relationships hold:

```
v1.getOwner ()   PROXY
v2.getOwner ()   NATIVE
v3.getOwner ()   PROXY
```

Native objects referenced by `v1` and `v2` will be deallocated when the corresponding Java proxy is garbage-collected.

You may change the owning attribute of a referenced native object, through the `setOwner` method of proxy classes. This should be used very carefully, as it may generate memory leaks or corruption. Doing the following:

```
T v1 = new T ();
v1.setOwner (Owner.NATIVE);
```

will deactivate the object deallocation on finalization. The Java programmer becomes responsible for explicitly deallocating the native object.

Note that, while moving the owning from `PROXY` to `NATIVE` is a relatively safe operation, raising the possibility of memory leaks but not object corruption, moving in the opposite direction should be done with great care. A object managed by the `NATIVE` side may be deallocated by the Ada application at any moment, or be declared as the field of another object or even on the native stack in situations such as callbacks. If it's not clear where the object is coming from, cloning it (resulting automatically into an object managed by the proxy) is often the best solution.

¹ As will be described below, cloning the proxy object has special semantics so that the native object gets copied, not only its reference.

As will be described below, only objects that are known to be managed through pointers can have their ownership changed. An object obtained, for example, from a global variable or field will need to retain its known ownership.

4.2.5 Object Allocators

The generated code keeps track of how native objects have been obtained, and restricts their operations accordingly. Three possible allocators are recognized:

DYNAMIC	Such objects are created either through bound constructors, accessed from native access types, or through an automatic copy from the bound code. The ownership of dynamic objects can be changed.
STATIC	Such objects are accessed from global variables, fields, or callback parameters that are not access types. The ownership of these objects is always <code>NATIVE</code> and cannot be changed.
UNKNOWN	In some cases, for example on copying values of tagged types, it is not possible to determine whether the object has been allocated statically or dynamically. In such cases you can access the allocator through which an object is referenced using the <code>IProxy.getAllocator()</code> function.

4.2.6 Restrictions on Proxy-Owned Objects Passed to Subprograms

Passing a (reference to a) native object as an actual parameter to an Ada subprogram whose corresponding formal parameter is either an access parameter or of an access type could in some cases result in a dangling reference to the native object, whereas in other situations it might be harmless. In order to provide the desired safety while allowing the needed generality, `ada2java`'s approach is similar to the way that the Ada standard supplies both the `'Access` and `'Unchecked_Access` attributes for composing pointers to data objects.

In summary, when the Java side is responsible for pointer memory management (i.e. the object is owned or it is accessed through `'Access` or `'Unrestricted_Access`), then an exception will be thrown in Java on an attempt to pass a reference to a native object as an actual parameter when the formal is of a named or anonymous access type. However, this check can be suppressed either locally (for a given formal parameter) or globally (for all calls).

Here is a summary of the troublesome scenario:

1. A reference to a proxy object is passed as a parameter to a Java method, and the reference to the corresponding native object is then passed to a native Ada subprogram whose formal is an access parameter or of an access type;
2. The Ada subprogram copies the parameter to a global variable and ultimately returns;

3. The proxy object ultimately becomes inaccessible and is garbage collected by the JVM;
4. The native object is deallocated as an effect of the proxy object's finalization;
5. The deallocated object is still designated by the global variable.

To prevent this, an exception is thrown at step 1 above.

Here is an example:

```
-- Original Ada package
package P is
  type T is
    record
      A, B : Integer;
    end record;

  type A_T is access all T;

  G : A_T;

  procedure Set_G (V : A_T);
end P;

package body P is
  procedure Set_G (V : A_T) is
  begin
    G := V;
  end Set_G;
end P;

// Java statements
T v = new T ();
P_Package.Set_G (v); // Throws an exception
```

If the invocation of `Set_G` did not throw an exception, `G` would become a dangling pointer when the proxy object referenced by `v` is garbage collected.

However, if the invocation of `Set_G` did not cause an assignment of its formal parameter `v` to a global variable, then passing `v` as an actual parameter would be harmless. `ada2java` allows the Ada API to enable such uses, and thus to deactivate the check:

```
pragma Annotate (AJIS, Assume_Escaped, Condition, Subprog, Formal_Param)
```

where:

- *Condition* is either `True` or `False`
- *Subprog* is the name of the Ada subprogram
- *Formal_Param* is the affected formal parameter, expressed as a `String` literal

When *Condition* is `False`, the Ada subprogram is responsible for ensuring that the formal parameter is not copied to a global variable, or used to set a field

of a structure that will be used outside of the scope of the subprogram. When *Condition* is `True`, an exception is thrown when a proxy object is passed as an actual parameter to the *Formal_Param* parameter of *Subprog*.

Here is an example:

```
package P is
  type T is
    record
      A, B : Integer;
    end record;

  type A_T is access all T;

  G : A_T;

  procedure Safe_Set_G (V : A_T);
  pragma Annotate (AJIS, Assume_Escaped, False, Safe_Set_G, "V");
end P;

package body P is
  procedure Safe_Set_G (V : A_T) is
  begin
    G := new T'(V.all);
  end Safe_Set_G;
end P;

// Java statements
T v = new T ();
P_Package.Safe_Set_G (v); // OK
```

The invocation `Safe_Set_G (v)` is safe since the Ada subprogram does not let the formal parameter “escape”.

The “escaped” checks can be globally activated through the `ada2java` switch ‘`--assume-escaped`’, and globally deactivated through the switch ‘`--no-assume-escaped`’. The default is ‘`--assume-escaped`’. An explicit `pragma` overrides the global configuration switch.

Note that canceling escape checks should be done with great care, as there is no way to ensure that no escaping has occurred. In many situations, the programmer’s intent is to create an object on the Java side and then store the object on the native side. In such a case, canceling proxy ownership through the `IProxy.setOwner` method will have the desired effect, for example:

```
// Java statements
T v = new T ();
v.setOwner (Owner.NATIVE);
P_Package.Set_G (v); // OK
```

4.3 Thread Safety

By default, the generated Java code is thread-safe, with locking logic that prevents multiple Java threads from accessing native methods at the same time. In summary, a Java method in a proxy object acquires a lock (a binary semaphore) before invoking the corresponding native method, and releases the lock when the native method returns (either normally or abnormally). The semaphore is global versus per Ada package; e.g. different Java threads are not allowed to invoke native methods simultaneously even if the native methods correspond to subprograms from different packages.

Even if the Java application does not explicitly create any threads, there are still two threads – the main (user) thread and the garbage collector; thus locking is needed in this case also.

The default locking behavior is not always appropriate, however. In particular, if the native code is using tasking, with mutual exclusion enforced on the Ada side, then there is no need for locking in the Java code (and indeed such locking could have undesirable consequences including deadlock).

User control over the generation of locking code is obtained through the following pragma:

```
pragma Annotate (AJIS, Locking, Subprogram, LockControl);
```

where *Subprogram* is the subprogram name, and *LockControl* is one of `Disable`, `Check` and `Protect`.

- | | |
|---------|---|
| Protect | Default setting. The generated code automatically brackets each invocation of the named native <i>Subprogram</i> within a <code>lock ... unlock</code> region. If a thread <code>t1</code> invokes <i>Subprogram</i> while some other thread <code>t2</code> holds the lock, then <code>t1</code> will be queued until the lock is released. Thus simultaneous calls of native methods are permitted but will entail queuing. |
| Disable | The generated code contains no locking around invocations of <i>Subprogram</i> , and one thread may invoke <i>Subprogram</i> while some other thread is executing a native method (even one whose <code>LockControl</code> is <code>Protect</code>). |
| Check | The generated code contains no locking around invocations of <i>Subprogram</i> , and the application must ensure that any such invocation is within a <code>lock ... unlock</code> region. An exception is thrown if a thread invokes <i>Subprogram</i> without holding the lock. |

The following examples illustrate the multithreading behavior and the effects of the *LockControl* argument to the pragma.

```
package P is
  procedure P1 (I : Integer);
  pragma Annotate (AJIS, Locking, P, Disable);

  procedure P2 (I : Integer);
  pragma Annotate (AJIS, Locking, P, Disable);
end P;
```

Concurrent invocations of P1 and P2 are allowed (i.e., locking code is not generated automatically).

Code protected by the lock can be provided by hand by the Java developer. The lock is created in the library class generated for the binding, *base_package.Ada2Java.Library*, under the identifier *lock*. Thus:

```
import base.P_Package;
import base.Ada2Java.Library;

public class Main {
  public static void main (String [] args) {
    Library.lock.lock ();

    try {
      P_Package.P1 (0);
      P_Package.P2 (0);
    } finally {
      Library.lock.unlock ();
    }
  }
}
```

Following standard Java coding style, the *lock/unlock* logic should always appear in a *try ... finally* block, in order to release the lock even if an unexpected exception is propagated.

In certain cases, locking is required but the logic is more complex than simply protecting each native method invocation with a semaphore. For example, it may be necessary to invoke a sequence of native methods as an atomic action. This effect can be achieved through the *Check* setting for *LockControl*:

```
package P is
  procedure P1 (I : Integer);
  pragma Annotate (AJIS, Locking, P1, Check);

  procedure P2 (I : Integer);
  pragma Annotate (AJIS, Locking, P2, Check);
end P;
```

The Java program has to acquire the library lock before attempting any native call. Invoking P1 or P2 outside of a section protected by the lock will throw a Java exception.

The locking behavior can be changed globally through the `ada2java` ‘`--[no-]locking`’ switch. More specifically, here are the permitted values for this switch:

```
--locking-protect
    Default setting; globally sets LockControl as Protect for all subprograms.

--locking-check
    Globally sets LockControl as Check for all subprograms.

--no-locking
    Globally disables locking (i.e., sets LockControl as Disable for all subprograms).
```

Note that an AJIS `Locking` pragma takes precedence over the global switch.

The `finalize` method invoked during garbage collection does not correspond to a subprogram from the Ada package that is input to `ada2java`, and it is not affected by the ‘`--[no-]locking`’ switch. Instead, the locking logic used for the `finalize` call of Java proxies during garbage collection is determined by the ‘`--[no-]locking-finalize`’ switch:

```
--locking-finalize-protect
    Default setting; sets LockControl as Protect for finalize.

--locking-finalize-check
    Sets LockControl as Check for finalize.

--no-locking-finalize
    Disables locking (i.e., sets LockControl as Disable for finalize)
```

A typical usage of these two switches would be to set `Check` as the *LockControl* for all subprograms, and `Protect` as the *LockControl* for `finalize` methods:

```
ada2java --locking-check --locking-finalize-protect p.ads
```

4.4 Proxies and Native Object Equality

Proxy classes are generated with an `equals` implementation that calls the corresponding Ada “=” operation. For example:

```
package P is
  type T is
    record
      F : Integer;
    end record;
end P;
```

The proxy class can be used as follows:

```
T v1 = new T ();
T v2 = new T ();
v1.F (0);  "setter" method for field F
v2.F (0);  "setter" method for field F
```

and now the result of `v1.equals (v2)` is `true`. This corresponds to “shallow” equality, in contrast with “`==`” which tests pointer identity.

A new proxy is created each time a native function returns a pointer. For example:

```
package P is
  type T is null record;
  type A_T is access all T;
  G : A_T := new T;
  function Return_G return A_T;
end P;

package body P is
  function Return_G return A_T is
  begin
    return G;
  end Return_G;
end P;

AT p1 = P_Package.Return_G ();
AT p2 = P_Package.Return_G ();
```

Now `p1 == p2` is `false`, since a new proxy is created by each function return, but `p1.equals (p2)` is `true`,

The association between a proxy and its native object is lost when the proxy is passed to a native method. For example:

```
package P is
  type T is null record;
  type A_T is access all T;
  function Return_This (This : A_T) return A_T;
end P;

package body P is
  function Return_This (This : A_T) return A_T is
  begin
    return This;
  end Return_G;
end P;

T v = new T ();
```

Now the result of `(v == P_Package.Return_This (v))` is `false`

Java reference equality has special semantics in the case of cross-language inheritance, due to the use of a shadow native object – see [Section 4.6.3.3 \[Shadow Object Equality\]](#), page 42 for more details.

4.5 Clone and Copy Semantics

A proxy class generated from a non-limited Ada type includes a `clone` method. The base class for all proxies, `AdaProxy`, implements the `Cloneable` interface, and defines a public method `clone`. The `clone` method performs a “shallow copy” of all the fields, except for the native object reference. The native object reference in the cloned proxy is a pointer to a newly allocated Ada object. This new native object is itself a shallow copy of the original native object.

Thus cloning a proxy does not result in the sharing of the native object by the original proxy and the cloned proxy. The latter points to a new native object. This behavior is needed to avoid a dangling reference (to the native object) when the original proxy is garbage collected.

If the proxy class corresponds to a limited type, then the generated clone method will throw an exception.

Additional semantics for clone and copy, in connection with cross-language inheritance, are covered below (see [Section 4.6.3.4 \[Shadow Object Cloning\]](#), page 43).

4.6 Cross-Language Inheritance

This section discusses a number of issues related to cross-language inheritance; i.e., defining a Java class as an extension of a proxy class for an Ada tagged type.

4.6.1 Inheriting from a Java Proxy

As explained in [Section 3.1.5 \[Tagged Types\]](#), page 19, `ada2java` maps an Ada tagged type to a non-final Java proxy class. You can extend this class in Java. For example:

```
package P is
  type T is tagged null record;

  procedure Prim (V : T);
end P;
```

results in the following Java class:

```
public class T {
  public void Prim () {
    ...
  }
  ...
}
```

which may be extended, with the instance method overridden:

```
public class T_Child extends T {
    public void Prim () {
        ...
    }
}
```

An object of the subclass `T_Child` is also a proxy; constructing such an object allocates a native object, referred to as a “shadow native object”. Its properties will be described below ([Section 4.6.3 \[The Shadow Native Object\], page 41](#)).

The Java program can then invoke the `Prim` method, with standard Java dispatching behavior. For example:

```
T v1 = new T ();
T v2 = new T_Child ();

v1.Prim (); // Will call the native Prim
v2.Prim (); // Will call the overridden Java Prim
```

4.6.2 Cross Language Dispatching from Ada

A method overridden in Java can be called in Ada using the usual Ada dispatching mechanism. For example:

```
package P is
    type T is tagged null record;

    procedure Prim (V : T);

    procedure Call_Prim (V : T'Class);
end P;

package body P is

    procedure Prim (V : T) is
    begin
        -- Native Prim implementation;
    end Prim;

    procedure Call_Prim (V : T'Class) is
    begin
        Prim (V); -- Dispatching call.
    end Call_Prim;

end P;
```

The invocation of `Prim` in `Call_Prim` is dispatching, and the `Prim` for the type of the actual parameter will be called. In Java, this procedure can be used in conjunction with a cross-language extension of `T`, e.g.:

```

class T_Child extends T {
  public Prim () {
    ...
  }
}

T v = new T_Child ();
P_Package.Call_Prim (v);

```

The Java program is invoking the native Ada `Call_Prim` procedure, which in turn dispatches to the Java method `Prim` in `T_Child`.

4.6.3 The Shadow Native Object

This section describes in more detail the semantics of the shadow native object.

4.6.3.1 Basic Properties

As seen above, cross-language dispatching is supported; an Ada dispatching call may result in the invocation of a Java method on a proxy object. This is possible because of the *shadow native object* concept.

For any tagged type `T` declared in a package spec that is input to `ada2java`, a new type is automatically generated that extends `T` with a component that references a proxy object. The Ada declaration is:

```

type Shadow_T is new T with
  record
    Link_To_Proxy : Java_Object;
  end record;

```

If `T_Child` is a Java class that extends `T`, then constructing an instance of `T_Child` will create a native `Shadow_T` object instead of a regular `T` object. This type overrides every controlling primitive of `T`, and delegates the dispatching to the Java side. If a method corresponding to an Ada primitive operation is not overridden in Java, then the subprogram from the parent Ada type will be automatically called.

The use of a shadow object introduces a tight relationship between a Java proxy and its Ada native object. Basically, there is a roundtrip dependency between the two, so that a Java proxy object is associated with a unique Ada shadow native object and vice versa.

This has several non-trivial implications that are described in the sections below.

4.6.3.2 Memory Management

The reference from the shadow native object to the Java object is called a *global reference*.² There are two kinds of global references: regular (usually simply referred to as global references) and weak. A regular global reference prevents the object from being garbage collected, whereas a weak global reference does not.

When a native object is owned by its proxy, the proxy is responsible for releasing the native memory. In this case, the reference from the native object to the proxy is a weak global reference: garbage collection will not be prevented, and both the proxy and then the native object will be released upon collection.

However, when the native object is owned by the native side, then the native object may continue to exist even if the proxy has become inaccessible from Java. In such a case, a global reference is used, so that the garbage collector is prevented from collecting the Java object. Such a global reference will be automatically released when the Ada object is actually deallocated.

Forcing the deallocation of the native object through `IProxy.deallocateNativeObject` will release the global reference as well.

Switching the owner of the native object between NATIVE and PROXY will switch the reference from a regular global reference to a weak one.

In order to avoid a potential memory leak, the link between the shadow native object and the Java proxy has to be broken manually when the shadow native object is no longer needed on the Ada side. This may be done in two ways:

- Deallocating the native shadow object, from Ada.
- Invoking the `unlink` method on the proxy, from Java

4.6.3.3 Shadow Object Equality

Java reference equality (`==`) is not consistent with native pointer equality, but Java object equality (`equals`) is. That is, if `A` and `B` are two native pointers where `A = B`, then on the Java side `A.equals(B)` is `true` but `a==b` is `false`.

However, since the association between shadow object and proxy is one-to-one, proxy equality is meaningful. For example:

² An explanation of global references may be found in Liang's *The Java Native Interface Programmer's Guide and Specification*, Chapter 5.

```

package P is
  type T is tagged null record;
  type T_Access is access all T'Class

  procedure Identity (Item : T_Access) return T_Access;
  -- Returns Item as its result
end P;

```

On return from `Identity`, the “glue code” between Java and Ada will check if the returned value is a shadow native object and, if so, will return the corresponding Java proxy instead of creating a new one. Hence the following fragment

```

public class T_Child extends T {
}

```

```

T v = new T_Child ();

```

will result in the expression `P_Package.Return_This (v) == v` delivering `true`.

4.6.3.4 Shadow Object Cloning

Cloning from Java will result in copying a native object. In the case of a shadow native object, there will be a new shadow object as well, referencing the newly created java proxy, thus preserving the one-to-one relationship between the shadow native object and the proxy.

However, there are cases where copies are made from Ada as well:

```

package P is
  type T is tagged null record;
  type T_Access is access all T'Class

  procedure Duplicate (Item : T_Access) return T_Access;
end P;

package body P is

  procedure Duplicate (Item : T_Access) return T_Access is
  begin
    First_Copy : T'Class := Item.all;
    Second_Copy : T_Access := new T'Class'(Item.all);
  begin
    return Second_Copy;
  end Duplicate;

end P;

```

Calling `Duplicate` from Java will lead to two shadow object copies, one on the stack (`First_Copy`) and one on the heap (`Second_Copy.all`). As explained earlier, the link between the shadow native object and the proxy will be deleted when the Ada object is deallocated, at the exit of `Duplicate` for the `First_Copy` variable.

The `Second_Copy` proxy created in the copy process will be returned by the Java method corresponding to the `Duplicate` procedure.

Note that a proxy created by such a copy does not own its Ada native object.

Proxy cloning from Ada does not involve a `clone` invocation from the Java extended object. If the Java code needs to perform a deep copy in the Java proxy, the method `void proxyCloned (IProxy initialObject)` should be overridden instead. The default `clone` implementation of the generated Java classes will call this method as well.

Limitation: when `proxyCloned` is called from Ada, the link between proxy and native object is not yet established. Thus `proxyCloned` is not allowed to invoke any native methods. The Ada type may be derived from `Finalization.Controlled`, and `Adjust` may be overridden, to work around this limitation.

Note that multiple copies – and thus repeated proxy creation – may be involved when a (non-access) shadow object is returned. For example:

```
package body P is

  function Identity (Item : T_Access) return T'Class is
  begin
    return Item.all;
  end Identity;

end P;
```

Due to internal machinery, three copies may be needed in the implementation of the return from this function when called from Java. If it is necessary to ensure that only one copy is performed, then the Ada function should be written as a wrapper for an Ada access-returning function.

4.6.4 Controlled Types

A class that is generated from an Ada controlled type may be extended in Java, with overriding versions of the `Adjust`, `Finalize`, and/or `Initialize` methods.

A current limitation is that `Initialize` cannot be called on a shadow object, e.g.:

```
package P is

  type T is new Controlled with null record;

  procedure Initialize (This : in out T);

end P;
```

can be used in:

```

class T_Child extends T {
    void Initialize () {
    }
}

T v = new T_Child ();

```

But the overridden `Initialize` will not be called by the constructor. This limitation will be removed in a future release.

4.7 Managing Attachment to Java Proxies

By default, subprograms (except controlling primitives) are not attached – they are placed in the default class `Ada_Package_Name._Package`. However, when all of the following conditions are met, a subprogram is attachable to the class corresponding to its first parameter:

- The first parameter of the subprogram has one of the following forms:
 - A private type or a record type,
 - An access type (of mode `in`) designating a private or record type, or
 - An access parameter designating a private or record type.
- The type of this parameter – or of the designated type if an access parameter – is declared in the same package spec as the subprogram.

A subprogram that meets these criteria can be mapped to a method defined in the class corresponding to the first parameter’s type, and the value of this first parameter will come from the hidden parameter `this`.

This attachment is activated by the annotation pragma `Attached`:

```
pragma Annotate (AJIS, Attached, Condition, Subprogram);
```

Example:

```

package Example is
  type Rec is null record;
  procedure Proc (V : Rec; I : Integer);
  pragma Annotate (AJIS, Attached, True, Proc);
end Example;

```

will map to the class

```

public class Rec extends com.adacore.ajis.internal.ada.AdaProxy {
    ...
    public void Proc (int I){...}
}

```

Attachment policies may be globally turned on / off using the following switches:

- [no-]attach-parameter
Activates or deactivates “best-effort” attachment to the class corresponding to the first parameter. When this is set, `ada2java` will try to perform subprogram attachment whenever possible. Default is `--no-attach-parameter`
- [no-]attach-access
Activates or deactivates attachment for access type. When activated, subprograms with an access type (named or anonymous) for their first parameter will be attached. (Note, however, that such attachment prevents passing a `null` value, since this is always the implicit parameter.) Default is `--no-attach-access`.
- [no-]attach-controlling
Activates or deactivates attachment of controlling primitives. This is required for cross language inheritance. Default is `--attach-controlling`.
- [no-]attach-ada2005
Activates or deactivates attachment based on applicability of Ada 2005 prefix notation. With the ‘`--attach-ada2005`’ switch, `ada2java` will attempt to attach a subprogram (define it in the class corresponding to the initial parameter) that would otherwise be placed in the default package, if it can be invoked via Ada 2005 prefix notation. Default is `--no-attach-ada2005`.

In the example below, attachment is requested for everything except noncontrolling initial access parameters:

```
ada2java --attach-parameter --no-attach-access p.ads
```

Pragma Attached takes precedence over the global switch.

4.8 Exceptions propagation

Exceptions raised from Ada are translated into instances of the relevant descendant of class `com.adacore.ajis.NativeException` (see [Section 3.4 \[Exceptions\], page 23](#)) and propagated to Java. Exceptions raised from a Java callback are translated back to the original Ada exception - or to `Java_Exception` declared in the `AJIS.Java` package, and propagated to Ada.

5 Using javastub to Generate Ada Package Specifications

The `javastub` utility program generates an Ada package specification from a Java class file that has native methods.

It is invoked as follows

```
$ javastub [filename.class] {filename.class}
```

where each '`filename.class`' is a class file for a Java class that has native methods.

For each class file argument, `javastub` generates an Ada package specification with subprograms corresponding to the native methods. The name of the generated file is '`filename_jni.ads`'.

The `javastub` utility is the Ada analog to the `javah -jni` command in the Java Development Kit, which takes a class file as input and produces a C header file with the correct function prototypes.

You can use `javastub` if you intend to use Ada, rather than C, to implement the native methods in a Java class. You will then be responsible for doing the necessary JNI programming, using the Ada binding to the C JNI library.

As an example, here is a simple Java class with a native method:

```
// Foo.java
class Foo{
    native void bar(int i);
}
```

Compile the Java source file and then invoke `javastub` on the class file:

```
$ javac Foo.java
$ javastub Foo.class
```

The following file '`foo_jni.ads`' is generated by `javastub`:

```
-- Stub for class Foo
with Interfaces.Java.JNI; use Interfaces.Java.JNI;
package Foo_JNI is
  -- Class:   Foo
  -- Method:  bar
  -- Signature: (I)V
  procedure bar (Env : JNI_Env_Access; Obj : J_Object; P1 : J_Int);
  pragma Export (C, bar, "Java_Foo_bar__I");
end Foo_JNI;
```


6 Using `jvm2ada` to Interface Ada with Java

The `jvm2ada` tool is under development and is not supported in the current GNAT-AJIS release.

Appendix A Using JNI Directly

This Appendix is currently under development and is not available in the GNAT-AJIS User's Guide for the current release.

Index

-
- '--assume-escaped' option (for ada2java) 11, 34
 - '--attach' option (for ada2java) 11
 - '--attach-access' option (for ada2java) .. 46
 - '--attach-ada2005' option (for ada2java) 46
 - '--attach-controlling' option (for ada2java) 46
 - '--attach-parameter' option (for ada2java) 46
 - '--java-enum' option (for ada2java) 11
 - '--link-method' option (for ada2java) 10
 - '--main-class' option (for ada2java) 10
 - '--monitor' option (for ada2java) 10
 - '--monitor-finalize' option (for ada2java) 10
 - '--no-assume-escaped' option (for ada2java) 11, 34
 - '--no-attach' option (for ada2java) 11
 - '--no-attach-access' option (for ada2java) 46
 - '--no-attach-ada2005' option (for ada2java) 46
 - '--no-attach-controlling' option (for ada2java) 46
 - '--no-attach-parameter' option (for ada2java) 46
 - '--no-java-enum' option (for ada2java) 11
 - '--no-monitor' option (for ada2java) 10
 - '--no-monitor-finalize' option (for ada2java) 10
 - '-b' option (for ada2java) 9
 - '-c' option (for ada2java) 9
 - '-fno-strict-aliasing' option (for gcc) .. 11
 - '-fPIC' option (for gcc) 11
 - '-fstack-check' option (for gcc) 11
 - '-h' option (for ada2java) 9
 - '-L' option (for ada2java) 10
 - '-M' option (for ada2java) 10, 12
 - '-o' option (for ada2java) 9
 - '-O' option (for gcc) 11
 - '-P' option (for ada2java) 10
- ## A
- ADA_PROJECT_PATH environment variable 5
 - ada2java 3
 - ada2java command 9
 - 'Ada2Java' directory 7
 - AdaProxy class 39
 - AJIS.Annotations package 14
 - Array types (mapping to Java) 16
 - ASIS 4
 - Assume_Escaped (argument to pragma Annotate) 33
 - Attachment (of entities to a class) 15, 22
- ## C
- Class-wide parameters (mapping to Java) .. 19
 - CLASSPATH environment variable 5
 - Clone and copy semantics 39
 - clone method (in AdaProxy class) 39
 - com.adacore.ajis.IProxy.Owner type 31
 - com.adacore.ajis.NativeException 46
 - Compatibility (of GNAT-AJIS and GNAT) .. 4
 - Constraint checks 16
 - Controlled types 44
 - Current limitations 25
- ## D
- Default class 15, 21, 45
- ## E
- Enumeration types (mapping to Java) 16
 - Equality 37
 - Exception propagation 46
- ## G
- getOwner method 31
 - Global reference 42
 - Global variables (mapping to Java) 21
 - gprbuild usage 7

I

Inheritance (cross-language).....	39
Installation of GNAT-AJIS.....	4
<code>IProxy.getAllocator</code> function.....	32

J

<code>javastub</code>	3
<code>javastub</code> command.....	47
JGNAT.....	3
<code>jvm2ada</code>	3
<code>jvm2ada</code> command.....	49

L

<code>LD_LIBRARY_PATH</code> environment variable.....	5
' <code>Library.java</code> ' file.....	7
<code>Locking</code> (argument to pragma <code>Annotate</code>)..	35
' <code>locking</code> ' option (for <code>ada2java</code>).....	37
' <code>locking-finalize</code> ' option (for <code>ada2java</code>)..	37

M

Memory management.....	42
Memory model.....	28

N

Name clashes.....	27
Native object equality.....	37
' <code>no-locking</code> ' option (for <code>ada2java</code>).....	37
' <code>no-locking-finalize</code> ' option (for <code>ada2java</code>)	37

O

Ownership (of native objects).....	30
------------------------------------	----

P

Parameters (mapping to Java).....	22
<code>PATH</code> environment variable.....	5
pragma <code>Annotate</code>	14, 27
Predefined environment (mapping to Java)	24
Proxy object.....	28, 31

R

Record types (mapping to Java).....	17
Regular global reference.....	42
Renamings (mapping to Java).....	24

S

SAL (Stand-Alone Library) project.....	12
Scalar types (mapping to Java).....	15
Shadow native object.....	41
Shared libraries.....	12
Strings (mapping to Java).....	17
Subprograms (mapping to Java).....	21

T

Tagged types (mapping to Java).....	19
Thread safety.....	35

U

UTF-16 encoding.....	17
UTF-8 encoding.....	17

W

Weak global reference.....	42
----------------------------	----

Table of Contents

About This Guide	1
What This Guide Contains	1
What You Should Know before Reading This Guide	1
1 Getting Started with GNAT-AJIS	3
1.1 Introduction.....	3
1.2 GNAT-AJIS Installation Structure.....	3
1.3 GNAT-AJIS / GNAT Compatibility.....	4
1.4 A Simple Example: Calling Ada from Java.....	4
1.4.1 Environment Setup	5
1.4.2 An Ada Package.....	6
1.4.3 Invoking <code>ada2java</code>	6
1.4.4 Compiling the Java class.....	7
1.4.5 Building the Application	7
1.4.6 Running the Program.....	8
2 Using <code>ada2java</code> to Generate Java Classes	9
2.1 Using the Tool.....	9
2.2 Compiling and Running the Generated Code	11
2.2.1 Issues with the Ada Generated Code.....	11
2.2.2 Compiling as an Ada Shared Library	12
2.2.3 Compiling as an Ada Main Subprogram	12
2.2.4 Compiling the Java Generated Classes	13
2.3 Pragma Annotate and <code>ada2java</code>	13
3 Mapping Ada to Java	15
3.1 Types	15
3.1.1 Scalar Types	15
3.1.2 Arrays	16
3.1.3 Strings.....	17
3.1.4 Simple Record Types.....	17
3.1.5 Tagged Types	19
3.1.5.1 General principles	19
3.1.5.2 Ada type hierarchies.....	20
3.1.5.3 Java class hierarchies.....	20
3.2 Global Variables and Constants.....	21
3.3 Subprograms.....	21
3.3.1 Method placement	21

3.3.2	Subprogram parameters	22
3.4	Exceptions	23
3.5	Renamings	24
3.6	Predefined Environment	24
3.7	Current Limitations	24
4	Advanced ada2java Topics	27
4.1	Dealing with Name Clashes	27
4.2	Memory Model	28
4.2.1	Requirements for Non-null Parameter Values	28
4.2.2	Allocating Ada Objects from Java	29
4.2.3	Automatic Creation of Native Objects	29
4.2.4	Native Ownership	30
4.2.5	Object Allocators	32
4.2.6	Restrictions on Proxy-Owned Objects Passed to Subprograms	32
4.3	Thread Safety	34
4.4	Proxies and Native Object Equality	37
4.5	Clone and Copy Semantics	39
4.6	Cross-Language Inheritance	39
4.6.1	Inheriting from a Java Proxy	39
4.6.2	Cross Language Dispatching from Ada	40
4.6.3	The Shadow Native Object	41
4.6.3.1	Basic Properties	41
4.6.3.2	Memory Management	41
4.6.3.3	Shadow Object Equality	42
4.6.3.4	Shadow Object Cloning	43
4.6.4	Controlled Types	44
4.7	Managing Attachment to Java Proxies	45
4.8	Exceptions propagation	46
5	Using javastub to Generate Ada Package Specifications	47
6	Using jvm2ada to Interface Ada with Java	49
Appendix A	Using JNI Directly	51
Index		53