

Laboratoire d'Architecture des Processeurs - EPFL
Semester Project - Transparent PLD use from Java
Final Report

Work: Christophe Dubach & Alok Menghrajani

Report: Alok Menghrajani

Assistant: Miljan Vuletic

Professor: Paolo Ienne

February 2004

Special thanks to the former students Cédric Gaudin, Stéphane Magnenat, Julien Pilet and Damien Baumann for their continuous support. To Miljan Vuletic and all the LAP professors without whom nothing would have happened.

Contents

1	About This Document	1
2	Introduction	1
2.1	Embedded Systems	1
2.2	RokEPXA	2
2.3	Past And Present Work	2
3	Project Description & Objectives	3
4	Development Environment	4
4.1	Linux Kernel 2.6.0	4
4.2	NFS Support	4
4.3	Native-compiler	4
4.4	Kaffe JVM	4
5	Research	5
5.1	Inside Java And It's Virtual Machine	5
5.1.1	Class Loader	5
5.1.2	Bytecode	6
5.1.3	BCEL	8
5.1.4	Java Native Interface	8
5.1.5	Garbage Collection	8
5.1.6	Java Hot Spot	9
5.1.7	Security Issues	9
5.2	IDEA	9
5.3	Generating VHDL From Java	9
5.4	Different Solutions To My Project	10
6	Implementation	13
6.1	XML Config File	13
6.2	Interfacing With JNI	13
6.3	Launcher Class	15
6.4	Custom Class Loader	16
7	Limitations / Improvements	17
8	Conclusions and Future work	17
9	Appendix	19
9.1	Code	19
9.1.1	Organization	19
9.1.2	Compiling & Running	19
9.1.3	VMWXMLParser.jj	20

9.1.4	VMWClassLoader.java	22
9.1.5	VMWLauncher.java	28
9.1.6	VMWConstants.java	29
9.1.7	ClassIdMethodSignature.java	30
9.1.8	IdMethodSignature.java	30
9.1.9	VMWSkel.java	31
9.1.10	VMWcfg.xml	31
9.1.11	VMWXMLcfg.dtd	32
9.1.12	IdeaPlus.java	33
9.1.13	IdeaTestSw.java	36
9.2	Bibliography	38

1 About This Document

I am a student at EPFL (Swiss Institute of Technology at Lausanne), Switzerland. I am doing a Master in Computer Science. This document is the report of my final year's semester project: "Transparent PLD use from Java". I worked at the Laboratoire d'Architecture des Processors.

Christophe Dubach was assigned a similar project ("Java virtual machine on FPGA-based platforms"), we therefore collaborated together and this is the result of our combined work.

The CD-ROM included with this report contains the source code of all our work. It also includes all the tools and libraries needed for development. Installation instructions are included in our report entitled: "RokEPXA Development Environment From Scratch" which is also being submitted simultaneously.

All the tools and libraries we used are available as part of the various open source projects, which are constantly evolving and being updated. We recommend using exactly the versions which we have used. We had to correct (patch) some of these tools. Newer version might not include these corrections.

All our work is being released under the GPL license, which means you can copy, read, use, modify our work, but you **MUST** keep it under the GPL license (which means you must also share any work based on our's). Christophe and I both strongly believe in open source development, we hope our contribution to the community will be helpful.

Part of my work is based on software developed by the Apache Software Foundation (<http://www.apache.org/>).

2 Introduction

2.1 Embedded Systems

The trend in the current society is to have smart chips in everyday objects. These chips are called embedded systems because they control one or more functions of the equipment. They are typically required to meet very different requirements than a general-purpose personal computer (they usually have small amount of memory, are cheap to produce and must consume low power).

These chips can be categorized in two types:

1. Processors (which can perform general tasks)
2. Specialized chips designed for specific tasks

Designing specialized chips is complex and expensive, but it is often

the only way to meet real-time constraints (eg: a dsp¹ chip or 3d graphics card can outperform a fast processor because it can be designed to perform multiple calculations simultaneously).

For research and development applications, there is a special kind of chip called FPGA² which can be programmed at the circuit level to perform a specific task. The programming is done in the same way as for designing a chip; that's is in a hardware description language (eg: VHDL³). The flexibility offered by FPGA chips comes at the cost of extra programming effort.

2.2 RokEPXA

<http://lapwww.epfl.ch/dev/arm/>

We used the RokEPXA board designed at the Laboratoire d'Architecture des Processeurs, EPFL. This board uses Altera's EPXA (*<http://www.altera.com/products/devices/arm/arm-index.html>*) integrated circuit which is a hybrid architecture. It combines a processor (ARM922T) with a FPGA.

The typical usage scenario for the RokEPXA is for robots; the FPGA controls the fundamental functions (sending signals to the motors, receiving signals from sensors, etc.) and the processor handles the higher level behaviour (artificial intelligence, navigation, etc.).

The FPGA can also be used to connect a camera or ethernet interface.

The FPGA on the RokEPXA can be programmed by the processor (by loading a `.map`⁴ design file that is obtained by compiling the VHDL code). This means that the processor and FPGA can share "objects". The processor can perform general computing while the FPGA is programmed to compute time critical sections (coprocessor).

One of the difficulties is communication between the processor and FPGA chip. Intensive work on this has already been done by Miljan Vuletic. He has designed a solution called VMW⁵ (formerly known as VIM), that allows the coprocessor to access the virtual memory of a user-space application.

2.3 Past And Present Work

The LAP has been working on the RokEPXA for a couple of years. When we began our project, we had a working RokEPXA card, with a minimal development environment (that we abandoned). We had the design files for IDEA (we'll get back to this later) and network support.

VMW was also in it's completion stage.

¹digital signal processing

²Field-Programmable Gate Array

³VHSIC (Very-High-Speed Integrated Circuit) Hardware Description Language

⁴In our specific environment, the `.map` files end with the `.sbi` extension

⁵Virtual Memory Window

We used a lot of open source projects such as the Linux Kernel, the Arm Linux project, etc. (we will introduce these projects and others later in this document). Some of these projects lacked testing on our specific hardware and lead us to spend a lot of time debugging.

3 Project Description & Objectives

We set ourselves the following objectives:

1. **Creating a complete development environment for the RokeEPXA.**

Christophe Dubach & Alok Menghrajani

Although there already was a development environment, we ran into trouble because it was incomplete, lacked documentation and had various bugs. We created a fresh environment that includes: linux 2.6.0, nfs support, a cross-compiler⁶, a native compiler, busybox system tools, a JVM⁷, dynamically linked library support and floating-point emulation.

2. **Porting IDEA⁸ to Java**

Christophe Dubach & Alok Menghrajani

IDEA is an encryption algorithm. We already had a C and VHDL version, so we ported it to Java in order to test our work and measure performances.

3. **Using the coprocessor from within Java.**

Christophe Dubach

Given a .map file, we want to run it from within Java. This way we can write most of the code in Java and only the time critical part in VHDL. Christophe worked on this, his solution is written in C and called from Java using JNI⁹.

4. **Transparent PLD¹⁰ use from Java.**

Alok Menghrajani

Given a Java method and a .map file that performs exactly the same calculation, we want the JVM to run the .map file instead of the Java version of the code. A PLD is the same as a FPGA.

“transparent” means that we don’t want to have the user modify the Java source code. We ignored the aspect of how to generate the .map file from the Java (we will talk about this in the conclusion of this report).

⁶A compiler that runs on our desktop environment but produces code for the ARM

⁷Java Virtual Machine

⁸International Data Encryption Algorithm

⁹Java Native Interface

¹⁰Programmable Logic Device

4 Development Environment

I will keep the discussion about the development environment to the bare minimum. You can refer to our other report “RokEPXA Development Environment From Scratch”, submitted simultaneously.

Here is the list of software we used with a little explanation. We unfortunately didn't find any complete distribution that met our needs, we therefore compiled everything on our own.

4.1 Linux Kernel 2.6.0

<http://www.kernel.org/>

We decided to switch to Linux 2.6.0 (the former development environment was based on Linux Kernel 2.4.19) for stability reasons. We applied the rmk1 ARM Linux patch (<http://www.arm.linux.org.uk/>) and then another patch called rokepxa (based on the former rokepxa patch by Cédric), that fixes problems with the ARM Linux patch and adds support for the specific features of RokEPXA.

We activated floating-point emulation at the kernel level, because when we tried to do it at the compiler level we ran into trouble. This method is also more optimal (saves code) and will scale better if we someday end up having a floating-point coprocessor in VHDL.

4.2 NFS Support

By having NFS support, we are able to transfer data between the RokEPXA and the desktop development machine at a much faster rate than with the parallel port. We also no longer need to flash the ROM each time we change or create files, because we can access them remotely.

The NFS uses the Ethernet support provided by the (rokepxa_extcam3d.sbi) coprocessor. The only problem with NFS is that we cannot have NFS and another coprocessor loaded at the same time.

4.3 Native-compiler

We created a native compiler “by accident”. We needed the ability to have dynamically linked libraries. At first we thought that this is done by “ld”. It turns out that “ld” is just the linker, the dynamic loader is in the kernel (and all it requires is that the LD_LIBRARY_PATH is set correctly).

4.4 Kaffe JVM

<http://www.kaffe.org/>

We worked with Kaffe JVM (ver 1.1.3). The main reason is because it is open source (Sun's code isn't) and there might be a need to modify it. It also scales well with our small system.

5 Research

This section documents the research we did. It should provide you with all the knowledge needed to understand our final implementation.

5.1 Inside Java And It's Virtual Machine

Java is a very rich programming language with many unique¹¹ features. When compared to other languages, Java has a layer of indirection called the JVM. This makes it easy to run the same Java code anywhere. The JVM is specifically designed for the Java programming languages, but can of course run any language (as long as someone writes the compiler).

A Java program is a collection of class files. Class files are created by compiling the java source code. The class file contains a constant pool and Java bytecode, which is the basic instructions that the JVM understands. The bytecode is similar to assembly, it is designed around a stack-based architecture.

The JVM is defined by "The Java Virtual Machine Specification". Besides running bytecode, it also performs verifications (static and dynamic) and garbage collection.

5.1.1 Class Loader

Code in Java is dynamically linked. Each time a class is referenced, a class loader is called to locate and load the bytecode. One of the interesting features of Java is the ability to write a custom class loader.

The JVM has a built-in primordial class loader, which it uses to load itself. Once the JVM has started, every time a class needs to be loaded, it is done with one of the class loader's written in Java (usually one of `java.lang.ClassLoader`, `java.net.URLClassLoader`, `sun.applet.AppletClassLoader`, etc.).

The tasks of the class loader are:

1. Locate the bytecode (on the disk or on the network).
2. Call `defineClass` to check that the bytecode is properly formatted. It also detects `ClassCircularityError` (when a class is itself the superclass of one of it's superclasses).
3. Load the superclass (and the superclass of the superclass, and so on).

¹¹Now being copied by others like .Net

4. Calls `resolveClass` so that the bytecode can be verified.

5.1.2 Bytecode

Java bytecode uses a stack-based architecture. So, for example, adding $2+3$ requires you to do the following:

```
bipush 2  /* Push the number 2 on the stack. */
bipush 3  /* Push the number 3 on the stack. */
iadd     /* Add the stack's top two elements, the arguments
          are popped and the result is on the stack. */
```




Figure 1: adding $2+3$

The bytecode also makes extensive use of strings. These strings are stored in the constant pool and are references by their entry number. For example an object creation or method invocation is done by its name (as a string). Example of a `System.out.println("Hello World");`

```
/* Get the out field from System. */
getstatic java/lang/System/out Ljava/io/PrintStream;

/* Push the constant string on the stack. */
ldc "Hello World"

/* Call the println method. */
invokevirtual java/io/PrintStream/println (Ljava/lang/String;)V
```

There are several types of method invocation:

- `invokevirtual`: invokes a non-static (instance) method.
- `invokestatic`: invokes a static method.
- `invokespecial`: used for constructors.
- `invokeinterface`: used for interfaces.

Methods are identified by their class name, method name and signature. A method can have several attributes:

- `public`: This method is available to all other classes.
- `private`: This method may be accessed only from within this class.

- **protected**: This method may be accessed from any class in the same package or any class that is a subclass of this class.
- **static**: This method is static.
- **final**: This method may not be overridden in subclasses.
- **synchronized**: The JVM will obtain a lock on this object before invoking the method. If this method is static, the JVM will obtain a lock on the Class object.
- **native**: This method is implemented in native code.
- **abstract**: This method has no implementation.

Two methods in a given class may not have the same name and signature (except if the argument part of the signature differ, a feature known as method overloading). It is interesting to note that even if the attributes differ (for eg: a static vs a non-static method) the name and signatures must differ.

The signature is composed of two parts, the argument types and the return type. The argument types is a list (enclosed in parenthesis). The signature has the following format:

$$(type_1 type_2 type_3 \dots) type_{return}$$

See Table 1 for the different formats of *type*. For example the signature of **void main(String argv[])** is $([Ljava/lang/String;)V$

Table 1: Signature formats

Type	Descriptor
array of <i>type</i>	$[type$
byte	B
boolean	Z
char	C
double	D
float	F
int	I
long	J
reference to <i>classname</i>	$Lclassname;$
short	S
void	V

5.1.3 BCEL

<http://jakarta.apache.org/bcel/index.html>

BCEL¹² is a library that is part of the Apache/Jakarta project. This library provides a way to manipulate Java bytecode, providing an amazing amount of power and flexibility (when I discovered this library I felt the same way as when I discovered Java reflexion a couple of years ago: I realized this library provides a whole new dimension to Java).

BCEL can be used after compile time (on *.class* files) or at load time (by writing a custom class loader). It has been used in many projects, including:

- compiler backends
- extending the Java language (multi-inheritance, parameterized classes¹³, preprocessor, custom bytecode, etc...)
- performing optimizations (like peep-hole)

BCEL includes a nice feature BCELifer, which takes a *.class* file and generates the Java code that will generate the *.class* file. This tool proves useful for testing BCEL and development.

The BCEL library (probably for historical reasons) has two parts: one set of classes to read bytecode and one set to modify and generate bytecode. This can be quite confusing at first.

5.1.4 Java Native Interface

The JVM provides a standard way of calling external code written in C (or any other language). The C code must be compiled as a shared library, and a JNI wrapper class must be created (that will wrap calls to Java methods into the shared library).

JNI also provides a way for the native method to create, access and update Java objects.

5.1.5 Garbage Collection

One of the features of Java is to have automatic memory management. This means that the programmer doesn't need to allocate and free memory, it is handled automatically by the JVM (by using a technique known as garbage collection).

Basically a garbage collector finds what objects cannot be referenced by a program and reclaims the storage used by those objects.

At first, we thought that garbage collection could be an issue for our project. If the coprocessor is running in parallel with the JVM, then the

¹²Byte Code Engineering Library

¹³Java 1.5 features this

JVM doesn't know what objects are still being used in the coprocessor. This can cause the garbage collector to delete something the coprocessor might need. Some implementations of garbage collectors also move data around, this can also create problems if a pointer to the data was passed to the coprocessor.

By using JNI, we were able to ignore this aspect in our projects. JNI keeps track of what references are being used by the native application.

5.1.6 Java Hot Spot

The nature of the Java language and its Virtual Machine has brought many people to do research in the field of "Java Hot Spot"¹⁴. The idea is to dynamically analyse a Java program execution (know which methods are invoked most often and where time is being spent), and to take optimization decisions as the program executes. It is also possible to recompile a class and dynamically replace it (HotSwap Class File Replacement).

Unfortunately, Kaffe doesn't have anything similar yet.

5.1.7 Security Issues

Java security is very complex. To put things in a simple way: by using a custom class loader or native methods, the security system is basically doomed. If security is really a matter, there are ways to digitally sign the class files in order to specify what is to be trusted. It is also possible to specify the security level of custom generated code (such as code generated by BCEL).

5.2 IDEA

IDEA is a cryptography algorithm. It takes 64 bit plaintext block and encrypts it using a 128 bit key. It is a symmetric algorithm, so the same key is used to decrypt the coded text.

We used IDEA to compare the time it takes to encrypt data with a software implementation (Java and C) as opposed to a hardware implementation (VHDL).

5.3 Generating VHDL From Java

In theory it is possible to automatically convert Java to VHDL, probably with under some constraints. I did not have enough time to delve into this topic, but it can offer interesting possibilities to our projects. This is definitely a future work.

¹⁴HotSpot is actually a trademark of Sun, but I couldn't find any generic term

5.4 Different Solutions To My Project

There are different ways in which I can get Java to invoke the coprocessor. I opted for a solution that is portable across JVM's. This was not a requirement I had, but it's an elegant method.

- **Modifying the JVM at invocation level**

Since we have the source code of Kaffe, we could modify the invocation method and take a decision whether to execute the coprocessor or normal Java code.

- **Catching “no such method” exceptions**

If we assume that methods implemented in the coprocessor don't have Java implementations, then we could catch “no such method” exceptions and try to run the coprocessor instead. One problem with this solution is how to continue execution from the instruction that caused the exception.

- **Creating a custom class loader**

I opted for this solution. The idea is to create a custom class loader. This class loader is used to load all the classes. If the current class it is loading has methods implemented in the coprocessor, the class loader will modify the bytecode and generate a method invocation (this solution uses Christophe's code). If the class doesn't need any special treatment, the default class loader is used. This method also has the advantage of being able to simultaneously have Java code and a coprocessor allowing the application user to choose which method to use.

Figure 2 illustrates how VMWClassLoader (the custom class loader) is loaded and how it takes over the loading thereafter.

- **Modifying .class files**

A similar method to the above, would be to patch the .class files before running the JVM. The advantage of this method is that it doesn't require a custom class loader. It does have the drawback of being less dynamic than by having a custom class loader.

- **Debugging framework**

Some JVM have debugging support, which means you can write code that intercepts method invocations. This solution is probably the best of all, because it also allows an extension to Java HotSpot technology. We could even imagine extending HotSpot so that the JVM can receive information about the coprocessor (method size in gates and timing).

Unfortunately I thought about this solution near the end of my project. As mentioned previously, Kaffe doesn't support such a feature.

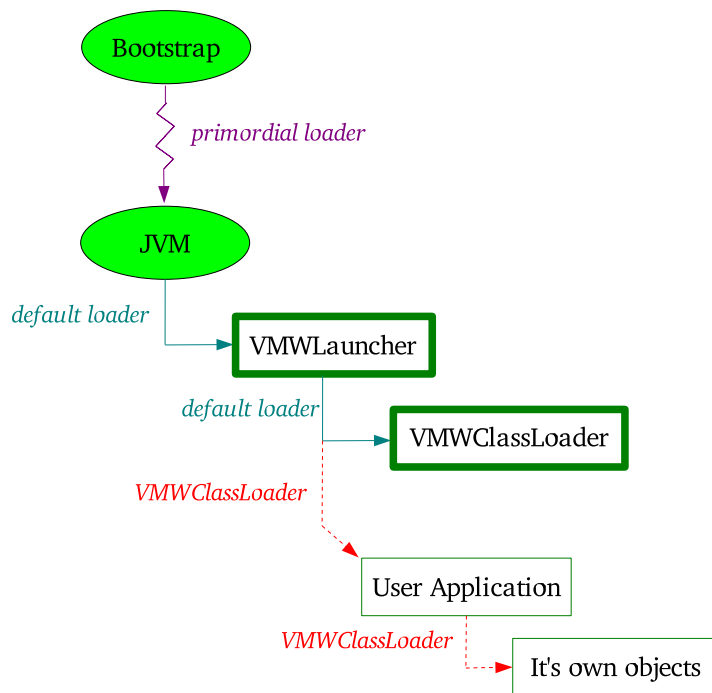


Figure 2: The different class loaders used at each stage

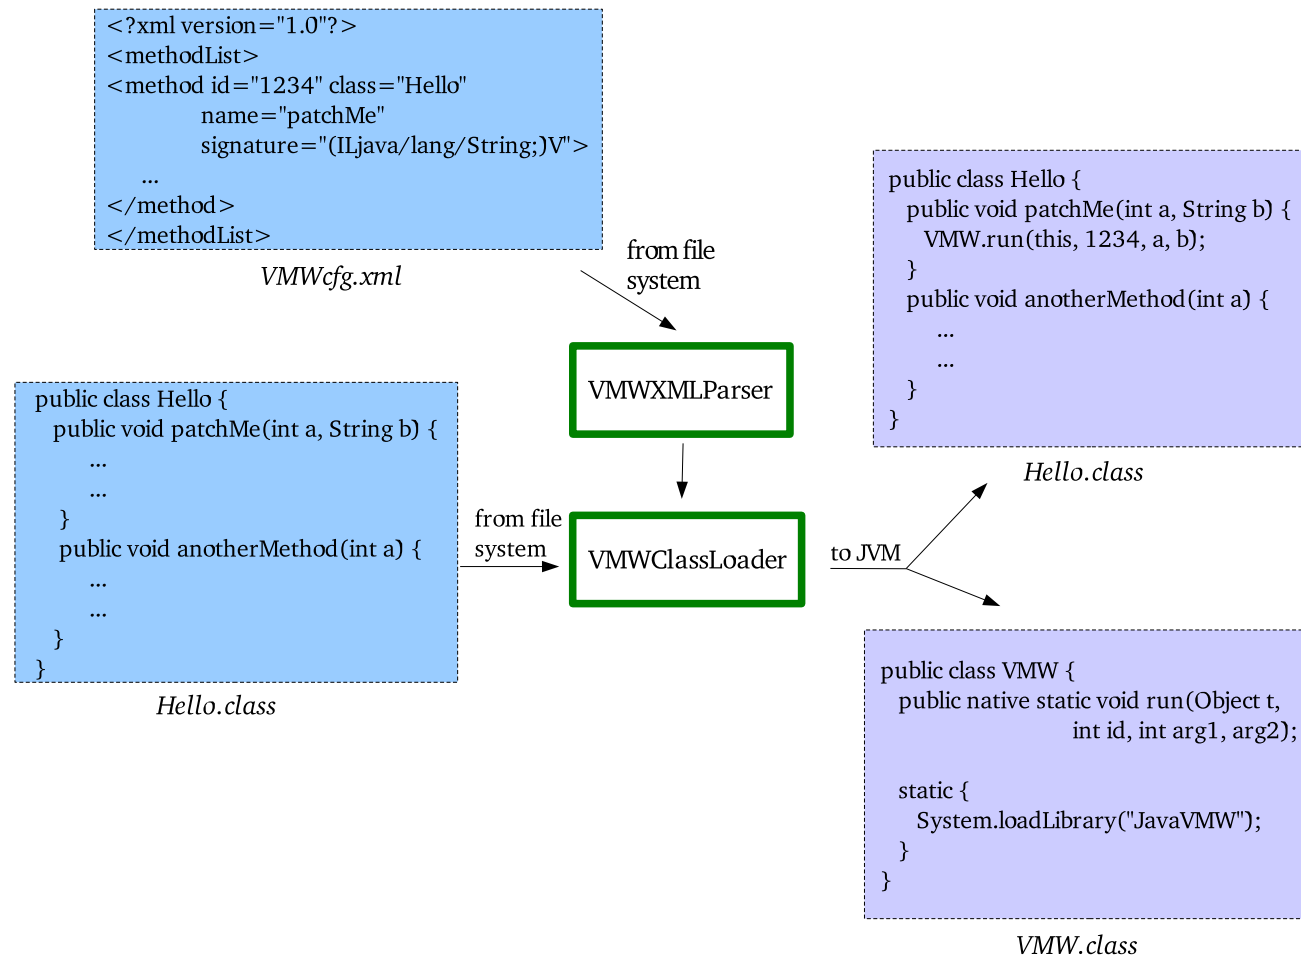


Figure 3: How VMWClassLoader handles the user classes and generates VMW.class

6 Implementation

As discussed above, I used a custom class loader *VMWClassLoader.java* to load all the classes. If the class loader needs to, it can delegate the loading to the default loader.

6.1 XML Config File

We used an XML (*VMWcfg.xml*) based configuration file. One could argue that the coprocessor can't be used in a transparent way if such a config file is necessary. I will not discuss this issue here, as it was necessary for me to find a simple way to know which methods are available as coprocessors and for Christophe's code to know which *.map* file corresponds to each method.

We opted for a XML format, because this allows us to use the same config file for both the projects. You might think this choice is bad because we are working with embedded systems. It turns out that the config file is parsed only once (actually twice, once in my code and once in Christophe's) and the format is simple enough so we can write very efficient parsers if need be.

Here is what it looks like:

```
<?xml version="1.0"?>
<methodList>
  <method id="1" class="AnotherClass"
        name="patchMe" signature="( [S[S[S]V]">
    <map>test.sni</map>
    ...
  </method>
</methodList>
```

There can be many other things inside the `<method>` and `/method` tags, but I only use the attributes of the `<method>` tag (id, class, name and signature).

I used `javacc` (<https://javacc.dev.java.net/>) to generate a parser (`VMWXML-Parser.jj`).

The information I retrieve from the configuration file is saved in `LinkedList` of `ClassIdMethodSignature.java` and `IdMethodSignature.java` objects (these classes are only wrapper classes).

See § 9.1.11 for the dtd of our config file.

6.2 Interfacing With JNI

The file *VMWimp.c* contains the C code written by Christophe that must be called from Java (this code gets compiled into a library called *libJavaVMW.so* on Linux). This code loads the coprocessor and runs it. The function to be called is `Java_Vim_run`:

```
JNIEXPORT void JNICALL Java_VIM_run(JNIEnv *env, jclass cls, jobject this, jint id, ...).
```

The `jobject` is a pointer (it's a real pointer; remember this is a native function) to the object that is calling the coprocessor. This is needed if the coprocessor wants to do callbacks to the JVM. Since we have a level of indirection (see Figure 4), the C code is actually called from a static Java context (the `VMW.class` wrapper); but we are interested in the object that contains the static call (the patched method). The value is `NULL` if the context is static.

It could have been possible to remove this level of indirection, but I think this design is better: less work is required at class load time and it is the only way to do things if you want to hot-swap classes in the future.

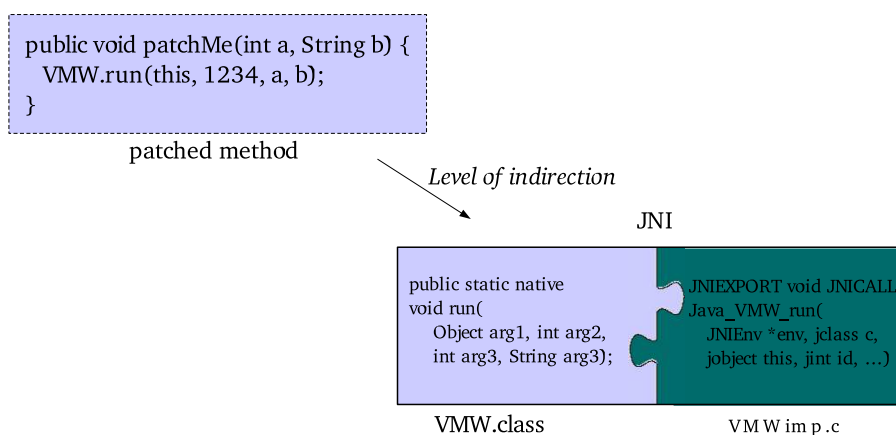


Figure 4: How the patched method invokes the library

At first we thought of passing the class name, method name and method signature to the native library. This worked, but presented a performance problem: for the native code to be able to manipulate the strings, it must perform callbacks to the JVM. Since both our projects share the same configuration file, we opted for an id that is attributed to each `<method>` tag in the config file. The reason the native code needs to know which method is calling it, is primarily so that it can locate the right `.map` file.

As you might have noticed, the C code takes a variable number of arguments. Unfortunately Java doesn't have variable number of arguments, but it has method overloading (methods with the same name but different signatures). On the other hand, C doesn't have method overloading. What JNI does is it maps every method called `run` in the class `VMW.class` to the C function `Java_VMW_run`. This means we must dynamically generate `VMW.class` (see Figure 3); using the XML config file, we will generate a method declaration for each possible signature. (`VMW.class` is actually generated using the `VMWSkel.java` skeleton file).

Here is an example of a *VMW.java* file. Of course this file never actually exists on the disk:

```
public class VMW {
    public native static void run(Object t, Int id, int a, String b);

    static {
        System.loadLibrary(VMWConstants.libFile);
    }
}
```

At first we had thought of using an array of Objects instead of the variable number of arguments. This has the main advantage of making the C code simpler because working with variable arguments is a *pain*. The array of Objects presents some problems though: primitives (int, byte, char, etc.) must be wrapped in an Object and then unwrapped because the native code will handle primitives differently than Objects (the native code can manipulate primitives, but it needs to make callbacks to use Objects). The other problem is related to performance, the C code must make multiple callbacks to the JVM in order to manipulate the array of Objects. It is interesting to note that Java 1.5 supports variable number of arguments, but doesn't solve the problem of wrapping and unwrapping primitives.

6.3 Launcher Class

Creating a custom class loader presents some problems. In order to be able to load all the classes with the custom class loader, we need to replace the default loader. Unfortunately this cannot be done using an environment variable when invoking the JVM. One way to solve this problem is to modify the class package that is included in Kaffe. There is another more elegant way of getting around: by using a launcher class (*VMWLauncher.java*).

The launcher class will be responsible for loading the user's main class using the custom class loader and then invoking the **static void main(Object[] argv)** method. The way the custom class loader is designed, it will be called on any subsequent class loading (see Figure 2).

The launcher class uses reflection to check that the user's main class has a static main method. It also shifts the argv array by one element (so the user's main class has "no idea" there was a level of indirection).

It turns out that Kaffe has a bug, it will not allow the custom class loader to inherit the classpath from the default loader. The work around involves two steps: first I declared a constant (*VMWConstants.kaffebug*) which contains the path where the user's class files reside. The second step is a fix at the custom class loader (see below).

6.4 Custom Class Loader

The custom class loader (*VMWClassLoader.java*) extends `URLClassLoader` which means it can handle `.jar` & `.zip` archives too.

The first thing the loader does is parse the XML config file.

The loader contains code from Kaffe's `URLClassLoader`; due to the class loader architecture, the *VMWClassLoader.java* cannot call `URLClassLoader`'s *myFindClass* or otherwise subsequent class loading won't call *VMWClassLoader*. (the `URLClassLoader`'s `findClass` calls `defineClass`).

The loader also contains code from BCEL due to a misdesign in the library (you can't directly create a *ClassGen* from an array of bytes).

BCEL contains a class loader, but it is not suitable for my project, because it only loads classes with specific name patterns.

Basically the custom loader has four cases:

- **The class is already loaded.** We ask the JVM (`findLoadedClass`) if the class is already loaded. If it is, we just return what `findLoadedClass` returned.
- ***VMW.class* needs to be dynamically created.** We iterate through the linkedlist (generated by the XML parser) and create a method declaration for every possible signature. We mustn't generate two declarations if two different methods have the same signature.
- **The class is listed in the XML config file.** We load the class and patch each of the methods listed in the config file. We replace the method's body by an invocation to:
`VMW.run(this, id, arguments_of_the_original_method)`. The first argument is `NULL` if the method being patched is static. If the method is synchronized, it will generate the `monitorenter` and `monitorexit` code.
- **The class is a normal class (none of the above 3 cases).** In the case of a normal class, the custom class loader must still load it manually (and not delegate to the system loader) in order to maintain control over what the normal class will load. If the custom class loader can't find the class (this can happen in Kaffe due to the bug discussed above), then the default System loader will be called. (This should be fixed as soon as the bug in Kaffe is solved).

It is important to notice that the original arguments are numbered from 1 to N for non-static methods (because argument 0 is 'this'), but from 0 to N-1:

```
public static void Hello(int a) {
    /* in this method, a is argument 0,
       it can be accessed with ILOAD(0) */
```

```
}  
  
public void Hello2(int a) {  
    /* in this method, a is argument 1,  
    it can be accessed with ILOAD(1) */  
}
```

Note: I discovered and fixed a bug in BCEL that prevents it from generating the right code for native methods. The fix has been accepted for the next release of BCEL, but until then, you must use my patched version (bcel-5.1b).

7 Limitations / Improvements

- Methods that are patched don't return anything. Currently this is not really a limitation, since we don't have a limit on the number of arguments (and you can always replace a return argument by an extra input argument). This might become a limitation in the future. We made this decision in order to simplify various parts of the code (VMWClassLoader doesn't need to generate different return instructions, the native method doesn't need to handle return values, etc.).
- We have used file names such as VMWClassLoader, VMWXMLParser, etc. It would be a wise idea to declare a package name for all these files and avoid having name space collisions.
- The config file is actually parsed twice, once in the Java code and once in the native code. Although each parser uses different information, it is possible to parse the file only once.
- The class files are also parsed twice. Once by BCEL before patching, and once by the JVM (VMWClassLoader calls `defineClass` which takes an array of bytes). This is a negligible thing.
- It would be interesting to be able to reload classes based on real-time statistics (HotSpot). This could be an interesting future work.
- It would be interesting to research how to convert Java code into VHDL automatically. This can also be an interesting future work.

8 Conclusions and Future work

The main goal of my project has been reached. I am successfully able to call coprocessors from Java, without modifying or recompiling the source code. The work I did can also be used to call C code from Java (thus

simplifying the coding of native libraries for Java). We could imagine that depending on the hardware available or security policies, a native code or Java code could be executed.

The overhead of the custom class loader is very negligible compared to the time gained in the coprocessor. The time required to load my code and bcel is very little as compared to the time required to load the JVM. The patching of methods is very fast, and is done only once. Currently the main problem is in the time required to load the coprocessor, but this time can be “hidden” by having the processor do something else meanwhile.

This project involved very interesting aspects of Java programming. I gained detailed insight at the internals of a JVM. I also spent a lot of time setting up the development environment. This process was time consuming and sometime frustrating, but it clearly highlighted how difficult and complex open source projects can sometimes be. By using hardware combinations that not a lot of people have tried, we ran into new problems on a day to day basis.

There are two open questions that I wish I would have had more time to work on: how to replace a class in the JVM in real-time and how to generate VHDL from Java bytecode. Both these problems seem complex and could very well be an entire project on their own, but these are the interesting problems that will provide new properties to my existing project.

An interesting possibility I did not cover in this report, is that the coprocessor can call the JVM (by using JNI). We did get this working (by displaying a “Hello World” back from the coprocessor), but we didn’t explore the possibilities of such a “reverse” system.

9 Appendix

9.1 Code

9.1.1 Organization

Here is the list of files and what they do:

- **VMWXMLParser.jj** This is the javacc file that will generate the XML parser. Once compiled by javacc, it will generate the following files: VMWXMLParser.java, TokenMgrError.java, ParseException.java, Token.java, SimpleCharStream.java.
- **VMWClassLoader.java** This is the custom class loader that does all the work.
- **VMWLauncher.java** The launcher used to set the custom class loader.
- **VMWConstants.java** File containing all the constants used in the program.
- **ClassIdMethodSignature.java** This file is just a wrapper.
- **IdMethodSignature.java** This file is just a wrapper.
- **VMWSkel.java** Used to generate VMW.class on the fly.
- **VMW.class** Generated on the fly, never exists as a file.
- **VMWcfg.xml** User configuration file.
- **libJavaVIM.so** Native library.
- **VWXMLcfg.dtd** DTD for config file.

9.1.2 Compiling & Running

Compiling the source code is straight forward, all you need is the bcel library:

```
javacc VMWXMLParser.jj
javac -classpath 'bcel-5.1b.jar' *.java
To run some code (eg: Hello) using my code:
java -cp 'bcel-5.1b.jar:.' VMWLauncher Hello parameters for Hello
```

9.1.3 VMXMLParser.jj

Listing 1: VMXMLParser.jj

```
1 /*
2  * Created on Jan 14, 2004
3  *
4  * Javacc file , to parse the xml config file .
5  * I didn't implement the entire dtd , only what
6  * interests me. The goal is to be the most
7  * optimized possible .
8  *
9  * When parseFile is called (from VMWClassLoader),
10 * the static LinkedList config will contain ClassIdMethodSignature
    objects .
11 *
12 */

14 /**
15  * @author Alok Menghrajani
16  *
17  */

19 options {
20     STATIC = true;
21     DEBUGPARSER = false;
22 }

24 PARSER_BEGIN(VMXMLParser)

26 import java.io.FileReader;
27 import java.util.*;

29 public class VMXMLParser {

31     public static LinkedList config;

33     /* For testing purpose only */
34     public static void main(String args[]) throws Exception {
35         FileReader in = new FileReader(args[0]);
36         VMXMLParser parser = new VMXMLParser(in);
37         parser.XMLDocument();
38         System.out.println("Ok");
39     }

41     public static void parseFile(String file) throws Exception {
42         /*
43          * This is the method that will be called from VMWClassLoader .
44          */
45         FileReader in = new FileReader(file);
46         VMXMLParser parser = new VMXMLParser(in);
47         parser.XMLDocument();
48     }
49 }

51 PARSER_END(VMXMLParser)

53 SKIP : {
54     " " | "\r" | "\t" | "\n"
55 }

57 TOKEN : {
```



```

58     < LT: "<" >
59     | < GT: ">" >
60     | < EQUAL: "=" >
61     | < SLASH: "/" >
62     | < QUOTE: "\"" >
63     | < XML_S: <LT> "?" >
64     | < METHODLIST_S: <LT> <METHODLIST> <GT> >
65     | < METHODLIST_E: <LT> <SLASH> <METHODLIST> <GT> >
66     | < METHOD_S: <LT> <METHOD> >
67     | < METHOD_E: <METHOD> <GT> >
68     | < ID: "id" >
69     | < CLASS: "class" >
70     | < NAME: "name" >
71     | < SIGNATURE: "signature" >
72     | < #METHOD: "method" >
73     | < #METHODLIST: "methodList" >
74     | < ANY: ~[] >
75     | < VAL: <QUOTE> (~["\""]) * <QUOTE> >
76 }

79 JAVACODE

81 /* Flush until a '>' is encountered. Useful when you want to
82 skip a tag that you don't care about. */
83 void flush() {
84     while ((getNextToken()).kind != GT)
85         ;
86 }

88 String string() : {} {
89     <VAL> {
90         String s = token.image;
91         return s.substring(1, s.length()-1);
92     }
93 }

95 void XMLDocument() throws Exception : {
96     String cl, n, sig;
97     short id;
98 } {
99     {
100         /* Initialize the Vectors */
101         config = new LinkedList();
102     }
103     <XML_S> {flush();} /* skip <?xml version ?> */
104     <METHODLIST_S>

106     (<METHOD_S>
107         <ID><EQUAL>{id=Short.parseShort(string());}
108         <CLASS><EQUAL>{cl=string();}
109         <NAME><EQUAL>{n=string();}
110         <SIGNATURE><EQUAL>{sig=string();}
111         {
112             /* Search to linkedlist to see if I don't already have this
113              class */
114             ListIterator i = config.listIterator(0);
115             boolean flag = false;
116             while (i.hasNext()) {
117                 ClassIdMethodSignature cims = (ClassIdMethodSignature) i.next
118                     ();
119                 if (cims.className.equals(cl)) {

```

```

118     /* Add the id, method and signature only.
119     (Perform check to see that this combination of
120     method/signature doesn't already exist)
121     */
122     ListIterator j = cims.idMethodSignature.listIterator(0);
123     while (j.hasNext()) {
124         IdMethodSignature ims = (IdMethodSignature) j.next();
125         if (ims.method.equals(n) && ims.signature.equals(sig)) {
126             // Check that method returns void !
127             // xml file is bad.
128             // Warning: I'm not going to check
129             // for something like this (totally illegal):
130             // int add(int, int) and
131             // String add(int, int) (two different
132             // signatures, but collision !)
133             throw new RuntimeException("Bad XML file");
134         }
135     }
136     cims.idMethodSignature.add(new IdMethodSignature(id, n, sig)
137     );
138     flag = true;
139 }
140 /* Add the class */
141 if (!flag) {
142     ClassIdMethodSignature cims = new ClassIdMethodSignature(cl);
143     cims.idMethodSignature.add(new IdMethodSignature(id, n, sig));
144     config.add(cims);
145 }
146 }
147 <GT>
148 { while((getNextToken()).kind != METHODE); } /* skip what's under
149    the <method> tag */
150 )*
151 <METHODLIST_E> <EOF>
152 }

```

9.1.4 VMWClassLoader.java

Listing 2: VMWClassLoader.java

```

1 /*
2  * Created on Jan 13, 2004
3  *
4  * The 'core' of my project.
5  * This is a special class loader which performs the following tasks:
6  *
7  * - reads XML config file.
8  *
9  * - loads classes, that aren't listed in XML, normally (using
10  *   VMWClassLoader if the
11  *   class can be found, otherwise using the system loader if the
12  *   class can't be found
13  *   (should only happen in Kaffe due to a little bug)).
14  *
15  * - loads classes specified in config file (classes that contain
16  *   optimized method) and 'patches' the optimized methods by calling
17  *   VMW.run
18  *
19  *

```

```

16 * - creates special class VMW (native library wrapper),
17 *   adding one run method for each possible signature (by looking in
18 *   the config file). This class is created from 'scratch' using
    VMWSkel.
19 *
20 * VMW's run methods take 2 extra parameters:
21 * this = a reference to the caller (or null if context is static).
22 * id = id allowing the native code to find the entry in the config
    file.
23 */
24
25 /**
26 * @author Alok Menghrajani
27 *
28 */
29
30 import java.net.*;
31 import java.io.*;
32 import java.util.*;
33
34 import org.apache.bcel.*;
35 import org.apache.bcel.classfile.*;
36 import org.apache.bcel.generic.*;
37
38 public class VMWClassLoader extends URLClassLoader {
39
40     static {
41         if (VMWConstants.debug)
42             System.out.println("VMWClassLoader loading...");
43         try {
44             /* Parse the XML config file. */
45             VMWXMLParser.parseFile(VMWConstants.configFile);
46             System.out.println("VMWClassLoader ready.");
47         } catch (FileNotFoundException e) {
48             System.out.println("XML config file "+VMWConstants.configFile+"
                not found.");
49             System.exit(-1);
50         } catch (Exception e) {
51             System.out.println(e.getMessage());
52             e.printStackTrace();
53             System.exit(-1);
54         }
55     }
56
57     public VMWClassLoader(URL[] url) {
58         super(url);
59     }
60
61     public Class loadClass(String name) throws ClassNotFoundException {
62         return loadClass(name, false);
63     }
64
65     public Class loadClass(String name, boolean resolve) throws
        ClassNotFoundException {
66         /* This method is called to load a class.
67         *
68         * There are 4 possibilities:
69         * 1) The class is already loaded
70         *    -> I test this using findLoadedClass.
71         * 2) The class is listed in the XML config file
72         *    -> I load it using myFindClass and then
73         *        patch it using BCEL.

```

```

74     * 3) The class to be loaded is VMW (java native class)
75     *     -> Build it from scratch.
76     * 4) The class is a normal/system class
77     *     -> I load it using myFindClass or findSystemClass.
78     *
79     */

81     Class result = findLoadedClass(name);
82     if (result != null) {
83         /* We are in case 1 - the class is already loaded. */
84         if (VMWConstants.debug)
85             System.out.println("Already loaded: " + name);
86         return result;
87     }

89     /* Check to see if it's VMW */
90     if (name.equals("VMW")) {
91         if (VMWConstants.debug)
92             System.out.println("Create VMW");
93         return createVMW();
94     }

96     /* Check the VMWXMLParser.config LinkedList to see if
97        it's a class that needs to be patched */
98     ListIterator i = VMWXMLParser.config.listIterator(0);
99     while (i.hasNext()) {
100         ClassIdMethodSignature cims = (ClassIdMethodSignature) i.next
101             ();
102         if (cims.className.equals(name)) {
103             if (VMWConstants.debug)
104                 System.out.println("Special class: " + name);
105             byte[] bytes = myFindClass(name);
106             result = patchClass(cims, bytes);
107             if (resolve) {
108                 resolveClass(result);
109             }
110             return result;
111         }
112     }

113     /* We are in case 2 - the class is a system/normal class */
114     try {
115         byte[] bytes = myFindClass(name);
116         if (VMWConstants.debug)
117             System.out.println("Normal Class: " + name);
118         Class cl = defineClass(name, bytes, 0, bytes.length);
119         if (resolve) {
120             resolveClass(cl);
121         }
122         return cl;
123     } catch (Exception e) {
124         if (VMWConstants.debug)
125             System.out.println("System Class: " + name);
126         return findSystemClass(name);
127     }
128 }

130 private byte[] myFindClass(String name) throws
131     ClassNotFoundException {
132     /* Mostly inspired/pasted from Kaffe sources. */
133     /* It seems Kaffe imposes a 1024 max file size for class */

```

```

134     URL url = findResource(name.replace('.', '/') + ".class");
135     if (url == null) {
136         throw new ClassNotFoundException(name);
137     }
138     try {
139         InputStream in = url.openStream();
140         ByteArrayOutputStream out = new ByteArrayOutputStream();
141         byte[] buf = new byte[1024];
142         for (int r; (r = in.read(buf)) != -1; ) {
143             out.write(buf, 0, r);
144         }
145         in.close();
146         buf = out.toByteArray();
147         return buf;
148     } catch (IOException e) {
149         throw new ClassNotFoundException(name + ": " + e);
150     }
151 }

153 private JavaClass createJavaClass(String classname, byte[] bytes) {
154     /* Inspired by BCEL sources. */
155     JavaClass clazz = null;
156     try {
157         ClassParser parser = new ClassParser(new ByteArrayInputStream(
158             bytes), "foo");
159         clazz = parser.parse();
160     } catch (Throwable e) {
161         e.printStackTrace();
162         return null;
163     }

164     ConstantPool cp = clazz.getConstantPool();
165     ConstantClass cl = (ConstantClass)cp.getConstant(clazz.
166         getClassNameIndex(), Constants.CONSTANT_Class);
167     ConstantUtf8 name = (ConstantUtf8)cp.getConstant(cl.getNameIndex
168         (), Constants.CONSTANT_Utf8);
169     name.setBytes(classname.replace('.', '/'));
170     return clazz;
171 }

171 private Class patchClass(ClassIdMethodSignature cims, byte[] bytes)
172 {
173     /* Patches Class cl, so that all invocations of methods
174     * that appear in XML config file are rerouted to VMW.run
175     *
176     * Here are the major steps performed:
177     * 1) Find every method (listed in cims) in Class
178     * 2) Convert signature from String to Type[].
179     * 3) Replace body by a call to VMW.run(this (or null),
180     *    id, param1, param2, ...)
181     */

182     JavaClass jc = createJavaClass(cims.className, bytes);
183     ClassGen clazz = new ClassGen(jc);
184     Method methods[] = clazz.getMethods();
185     for (int i=0; i<methods.length; i++) {
186         /* Search in cims.methodSignature LinkedList for
187         * a match with methods[i].getName() and methods[i].
188         * getSignature().
189         */
190         ListIterator j = cims.idMethodSignature.listIterator();
191         while (j.hasNext()) {

```

```

191     IdMethodSignature ims = (IdMethodSignature)j.next();
192     if (ims.method.equals(methods[i].getName()) &&
193         ims.signature.equals(methods[i].getSignature())) {
194         /* Set the flag, so we can show a warning if
195          * a method isn't present in the class but
196          * is in the XML config file.
197          */
198         if (ims.flag) {
199             throw new RuntimeException("flag already set");
200         }
201         ims.flag = true;
202         /* Patch this method */
203         Method patch = patchMethod(clazz, methods[i], ims.id);
204         clazz.replaceMethod(methods[i], patch);
205     }
206 }
207
208
209 /* Check the flags */
210 ListIterator i = cims.idMethodSignature.listIterator();
211 while (i.hasNext()) {
212     IdMethodSignature ims = (IdMethodSignature) i.next();
213     if (!ims.flag && VMWConstants.debug) {
214         System.out.println("Warning: method "+cims.className+"."+ims.
215             method+"{"+
216                 ims.signature+"} not found in class file but used in
217                 config file.");
218     }
219 }
220
221 /* Register new class */
222 byte[] output = clazz.getJavaClass().getBytes();
223 Class cl = defineClass(cims.className, output, 0, output.length);
224 return cl;
225 }
226
227 private Method patchMethod(ClassGen clazz, Method o, short id) {
228     /* Replace the body of this method by a call to VMW.run */
229     ConstantPoolGen cpg = clazz.getConstantPool();
230     MethodGen mg = new MethodGen(o, clazz.getClassName(), cpg);
231
232     /* We will call a run method of VMW that has the same signature.
233      * We still need to create a new signature, because VMW will take
234      * a few extra parameters (this, id)
235      */
236     Type[] arguments = Type.getArgumentTypes(o.getSignature());
237     Type[] newtype = new Type[2+arguments.length];
238     newtype[0] = Type.OBJECT; // reference to caller
239     newtype[1] = Type.INT;
240     for (int i=0; i<arguments.length; i++) {
241         newtype[i+2] = arguments[i];
242     }
243     String newsignature = Type.getMethodSignature(Type.VOID, newtype);
244     if (VMWConstants.debug)
245         System.out.println("NEW SIGNATURE="+newsignature);
246
247     int vmwrun = cpg.addMethodref("VMW", "run", newsignature);
248     if (vmwrun==-1) {
249         System.out.println("Fatal error in patchMethod (vmwrun=-1)");
250         System.exit(-1);
251     }
252 }

```

```

251     /* Start creating new instructions */
252     InstructionList il = new InstructionList();

254     /* Check if this method is in a static context. */
255     /* in static context, var 0 = param 1, var 1 = param 2, etc..
256     in instance context, var 0 = this, var 1 = param 1, etc...
257     */
258     int first_param;

260     if (o.isStatic()) {
261         il.append(new ACONST_NULL());
262         first_param = 0;
263     } else {
264         il.append(new ALOAD(0));
265         first_param = 1;
266     }

268     il.append(new SIPUSH(id));

270     for (int i=0; i<arguments.length; i++) {
271         il.append(new ILOAD(i+first_param));
272     }
273     il.append(new INVOKESTATIC(vmwrun));
274     il.append(new RETURN());

276     mg.setInstructionList(il);
277     mg.removeLineNumbers();
278     mg.setMaxStack();

280     return mg.getMethod();
281 }

283 private Class createVMW() {
284     /* Create VMW Class using VMWSkel */
285     ClassGen clazz = new ClassGen(Repository.lookupClass("VMWSkel"));
286     clazz.setClassName("VMW");

288     ConstantPoolGen cpg = clazz.getConstantPool();
289     InstructionFactory factory = new InstructionFactory(clazz, cpg);

291     InstructionList il = new InstructionList();
292     /* Create a method declaration for each possible signature */
293     LinkedList signatures = new LinkedList();
294     ListIterator i = VMWXMLParser.config.listIterator(0);
295     while (i.hasNext()) {
296         ClassIdMethodSignature cims = (ClassIdMethodSignature) i.next();
297         ListIterator j = cims.idMethodSignature.listIterator(0);
298         while (j.hasNext()) {
299             IdMethodSignature ims = (IdMethodSignature) j.next();
300             /* Check to see if we haven't already handled this signature
301             */
301             if (!signatures.contains(ims.signature)) {
302                 signatures.add(ims.signature);
303                 Type[] arguments = Type.getArgumentTypes(ims.signature);
304                 Type[] newtype = new Type[2+arguments.length];
305                 newtype[0] = Type.OBJECT; // reference to caller
306                 newtype[1] = Type.INT; // id
307                 for (int k=0; k<arguments.length; k++) {
308                     newtype[k+2] = arguments[k];
309                 }

311                 String[] strings = new String[newtype.length];

```

```

312         for (int k = 0; k<strings.length; k++) {
313             strings[k] = "arg"+k;
314         }
315         /* We can assume the return type is void */
316         if (VMWConstants.debug)
317             System.out.println("ADDING: " + ims.signature);
318
319         MethodGen mg = new MethodGen(Constants.ACC_PUBLIC |
320             Constants.ACC_STATIC | Constants.ACC_NATIVE, Type.VOID,
321             newtype, strings, "run", "VMW", il, cpg);
322
323         clazz.addMethod(mg.getMethod());
324     }
325 }
326
327 /* Register new class */
328 byte[] output = clazz.getJavaClass().getBytes();
329 Class cl = defineClass(clazz.getClassName(), output, 0, output.
330     length);
331 return cl;
332 }

```

9.1.5 VMWLauncher.java

Listing 3: VMWLauncher.java

```

1  /*
2  * Created on Jan 13, 2004
3  *
4  * This is a wrapper to launch Java applications (name passed in arg
5  * [0]).
6  * The goal is the have AlokClassLoader load the classes for this
7  * Application.
8  * The remaining arguments will be passed to the main Java application
9  *
10 *
11 */
12 /**
13 * @author Alok Menghrajani
14 *
15 */
16
17 import java.lang.reflect.*;
18 import java.net.*;
19
20 public class VMWLauncher {
21
22     public static void main(String args[]){
23         if (args.length==0) {
24             System.out.println("Usage: java AlokLauncher appName extraargs")
25             ;
26             System.exit(-1);
27         }
28
29         // Check to see if there are any arguments to be
30         // passed to the main Java app.

```



```

30     String args2[] = new String[0];

32     if (args.length>1) {
33         args2 = new String[args.length-1];
34         for (int i=1; i<args.length; i++) {
35             args2[i-1] = args[i];
36         }
37     }

39     // Recover current search path for custom class loader.
40     ClassLoader cl = Thread.currentThread().getContextClassLoader();

42     URL urls[] = ((URLClassLoader) cl).getURLs();
43     if (urls.length == 0) {
44         urls = new URL[1];
45         try {
46             urls[0] = new URL(VMWConstants.kaffebug);
47         } catch (MalformedURLException e) {
48             System.out.println("Malformed filename in VMWConstants (String
49                 kaffebug)");
50             System.exit(-1);
51         }
52     }
53     VMWClassLoader customLoader=new VMWClassLoader(urls);

54     try {
55         Class mainClass = customLoader.loadClass(args[0]);

56         // Check using reflection that main Java
57         // app has a function main(String[])
58         Class parameters[] = new Class[1];
59         parameters[0] = String[].class;

60         Method method = mainClass.getMethod("main", parameters);
61         Object[] arguments = new Object[1];
62         arguments[0]=args2;

63         method.invoke(null, arguments);
64     } catch (ClassNotFoundException e) {
65         // The requested main Java app wasn't found.
66         System.out.println("The class "+args[0]+" was not found.");
67         System.exit(-1);
68     } catch (NoSuchMethodException e) {
69         // Reflection didn't find main function.
70         System.out.println("It seems "+args[0]+" doesn't have a main.");
71         System.exit(-1);
72     } catch (Exception e) {
73         System.out.println(e.getMessage());
74         e.printStackTrace();
75     }
76 }
77 }
78 }
79 }
80 }

```

9.1.6 VMWConstants.java

Listing 4: VMWConstants.java

```

1 public class VMWConstants {
2     /* path to config file */
3     public static final String configFile = "/etc/VMWcfg.xml";

```

```
5  /* name of VMW library file */
6  /* on linux the name will be automatically converted to "lib"+
   libFile+".so"
7  and on windows the name will be libFile+".dll" */
9  public static final String libFile = "JavaVMW";
11 /* solve a little bug in Kaffe */
12 public static final String kaffebug = "file:///";
13 // public static final String kaffebug = "file:///home/menghraj
   /undergrad/epfl_4/sem_proj_epxa/vmw/";
14 public static final boolean debug = true;
15 }
```

9.1.7 ClassIdMethodSignature.java

Listing 5: ClassIdMethodSignature.java

```
1  /*
2  * Created on Jan 15, 2004
3  *
4  * Class to encapsulate a String (representing a class) and
5  * a LinkedList of IdMethodSignatures (which itself is a Class to
6  * encapsulate two Strings and an int).
7  *
8  * I have just proven how intrusive Java can sometimes become.
9  *
10 /*
12 /**
13 * @author Alok Menghrajani
14 *
15 */
17 import java.util.*;
19 public class ClassIdMethodSignature {
20     public String className;
21     public LinkedList idMethodSignature;
23     public ClassIdMethodSignature(String c) {
24         this.className = c;
25         this.idMethodSignature = new LinkedList();
26     }
27 }
```

9.1.8 IdMethodSignature.java

Listing 6: IdMethodSignature.java

```
1  /*
2  * Created on Jan 15, 2004
3  *
4  * Class to encapsulate two Strings (the first representing a method,
5  * the second representing a signature) and an int (representing the
   id in the code).
```

```
6 * This int is used for optimization (don't need to pass Strings to
   * the native library).
7 *
8 * I wish there was a way to avoid the creation of such silly classes
   * ...
9 *
10 */
11
12 /**
13 * @author Alok Menghrajani
14 *
15 */
16
17 public class IdMethodSignature {
18
19     public String method;
20     public String signature;
21     public short id;
22     public boolean flag;
23     /* This flag is used to display a warning if
24      * a method is specified in the XML config file but not found
25      * in the actual class.
26      */
27
28     public IdMethodSignature(short id, String m, String s) {
29         this.method = m;
30         this.signature = s;
31         this.id = id;
32         this.flag = false;
33     }
34 }
```

9.1.9 VMWSkel.java

Listing 7: VMWSkel.java

```
1 /*
2 * Created on Jan 15, 2004
3 *
4 * Skeleton Class to build VMW.class.
5 * I could do things without this class, but it's easier this
6 * way (and it is more flexible since changing this class is easy).
7 *
8 */
9
10 /**
11 * @author Alok Menghrajani
12 */
13
14 public class VMWSkel {
15     static {
16         System.loadLibrary(VMWConstants.libFile);
17     }
18 }
```

9.1.10 VMWcfg.xml

Listing 8: VMWcfg.xml

```

1 <?xml version="1.0"?>
2 <methodList>
3   <method id="1" class="AnotherClass" name="patchHim" signature="( [S[S
   [S)V">
4     <map>test.sni</map>
5     <methods maxid="4">
6       <lookupInfo id="1">
7         <kind>instance</kind>
8         <class>java/lang/String</class>
9         <name>&lt ; init ></name>
10        <signature >()V</signature >
11      </lookupInfo >
12      <lookupInfo id="2">
13        <kind>static</kind>
14        <class>java/lang/System</class >
15        <name>gc</name>
16        <signature >()V</signature >
17      </lookupInfo >
18      <lookupInfo id="3">
19        <kind>instance</kind>
20        <class>java/io/PrintStream</class >
21        <name>println</name>
22        <signature >()V</signature >
23      </lookupInfo >
24      <lookupInfo id="4">
25        <kind>instance</kind>
26        <class>java/io/PrintStream</class >
27        <name>println</name>
28        <signature >(Ljava/lang/String ; )V</signature >
29      </lookupInfo >
30    </methods >
31    <fields maxid="1">
32      <lookupInfo id="1">
33        <kind>static</kind>
34        <class>java/lang/System</class >
35        <name>out</name>
36        <signature >Ljava/io/PrintStream;</signature >
37      </lookupInfo >
38    </fields >
39  </method >
41 </methodList >

```

9.1.11 VMXMLcfg.dtd

Listing 9: VMWcfg.dtd

```

1 /*
2  * Created on Jan 14, 2004
3  *
4  */
5
6 /**
7  * @author Alok Menghrajani
8  *
9  */
11 <ELEMENT methodList (method*)>

```

```

12 <ELEMENT method (map, methods*, fields*)>
13 <ELEMENT map (#PCDATA)>
14 <ELEMENT methods (lookupInfo+)>
15 <ELEMENT fields (lookupInfo+)>
16 <ELEMENT lookupInfo (kind, class, name, signature)>
17 <ELEMENT kind (#PCDATA)>
18 <ELEMENT class (#PCDATA)>
19 <ELEMENT name (#PCDATA)>
20 <ELEMENT signature (#PCDATA)>

22 <!ATTLIST method
23     id CDATA #REQUIRED
24     class CDATA #REQUIRED
25     name CDATA #REQUIRED
26     signature CDATA #REQUIRED>
27 <!ATTLIST methods maxid CDATA #REQUIRED>
28 <!ATTLIST fields maxid CDATA #REQUIRED>
29 <!ATTLIST lookupInfo id CDATA #REQUIRED>

```

9.1.12 IdeaPlus.java

Listing 10: IdeaPlus.java

```

1 public class IdeaPlus {
2
3     /* Number of IDEA rounds */
4     private static final short IDEA_ROUNDS = 8;
5
6     /* Number of IDEA subkeys */
7     public static final short IDEA_SK_NUM = (6 * IDEA_ROUNDS + 4);
8
9     /* low significant 16-bit */
10    private static short lsw16(int y){
11    return (short)(y & 0xffff);
12    }
13
14    /* most significant 16-bit */
15    private static short msw16(int y){
16    return (short)((y>>16) & 0xffff);
17    }
18
19    /* 2**16 + 1 */
20    private static final int MULMOD = (1<<16) | 1;
21
22    // Multiplication modulo 2**16 + 1
23    private static short mul(short x, short y)
24    {
25    short t16;
26    int t32;
27
28    x = lsw16(x-1);
29    t16 = lsw16(y - 1);
30    t32 = (int) (((char)x * (char)t16 + (char)x + (char)t16 + 1);
31    x = lsw16(t32);
32    t16 = msw16(t32);
33    x = (short)(((char)x - (char)t16) + (((char)x)<=((char)t16) ? 1 : 0)
34    );
35    return x;
36    }
37 }

```

```

37     // Compute multiplicative inverse of x by Euclid's GCD algorithm.
38     private static short mul_inv(short x)
39     {
40     int n1 = MULMOD;
41     int n2 = (int)(char)x;
42     int b1 = 0;
43     int b2 = 1;
44     int q, r, t;

46     if ((char)x<=1)
47         return x;

49     while (true) {
50         r = n1 % n2;
51         q = n1 / n2;
52         if (r==0) {
53             if (b2<0) {
54                 b2 += MULMOD;
55             return lsw16(b2);
56             }
57             else {
58                 n1 = n2;
59                 n2 = r;
60                 t = b2;
61                 b2 = b1-(q*b2);
62                 b1 = t;
63             }
64         }
65     }

68     // Computes IDEA encryption subkeys.
69     public static void idea_encrypt_subkeys
70     (short [] key, short [] subkeys)
71     {
72     int i;

74     for (i=0;i<key.length;i++)
75         subkeys[i] = key[i];

77     for (;i<IDEA_SK_NUM;i++)
78         subkeys[i] = lsw16(((char)(subkeys[((i+1) % 0x8)!=0 ? i-7 : i
79             -15]) << 9) |
80             ((char)(subkeys[((i+2) % 0x8) < 2 ? i-14 : i-6]) >> 7));

82     // Computes IDEA decryption subkeys from encryption subkeys.
83     public static void idea_decrypt_subkeys
84     (short [] encrypt_subkeys, short [] decrypt_subkeys)
85     {
86     int pen = 0;
87     int pde = 0;
88     short [] t = new short[encrypt_subkeys.length];
89     int pt = 0;
90     int i;

92     t[6 * IDEA_ROUNDS] = mul_inv(encrypt_subkeys[pen++]);
93     t[6 * IDEA_ROUNDS + 1] = lsw16(-encrypt_subkeys[pen++]);
94     t[6 * IDEA_ROUNDS + 2] = lsw16(-encrypt_subkeys[pen++]);
95     t[6 * IDEA_ROUNDS + 3] = mul_inv(encrypt_subkeys[pen++]);

97     for (i=6*(IDEA_ROUNDS-1);i>=0;i-=6) {

```

```

98     t[i + 4] = encrypt_subkeys[pen++];
99     t[i + 5] = encrypt_subkeys[pen++];
100    t[i] = mul_inv(encrypt_subkeys[pen++]);
101    if (i!=0) {
102    t[i + 2] = lsw16(-encrypt_subkeys[pen++]);
103    t[i + 1] = lsw16(-encrypt_subkeys[pen++]);
104    }
105    else {
106    t[1] = lsw16(-encrypt_subkeys[pen++]);
107    t[2] = lsw16(-encrypt_subkeys[pen++]);
108    }
109    t[i + 3] = mul_inv(encrypt_subkeys[pen++]);
110 }

112 for (i=0;i<IDEA.SK_NUM;i++) {
113     decrypt_subkeys[pde++] = t[pt];
114     t[pt++] = 0;
115 }
116 }

119 //IDEA encryption/decryption algorithm.
120 // Note: block_in and block_out can be the same block.
121 public static void idea_cipher
122 (short [] block_in , short [] block_out , short [] key)
123 {
124 int pin = 0;
125 int pout = 0;
126 int pk = 0;
127 short word1 , word2 , word3 , word4;
128 short t1 , t2;
129 int i;

131 word1 = block_in[pin++];
132 word2 = block_in[pin++];
133 word3 = block_in[pin++];
134 word4 = block_in[pin++];

136 for (i=IDEA_ROUNDS;i>0;i--) {
137     word1 = mul(word1 , key [pk++]);
138     word2 = lsw16((char)word2 + (char)key [pk++]);
139     word3 = lsw16((char)word3 + (char)key [pk++]);
140     word4 = mul(word4 , key [pk++]);

142     t2 = (short)(word1 ^ word3);
143     t2 = mul(t2 , key [pk++]);
144     t1 = lsw16((char)t2 + (char)(word2 ^ word4));
145     t1 = mul(t1 , key [pk++]);
146     t2 = lsw16((char)t1 + (char)t2);

148     word1 ^= t1;
149     word4 ^= t2;

151     t2 ^= word2;
152     word2 = (short)(word3 ^ t1);
153     word3 = t2;
154     //System.out.println("Round "+i+" :"+Integer.toHexString((char)
155     word1)+" "+Integer.toHexString((char)word2)+" "+Integer
156     .toHexString((char)word3)+" "+Integer.toHexString((char)word4
157     ));
158 }

```

```

157 word1 = mul(word1, key [pk++]);
158 block_out [pout++] = word1;
159 block_out [pout++] = lsw16 ((char) word3 + (char) key [pk++]);
160 block_out [pout++] = lsw16 ((char) word2 + (char) key [pk++]);
161 word4 = mul(word4, key [pk]);
162 block_out [pout] = word4;
163 //System.out.println("Round 0:" + Integer.toHexString((char) word1)
      + " " + Integer.toHexString((char) word2) + " " + Integer.toHexString((
      char) word3) + " " + Integer.toHexString((char) word4));
164     }
166 }

```

9.1.13 IdeaTestSw.java

Listing 11: IdeaTestSw.java

```

1 public class IdeaTestSw {
2
3     public static void main(String [] args)
4     {
5
6         if (args.length < 1) {
7             System.out.println("Usage: java IdeaTestSw <integer>");
8             return;
9         }
10
11        int i, j;
12        int fd;
13        int data_in;
14        int data_out;
15
16        int [] param = new int [4];
17        int [] obj0 = new int [2];
18        int n, n64;
19
20        try {
21            n64 = (new Integer(args [0])).intValue ();
22        }
23        catch (NumberFormatException nfs) {
24            System.out.println("You must provide an integer");
25            return;
26        }
27        n = n64 * 2;
28
29        short [][] key = new short [3][8];
30        short [] enkey = new short [IdeaPlus.IDEA_SK_NUM];
31        short [] dekey = new short [IdeaPlus.IDEA_SK_NUM];
32        short [][] rdata = new short [n64][4];
33        short [][] rdata_out = new short [n64][4];
34
35        /*
36         =====
37         IDEA reference code
38         =====
39        */
40
41        for (i=0; i<8; i++)
42            key [0][i] = (short) (i+1);

```



```

44  for (i=0; i<8; i++)
45      System.out.print((Integer.toHexString(key[0][i]))+" ");

47  /* Test with only one key */
48  IdeaPlus.idea_encrypt_subkeys(key[0],enkey);
49  System.out.println();
50  for (i=0; i<8; i++)
51      System.out.print((Integer.toHexString((char)enkey[i]))+" ");
52  IdeaPlus.idea_decrypt_subkeys(enkey,dekey);
53  System.out.println();
54  for (i=0; i<8; i++)
55      System.out.print((Integer.toHexString((char)dekey[i]))+" ");

57  System.out.println();
58  System.out.println("Input blocks:");
59  for (j = 0; j < n64; j++) {
60      for (i=0 ; i<4; i++) {
61          rdata[j][i]= 0;//(short)(j*10 + i);
62          System.out.print((Integer.toHexString((char)rdata[j][i]))+"\t");
63      }
64      System.out.println();
65  }

67  /* Encryption starts */
68  /*if ((fd = open("/dev/fpga",O_RDONLY)) == -1 )
69      {
70          printf(" Can't open fpga");
71          exit(1);
72      }*/

74  /*if (ioctl(fd,FPGA_TIMER_INIT,1) != 1)
75      {
76          printf(" Error initalizing timer...\n");
77          return 1;
78      }*/

80  for (j = 0; j < n64; j++) {
81      IdeaPlus.idea_cipher(rdata[j], rdata_out[j], enkey);
82  }
83  System.out.println();
84  System.out.println("Output blocks:");
85  for (j = 0; j < n64; j++) {
86      for (i=0 ; i<4; i++)
87          System.out.print((Integer.toHexString((char)rdata_out[j][i]))+"\t"
88              );
89      System.out.println();
90  }

91  /*if (ioctl(fd,FPGA_TIMER_STOP,1) != 1)
92      {
93          printf(" Error stopping timer...\n");
94          return 1;
95      }*/
96  System.out.println("This is the SW time measured for encryption of "
97      +n64+" 64-bit blocks");

98  /* Encryption ends */

100 /*
101  =====
102  IDEA reference code END
103  =====

```

```
104  */
106  }
108 }
```

9.2 Bibliography

References

- [1] Joshua Engel, *Programming for the Java Virtual Machine*, Addison Wesley, second edition, 1999.
- [2] Java API, <http://java.sun.com/j2se/1.4.2/docs/api/>
- [3] BCEL API, <http://jakarta.apache.org/bcel/apidocs/index.html>
- [4] BCEL browsable source, <http://jakarta.apache.org/bcel/xref/index.html>
- [5] Miljan Vultic, Ludovic Righetti, Laura Pozzi, Paolo Ienne, *Portable Reconfigurable Coprocessors through Interface Virtualization*, EPFL.
- [6] Ludovic Righetti, Miljan Vuletic, *Operating System Support for Reconfigurable System on Chip*, EPFL semester project, 2003.
- [7] Christophe Dubach, Miljan Vuletic, *Java Virtual Machine on FPGA based platforms*, 2004.