



## SCIENTIFIC COMMUNICATION

Centre of Innovation in Telecommunications and Integration of services

### **JooFlux: Hijacking Java 7 InvokeDynamic To Support Live Code Modifications**

*Julien Ponge & Frédéric Le Mouél*

October 2012

Changing functional and non-functional software implementation at runtime is useful and even sometimes critical both in development and production environments. JooFlux is a JVM agent that allows both the dynamic replacement of method implementations and the application of aspect advices. It works by doing bytecode transformation to take advantage of the new `invokedynamic` instruction added in Java SE 7 to help implementing dynamic languages for the JVM. JooFlux can be managed using a JMX agent so as to operate dynamic modifications at runtime, without resorting to a dedicated domain-specific language. We compared JooFlux with existing AOP platforms and dynamic languages. Results demonstrate that JooFlux performances are close to the Java ones — with most of the time a marginal overhead, and sometimes a gain — where AOP platforms and dynamic languages present significant overheads. This paves the way for interesting future evolutions and applications of JooFlux.



University of Lyon  
INSA Lyon  
INRIA



# JooFlux: Hijacking Java 7 InvokeDynamic To Support Live Code Modifications

Julien Ponge and Frédéric Le Mouél

University of Lyon

INSA-Lyon, CITI-INRIA F-69621, Villeurbanne, France

`julien.ponge@insa-lyon.fr` and `frederic.le-mouel@insa-lyon.fr`

October 2, 2012

## Abstract

Changing functional and non-functional software implementation at runtime is useful and even sometimes critical both in development and production environments. JooFlux is a JVM agent that allows both the dynamic replacement of method implementations and the application of aspect advices. It works by doing bytecode transformation to take advantage of the new `invokedynamic` instruction added in Java SE 7 to help implementing dynamic languages for the JVM. JooFlux can be managed using a JMX agent so as to operate dynamic modifications at runtime, without resorting to a dedicated domain-specific language. We compared JooFlux with existing AOP platforms and dynamic languages. Results demonstrate that JooFlux performances are close to the Java ones — with most of the time a marginal overhead, and sometimes a gain — where AOP platforms and dynamic languages present significant overheads. This paves the way for interesting future evolutions and applications of JooFlux.

# 1 Introduction

Most applications are — to some degree — relatively static. This is especially true for software developed using statically typed languages. Most dynamically-typed languages offer the ability to replace methods at runtime, also coined as *monkey patching*. Even so, this is rarely possible to perform from outside the program: such changes can be performed by the program, not by some out of process management tool. Very few languages make it easy to do so from the outside, even purely functional languages with *read-eval-print loops*.

Dynamically changing code implementation at runtime has many virtues though [7, 9]. In development phases, a significant share of time is lost due to the need to restart applications to see the effect of source code changes. In production, applying an important security or bug fix requires the same restart procedure, this time impacting the deployed application availability and causing issues to end-user. A full restart can take several minutes to be validated and performed, while in such cases the fix may simply consist in replacing some method implementation without introducing any side-effect like changing methods signatures. Adding cross-cutting concerns, such as logging or security aspects, can also be mostly useful for capturing, monitoring and managing non-functional behaviours. The ability to dynamically apply and/or remove aspects in a program is especially appealing, again both in development and production contexts.

In the case of adaptive execution environments like a Java virtual machine, this however involves loosing potential optimizations, that can only be obtained and stabilized by a long-running process, leading to performance deterioration.

**Our proposal** This paper describes a new approach to dynamically patch and weave aspects to a Java application at run-time by transforming the whole application and making it dynamic.

Our contributions are as follows.

1. We designed and developed a Java agent that intercepts and modifies the bytecode at runtime to replace all method invocations by the new Java 7 `invokedynamic` bytecode.
2. We propose a control API available through a JMX agent that allows to dynamically manage the modifications at runtime without the need for a dedicated language or annotation/pointcut inserts.
3. We present a comparison between our JooFlux prototype, existing AOP platforms and dynamic programming languages. Results demonstrate that JooFlux performances are close to the Java ones — with most of the time a marginal overhead, and sometimes a gain — while AOP platforms and dynamic languages present significant overheads.

**Paper structure** This paper is structured as follows. We start with a contextual overview with brief recalls on aspect-oriented programming, the Java virtual machine and the new `invokedynamic` instruction. Next, we give details on how JooFlux works before comparing its performances with other platforms and programming languages with dynamic-dispatch. Finally, we discuss the related work and give perspectives for future work.

## 2 Context

Let us present the context into which this work was conducted. This section contains 3 parts. We start with an informal recall on aspect-oriented programming. We then continue with an overview of the Java Virtual Machine. Finally, we focus on a significant evolution made starting from Java SE 7 to ease the support of dynamic languages, and that is at the core of our approach.

### 2.1 Aspect-oriented programming

Aspect-oriented programming (AOP) aims at modularizing the cross-cutting concerns in programs with most of these being related to non-functional code, e.g. security, validation, logging or transactions [11]. Indeed, such code often needs to be repeated across several application layers, increasing the maintenance costs as both functional and non-functional code get mixed together.

AOP attempts to solve this problem as follows:

1. an *aspect* captures a cross-cutting concern as a single code unit, and
2. a *point-cut* specifies a set of precise points in a program, and
3. an *advice* materializes the application of an aspect at a point-cut.

As an example, suppose that we would like to log each invocation to a getter of the classes in package `foo`. To do that, an aspect would be the code performing the logging logic, the point-cut would be the set of calls to public methods of the classes in package `foo` whose names start with `get`. An AOP processor, also called a *code weaver*, would apply the aspect to every matching pointcut specified for the advice application.

In practice, the Eclipse AspectJ project is an established AOP framework for Java that features its own language [22]. Other approaches exist too [19, 5, 6]. An AspectJ source code unit can define aspects, point-cuts and advices. It can also define *inter-type declarations*, that is, the ability to add new field and method members to existing types.

Technically, AspectJ works by modifying compiled bytecode either statically by transforming `.class` files, or at runtime through the use of a dedicated class loader or JVM agent. In turn, aspects which are essentially “just Java code”, get compiled into separate classes. The AspectJ weaver then applies advices by inserting method calls from the bytecode matching a point-cut to the corresponding aspect. Back to our example, the logging aspect would be a compiled class, and each invocation to a matched getter method would be preceded with a invocation of the method holding the logging logic in the generated aspect.

## 2.2 The Java Virtual Machine

The Java Virtual Machine (JVM) specification has been very stable since its inception [13]. A JVM consumes bytecode that is traditionally enclosed into `.class` files that represent Java classes. The format of a compiled class was designed to mimic a Java source file: a compiled `.class` file contains a single Java class, a set of fields and methods. It also contains a *constant pool*, that is, a set of indexed constant values that can be referred to by number in the bytecode, thus reducing the compiled bytecode footprint by avoiding duplicates.

The execution model of the JVM lies around a stack. Opcodes may manipulate the stack, consume element as operands, and push new elements. As an example, invoking the `int java.lang.String::indexOf(String)` instance method consumes 2 elements from the stack: a `java.lang.String` instance to invoke the method followed by another instance that corresponds to the sole parameter. In return, it pushes a primitive `int` value as a return value.

Consider the following Java class:

```
public class Hello {
    public static void main(String... args) {
        System.out.println("Hello!");
    }
}
```

Using a decompiler tool such as `javap`, the bytecode corresponding to the `main()` method is as follows:

```
public static void main(java.lang.String...)
  flags: ACC_PUBLIC, ACC_STATIC, ACC_VARARGS
  Code:
    stack=2, locals=1, args_size=1
      0: getstatic      #2
      3: ldc            #3
      5: invokevirtual #4
      8: return
```

The `getstatic` opcode is used to lookup the `out` field in the `java.lang.System` class which was the second entry from the class constant pool. The `ldc` opcode then loads the third value: the "Hello" string. Finally, the `println(String)` virtual method is invoked before returning, with the method reference corresponding to the fourth entry in the constant pool.

While the Java bytecode class format was designed specifically with the requirements of the Java programming language in mind, it is still a general-purpose bytecode. In fact, the Java bytecode is strictly more expressive than its counterpart language, a property that has been widely exploited in applications like obfuscation [3].

The JVM does not limit itself to running applications written in Java. There is a vibrant ecosystem of programming languages targeting the JVM. This includes ports of previously-existing languages to the JVM (e.g., JRuby, Jython, Rhino) or original languages that appeared first on the JVM (e.g., Groovy, Scala, Clojure). Indeed, there has been extensive virtual machine optimization research supported by industrial validation that make the JVM a compelling target runtime for a language [15, 12, 8].

## 2.3 Java SE 7 and the new `invokedynamic` opcode

The JVM specification provides 4 opcodes for invoking methods. `invokestatic` is used for static methods. `invokevirtual` is used when dispatching shall be performed based on the receiver type. This corresponds to `public` and `protected` methods in Java as they may be overridden in subclasses. `invokespecial` is used to dispatch to the other types of methods such as `private` methods and constructors. Finally, `invokeinterface` is used to dispatch a method to a receiver instance implementing the said interface. This set of method invocation opcodes remained stable until the release of Java SE 7 when `invokedynamic` was introduced [20].

**A new opcode** The motivation for a new method invocation opcode was to make it easier for dynamic languages to be implemented on top of the JVM. Indeed, dynamic languages often need to resolve types, symbols and invocation targets at runtime. Using either of the previous method invocation opcodes, dynamic language implementors had to rely on reflection and dynamic proxy generations to delay such tasks to the runtime. This proved to be slow in practice, as just-in-time optimizations were rarely picked up efficiently by the virtual machines.

`invokedynamic` is very similar to `invokeinterface` and `invokevirtual` in the sense that the method dispatch is performed at runtime. However, both still require a receiver type, that is, a base class or interface declaring the target method signature. `invokedynamic` relaxes this and is closer to function pointers semantics as, say, in C. More specifically, `invokedynamic` opcodes work with:

1. a symbolic name to designate the invocation, and
2. a type signature for the parameters and return type, and
3. a bootstrap instruction that is invoked the first time that a given `invokedynamic` opcode is encountered (a call site).

**A runtime support API** The role of the bootstrap instruction is to bind the call site with a target to handle invocations. To do that, there is a new API found as part of the `java.lang.invoke` package, and which defines 2 useful types. `CallSite` represents a call site, and is the return value of the bootstrap process. It points to a `MethodHandle` which is either a direct reference to a class method or field, or a chain of method handles called *combinators* [20]. Here is an usage example of the new API:

```
public static MethodHandle replaceSpaces(Lookup lookup) throws Throwable {
    return insertArguments(lookup.findVirtual(String.class, "replaceAll",
        methodType(String.class, String.class, String.class)), 1, "%20", " ");
}

public static void main(String... args) throws Throwable {
    MethodHandle mh = replaceSpaces(lookup());
    System.out.println((String) mh.invokeExact("A%20B%20C%20"));
}
```

It shows how to obtain a method handle over the `String.replaceAll(String, String)` method. Once this is done, we insert an `insertArguments` combinator whose role is to pre-bind arguments to constant values. In this case we bind both arguments of `replaceAll()` so that when invoked, the method replaces occurrences of `"%20"` with `" "`. As this is a virtual method, the argument at index 0 is actually the receiver object. Finally, the method handle can be invoked, in this case printing `"A B C"`. Of course, a wider range of combinators exists, but the point is that method handles can be manipulated to point to concrete methods or some transformation / adaptation code as we will see later in this paper.

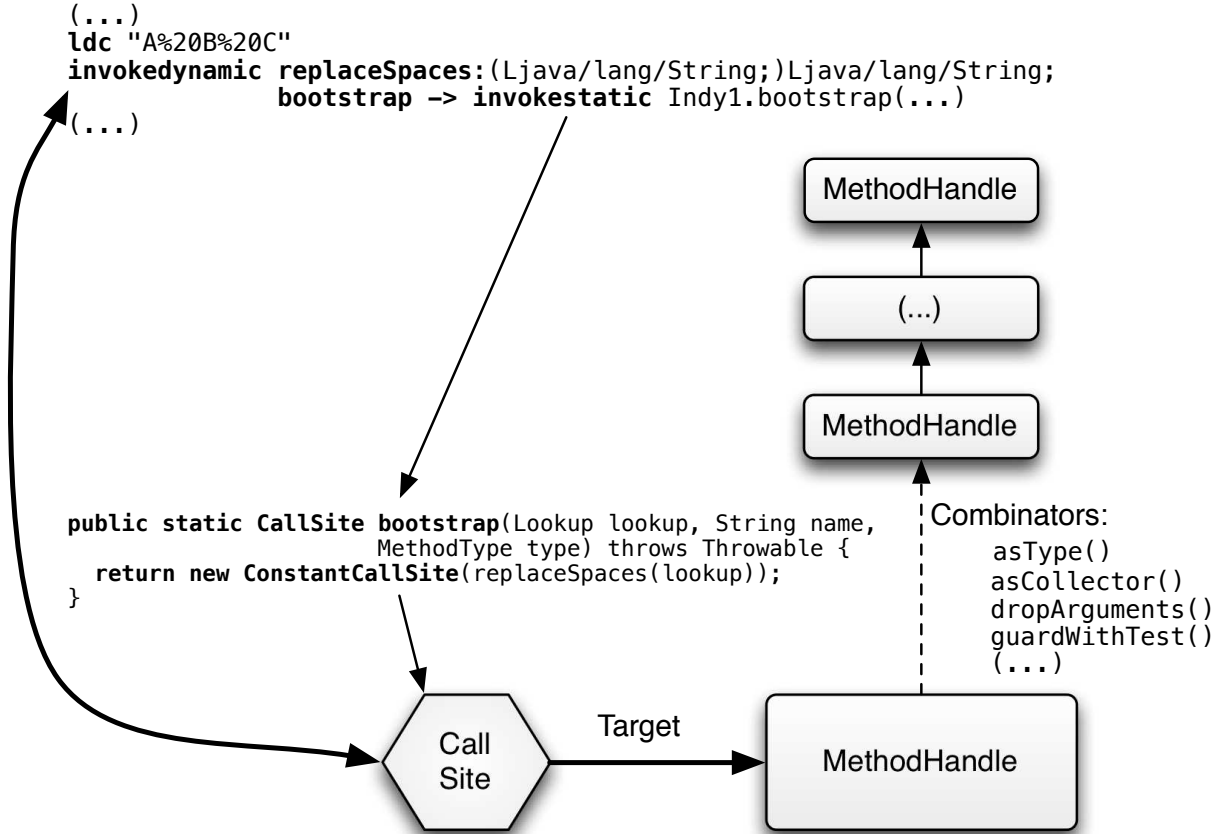


Figure 1: Bootstrapping an invokedynamic call site.

The link between an `invokedynamic` opcode and the bootstrap method is usually done by invoking a static method that obtains a method handle whose type matches those of the call site, then instantiates a `CallSite` and returns it. This is illustrated in Figure 1. The call site instance is only assigned once by invoking the bootstrap method the first time the invocation is made.

One of the main improvement in using `invokedynamic` is that combinator chains are assigned to identifiable call sites and have plugs into the JVM internals. This way, optimizations can be put in place more easily than with traditional reflective approaches [21]. Another point is that call sites may be re-targeted to new method handles, which is useful to, say, implement instance method redefinition at runtime as allowed in languages such as Python or Ruby. Also, method handles are type-checked only at creation time, unlike reflection objects that need to perform such checking for each invocation. This yields better method dispatching performance.

### 3 JooFlux

This section gives technical details on how JooFlux works. Specifically, we explain how it introduces a method call indirection through bytecode rewriting and the `invokedynamic` instruction. Then, we explain how aspect advices can be attached to methods, both before and after they are being executed.

Finally, we present the management layer of JooFlux that allows methods to be replaced, and aspect advices to be injected, all at runtime.

### 3.1 Introducing a method call indirection

JooFlux works by introducing an indirection on method invocations, so that method replacement and application of aspects can be performed at runtime. While JooFlux primarily focuses on bytecode emitted by a compiler for the Java programming language, it can theoretically work with any valid JVM bytecode produced by another language such as Scala [14].

In JooFlux, we took the approach of taking advantage of `invokedynamic`, as the possibility of dynamically rebinding call sites to new method handles chains effectively permits method implementations to be changed. Also, the range of method handle combinators includes what to perform additional processing on both invocation arguments and return values. As we will see, this effectively allows us to implement aspect-oriented programming.

As of version 7, Java (the language) does not rely on `invokedynamic`. Compiled Java bytecode perform method invocations using the original `invokestatic`, `invokevirtual`, `invokespecial` and `invokeinterface` opcodes.

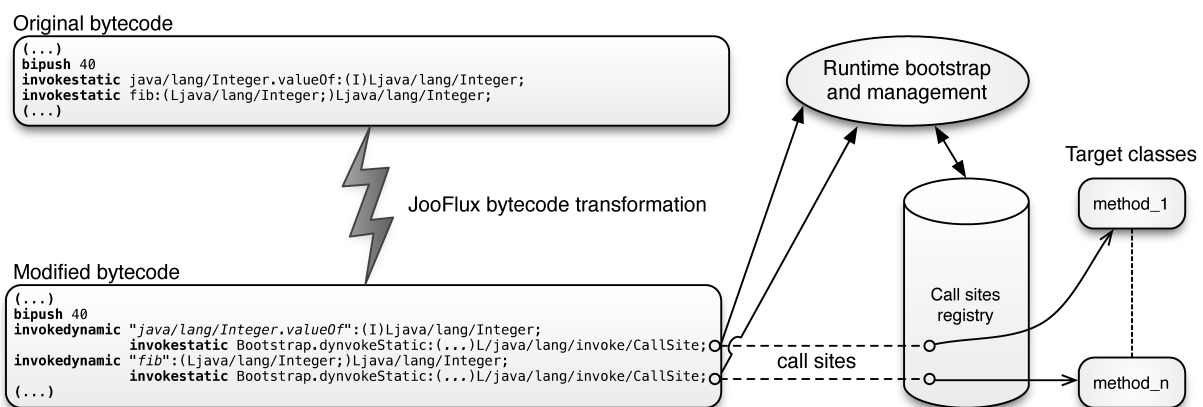


Figure 2: Overview of the JooFlux JVM agent.

Java virtual machines offer the possibility to attach *agents*. Once plugged into a JVM, an agent is able to perform many operations, including the ability to define itself as a bytecode transforming agent. By doing so, an agent can intercept bytecode as it is being loaded into the JVM, and it can rework it by adding and removing instructions.

JooFlux works as a JVM agent for which Figure 2 gives an overview. When classes are being loaded, it looks for occurrences of `invokestatic`, `invokevirtual`, `invokespecial` or `invokeinterface` and replaces them by a semantic-preserving `invokedynamic` instruction. By doing so, original call sites get bound at runtime to method handles. The bytecode transformations are being made using the ASM library<sup>1</sup>.

Initially, the call sites point to method handles to the original classes methods. For instance, when a `invokestatic` instruction is in the bytecode, it is replaced by a `invokedynamic` instruction whose symbolic name is based on the original `invokestatic` target signature. This is useful to have a uniform naming scheme when we later want to perform replacements and aspect injection. The type of the original target is preserved too. The JooFlux runtime call site bootstrap class provides several static methods depending on the original invocation type, i.e. static, virtual, special and interface. By doing so, we only introduce a thin indirection layer.

### 3.2 Aspect advices using method handle combinators

In existing tools such as AspectJ, an aspect advice is injected by adding new instructions to the original bytecode. For instance, an AspectJ advice whose pointcut is before a method invocation will be applied

<sup>1</sup>ASM 4.0: <http://asm.ow2.org/>

just before the said method call site. To do so, a first static method is invoked to fetch the aspect class instance, then a second method invocation is performed on it to the method that corresponds to the advice. This can be easily checked by decompiling bytecode. Other tools such as JBoss Byteman work in a similar fashion, while in the case of the later new rules / advices can be injected and removed at runtime by reloading class definitions. This is allowed by the JVM as long as the reloaded classes only change method implementations but do not add or remove methods and fields.

JooFlux does not require adding new instructions into the transformed bytecode. Also, it does not require reloading classes as new advices are being injected, or when method implementations are being replaced. Instead, it mainly relies on 2 method handle combinators found in the `java.lang.invoke.MethodHandles` class:

1. `filterArguments` takes a target method handle, an argument position and an array of filter function method handles, and
2. `filterReturnValue` takes a target method handle and a filter function method handle.

The filter method handle function types must match those of their target methods. In our case, we opted for a generic approach that would fit all kinds of target methods. Hence, JooFlux advices work with arrays of objects matching the arguments when intercepting invocations, and they use objects when intercepting returns. The following is an example class providing 2 static methods that can be used as logging advices:

```
public class Dumpers {
    public static Object[] onCall(Object[] args) {
        System.out.println(">>> " + Arrays.toString(args));
        return args;
    }
    public static Object onReturn(Object retval) {
        System.out.println("<<< " + retval);
        return retval;
    }
}
```

In this example the advice methods simply log the invocation arguments and return values. They could as well do other things such as throwing an exception if some argument value is invalid. Also, because they act as *filters*, they can modify argument values before they are given to the target, and they can modify the return value before the invocation client obtains it.

Advices can be stacked, that is, several of them may be applied to a given target method. For instance, we could stack a validation and a logging advice to be executed before a method is being invoked. What an advice can do also depends on the target method original invocation mode. In the case of static methods, the arguments array simply contains the values corresponding to the static method signature. However in the case of instance methods, the first argument is the receiver, that is, the instance of the object that the method shall be invoked on.

Finally, we mentioned that target and filter method handles must match. However we opted for a generic solution that can work with any method signature. To make this work, we took advantage of 2 further method handle combinators. The first one is `asSpreader`, which can wrap some arguments into an array. The second one is `asCollector`, which performs the reverse operation by mapping the values of an array to parameters. For method invocation interception, we first wrap all arguments into a `Object[]` array using `asSpreader`. We can then pass it to the advice method, and extract the filtered values from the resulting array back to the target method parameters using `asCollector`. The case of method return interception is simpler as it does not require using such combinator chain. Instead, we just need to loosen the return value type to `Object` before passing it to the advice, then narrow it again to its original type before returning to the method invoker. Such type transformations happen with the `asType` combinator.



### 3.3 Managing JooFlux: live method replacement and aspects injection

Call sites are being put in a central registry at bootstrap time. This registry sports no role while application code is being executed, and there is also no lookup penalty in having it. It is used by the management layer of JooFlux in several ways. To do so, JooFlux offers a JMX<sup>2</sup> agent for querying and interacting with itself.

This agent offers the following set of remotely-accessible operations:


1. replacing method implementations by changing their call site targets, and
2. plugging an aspect advice before or after certain call sites, and
3. querying for various metrics, including the number of call sites monitored by JooFlux and a list of them.

As an example, let us see the interaction with the method replacement operation, whose signature as a JMX interface method is

```
void changeCallSiteTarget(String methodType,  
                          String oldTarget, String newTarget)
```

The `methodType` parameter specifies the type of method invocation in the original bytecode: static, virtual, interface and special. The last 2 parameters specify an identifier for a call site to be replaced and a new method handle to use as a target.

Before button handler replacement



After button handler replacement



JMX console (*jconsole*)

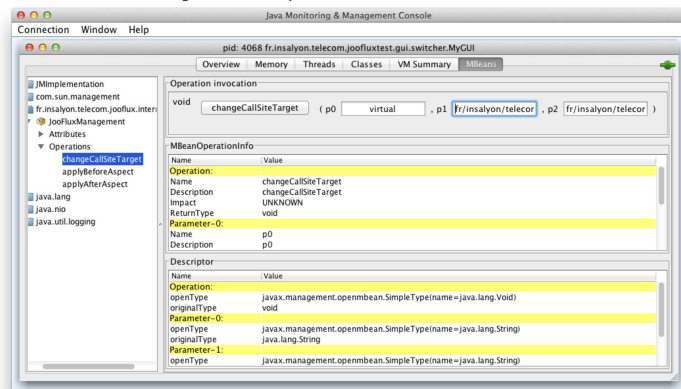


Figure 3: Interacting with an application through the JooFlux JMX agent.

Figure 3 features a graphical application with a button on the left pane and a text label on the right pane. By connecting to the JooFlux JMX agent of the application using a tool such as *jconsole*, we can replace the button action handler so that the next time it is clicked, a picture replaces the label. In that specific demonstration case, we used the following input to the `changeCallSiteTarget` JMX operation:

```
virtual, // virtual method  
fr/insalyon/telecom/joofluxtest/gui/switcher/  
    MyActionListener.counterIncrement:(MyActionListener)void, // handler id  
fr/insalyon/telecom/joofluxtest/gui/switcher/  
    MyActionListener.pictureSwitch:()V // our new handler
```

<sup>2</sup>Java Management Extensions: <http://jcp.org/aboutJava/communityprocess/final/jsr003/index3.html>

Injecting an aspect into an existing call site works in a similar fashion:

```
void applyBeforeAspect(String callSitesKey,
                      String aspectClass, String aspectMethod)

void applyAfterAspect(String callSitesKey,
                      String aspectClass, String aspectMethod)
```

## 4 Experimental comparison with AOP platforms and dynamic languages

To assess our approach, we compare the performance of the Java Virtual Machine with our JooFlux agent to other dynamic method call dispatching approaches: AOP platforms (section 4.1) and dynamic programming languages (sections 4.2-4.3).

All these approaches are JVM-based and have been tested with micro- and macro-benchmarks on a MacBook Pro 2,3 GHz Intel Core2 (i5), 4 Go 1333 MHz DDR3, running Mac OS X Lion 10.7.4 (11E53) and the OpenJDK Runtime Environment 1.7.0-u10-b06-20120906<sup>3</sup>. These tests were renewed 10 times to constitute the result set. Quartiles and median overhead were calculated on this result set.

### 4.1 Micro-benchmarks and AOP platforms

We compare JooFlux aspect injection functionality with two other AOP platforms: AspectJ<sup>4</sup>[10] and Byteman<sup>5</sup>[6]. We use a classic recursive Fibonacci micro-benchmark (`classicfibo`)<sup>6</sup> where we inject an empty aspect - i.e. a redirection call to an empty method - before or/and after the Fibonacci method. Results are presented in Table 1.

Execution times of the AspectJ generated bytecodes are close to the JVM ones, but it weaves the aspects statically during the compilation and no modification can be performed at runtime. Byteman allows the aspect injection at runtime by unloading the class, and modifies the bytecode during the reloading. This technique degrades significantly the performance and implies that all optimizations achieved by the JIT are lost. The bytecode generated by JooFlux already includes dynamic calls, so the cost of weaving an aspect is just the cost of including a new combinator in the method handle chains, recopying the arguments and transferring the return value. By keeping the method handle chain almost intact, JooFlux protects the JIT optimizations.

### 4.2 Micro-benchmarks and dynamic languages

Several programming languages propose the dynamic method dispatch directly in the language, e.g. with reflective APIs. To test the JooFlux dynamic method dispatch functionality, we compare it to the main JVM-based programming languages: Java<sup>7</sup>, Clojure<sup>8</sup>, JRuby<sup>9</sup>, Groovy<sup>10</sup>, Rhino JavaScript<sup>11</sup> and Jython<sup>12</sup>.

We first use the same Fibonacci micro-benchmark to test the performance. According to the programming languages, the Fibonacci method can be written in different manner: the `classicfibo` that manipulates objects, the `fastfibo` that manipulates long-typed parameters, the `fastestfibo` that manipulates long-typed parameters and return long-typed results and finally the `reflectivefibo` where

<sup>3</sup>OpenJDK Runtime Environment 1.7.0-u10-b06-20120906: <http://code.google.com/p/openjdk-osx-build/>

<sup>4</sup>AspectJ 1.7.1: <http://www.eclipse.org/aspectj/>

<sup>5</sup>Byteman 2.1.0: <http://www.jboss.org/byteman/>

<sup>6</sup>Recursive Fibonacci: [http://en.wikipedia.org/wiki/Recursion\\_\(computer\\_science\)#Fibonacci](http://en.wikipedia.org/wiki/Recursion_(computer_science)#Fibonacci)

<sup>7</sup>Java 7: <http://docs.oracle.com/javase/7/docs/api/>

<sup>8</sup>Clojure 1.4.0: <http://clojure.org>

<sup>9</sup>JRuby 1.6.7.2 (-indy), 1.7.0.preview2 (+indy): <http://jruby.org>

<sup>10</sup>Groovy 2.0.2 (-+indy): <http://groovy.codehaus.org>

<sup>11</sup>Rhino Javascript 1.7R4: <https://developer.mozilla.org/en-US/docs/Rhino>

<sup>12</sup>Jython 2.5.3: <http://www.jython.org>

Exec. Plat.	Weav. Time	Impl.	Q1-min	Q2-25%	Q3-median	Q4-75%	Q5-max	Overhead
JVM	-	classicfibonacci	611	613	615	616	636	-
JVM + AspectJ	compilation	classicfibonacci + before aspect	609	613	620	621	650	+0,8%
		classicfibonacci + after aspect	609	611	620	621	693	+0,8%
		classicfibonacci + before aspect & after aspect	609	613	620	622	717	+0,8%
JVM + JooFlux Agent	runtime	classicfibonacci + before aspect	3007	3028	3051	3070	3525	+396%
		classicfibonacci + after aspect	1269	1273	1280	1295	1356	+108%
		classicfibonacci + before aspect & after aspect	3804	3830	3861	3893	3950	+528%
JVM + ByteMan	runtime	classicfibonacci + before aspect	79309	84998	85371	85695	86397	+13781%
		classicfibonacci + after aspect	82561	82780	82911	83040	83406	+13381%
		classicfibonacci + before aspect & after aspect	179984	188088	188187	188321	188971	+30499%

Table 1: Fibo(40) micro-benchmarks for AOP platforms (in ms)

each method invocation use the reflective API. These different implementations are important because they influence the generated bytecode and hence performance.

Java results are presented in Table 2. JooFlux overhead is insignificant for `classicfibonacci` and presents a slowing factor of 2 for `reflectivefibonacci`. This implementation is indeed our worst case as the JVM can hardly inline the method calls.

P. Lang.	Exec. Plat.	Impl.	Q1-min	Q2-25%	Q3-median	Q4-75%	Q5-max	Overhead
Java	JVM	classicfibonacci	611	613	615	616	636	-
		reflectivefibonacci	1758	1762	1782	1803	4121	-
Java	JVM + JooFlux Agent	classicfibonacci	611	613	616	618	690	+0,001%
		reflectivefibonacci	3668	3686	3717	3743	4273	+108%

Table 2: Fibo(40) micro-benchmarks for Java programming language (in ms)

The others programming languages results are presented in Table 3. The different `classicfibonacci`, `fastfibonacci` and `fastestfibonacci` overheads are calculated compared to JVM+JooFlux `classicfibonacci`. `reflectivefibonacci` implementations are still compared to JVM+JooFlux `reflectivefibonacci` one.

Even if Clojure strongly-typed implementations presents only a slowdown factor of 1.2-1.4, most other languages are from 3 to 18 times slower than our Java+JooFlux prototype. JRuby and Groovy languages propose earlier versions using the `invokedynamic` bytecode (*+indy, figures in italic in the Table 3*) but even if they gain a slowing factor, they remain significantly slower than Java+JooFlux.

For AOP and dynamic languages Fibonacci micro-benchmarks, our JooFlux agent rewrites the bytecode in 75-100ms for a transformation of 2 classes (`Fibonacci` and `InvokeBootstrap` - the bootstrap itself is dynamic) and 1 method (`classicfibonacci`). This delay is introduced before the application launch, so it is quite insignificant and comparable to other platforms launching times.

### 4.3 CPU-intensive and interception/rewrite-intensive macro-benchmarks

To test more intensively the JooFlux dynamic method dispatch, we have used 3 macro-benchmarks:

P. Lang.	Exec. Plat.	Impl.	Q1-min	Q2-25%	Q3-median	Q4-75%	Q5-max	JooFlux Diff
Clojure	JVM	fastestfibo	722	732	734	740	742	+19%
		fastfibo	859	862	864	875	892	+40%
		classicfibo	4105	4118	4171	4265	4326	+577%
JRuby	JVM (+indy)	classicfibo	6290 (3982)	6333 (4006)	6382 (4020)	6486 (4069)	7043 (4323)	+936% (+552%)
		reflectivefibo	10226 (7545)	12020 (7561)	12060 (7581)	12076 (7621)	12288 (7737)	+224% (+104%)
Groovy	JVM (+indy)	fastestfibo	1383 (3061)	1388 (3077)	1394 (3092)	1401 (3150)	1417 (3165)	+126% (+402%)
		fastfibo	2709 (2513)	2721 (2519)	2725 (2528)	2749 (2540)	2766 (2583)	+342% (+310%)
		classicfibo	8660 (4461)	8691 (4488)	8716 (4522)	8726 (4584)	9066 (4656)	+1315% (+634%)
		reflexivefibo	57734 (8366)	57892 (8378)	58009 (8386)	58182 (8405)	58364 (8697)	+1460% (+125%)
Javascript	JVM	classicfibo	9052	9208	11275	11441	11764	+1730%
Jython	JVM	classicfibo	29053	29258	29675	30202	31871	+4717%

Table 3: Fibo(40) micro-benchmarks for JVM-based dynamic languages (in ms)

1. SCImark 2.0<sup>13</sup>: SciMark 2.0 is a CPU- and memory-intensive Java benchmark for scientific and numerical computing. It measures several computational kernels and reports a composite score in approximate Mflops.
2. A parallel file wordcounter with Fork/Join [17]: This file wordcounter is a memory- and IO-intensive Java benchmark that looks down a repository hierarchy and counts words inside each file.
3. The execution of Clojure runtime over JooFlux: We have used the Clojure language runtime – written in Java – to have an Interception-intensive benchmark.

**SCImark 2.0** The benchmark performs Fast Fourier Transformations (FFT), Jacobi Successive Over-relaxation (SOR), Monte Carlo integration (MC), Sparse matrix multiply (SM), dense LU matrix factorization (LU) and computes a composite score (CS). These calculations can be done with a small- or large-memory data set.

Performance results with and without-JooFlux are presented in Table 4. Almost all calculations have a positive or negative marginal overhead, except the Monte Carlo integration that loses 20% performance. The Monte Carlo algorithm exercises random-number generators, synchronized function calls, and function inlining that can explain the JIT difficulties.

	FFT		SOR		MC	SM		LU		CS	
	2 <sup>10</sup>	2 <sup>20</sup>	10 <sup>2</sup> ×10 <sup>2</sup>	10 <sup>3</sup> ×10 <sup>3</sup>	-	N=10 <sup>3</sup> nz=5.10 <sup>3</sup>	N=10 <sup>5</sup> nz=10 <sup>6</sup>	10 <sup>2</sup> ×10 <sup>2</sup>	10 <sup>3</sup> ×10 <sup>3</sup>	Small	Large
JVM	827	178,7	1196,5	1077,6	635,4	1180,8	1227,8	2467,8	1490,2	1261,5	915,4
JVM + JooFlux	824,3	178,5	1193,2	1080,3	507,2	1172,9	1249,5	2325,9	1489,7	1204,9	900,2
Perf.	-0,3%	-0,1%	-0,3%	+0,3%	-20%	-0,7%	+1,8%	-5,8%	-0,03%	-4,5%	-1,7%

Table 4: SCImark2.0 small and large macro-benchmarks (in Mflops)

**Parallel file wordcounter with Fork/Join** We apply the file wordcounter to the HotSpotVM source repository (2.10<sup>3</sup> files) with a mono-thread process and with a 2-threads on 2-cores parallel execution.

<sup>13</sup>SCImark 2.0: <http://math.nist.gov/scimark2/>

As shown in Table 5, surprisingly, the bytecode modified by JooFlux betters the performances. This result demonstrates that synchronized method calls are not affected by the `invokedynamic` indirection and the JIT can even apply a better inlining.

Exec. Plat.	Impl.	Q1-min	Q2-25%	Q3-median	Q4-75%	Q5-max	Overhead
JVM	wordcounter single thread	3707	3727	3733	3752	4769	-
	wordcounter fork/join 2 threads	1826	1934	2161	2235	5949	-
JVM + JooFlux Agent	wordcounter single thread	3634	3646	3658	3689	4829	-2%
	wordcounter fork/join 2 threads	1825	1916	2029	2070	2407	-6%

Table 5: Parallel file wordcounter macro-benchmarks for Java programming language (in ms)

**Clojure over JooFlux** The Clojure language runtime is a quite large Java application ( $3.10^3$  classes) and we use it as an interception-intensive benchmark. We have retested the Fibonacci micro-benchmark, but this time with a total Clojure rewritten bytecode: 1325 transformed class, 26866 transformed methods and 19646 initial interceptions. We do not log the total number of dynamic method call interceptions as continuous loggers dramatically slow down the performance.

Table 6<sup>14</sup> presents the results. The overhead introduces is still insignificant compared to the execution time. We can say so that having a dynamic Clojure is effectively nice<sup>15</sup> and has no cost at runtime ! The 4s bytecode transformation time is not any more negligible but this cost is only applied one time at launch time and it's a low price to let the entire application becomes dynamic !

P. Lang.	Exec. Plat.	Impl.	Q1-min	Q2-25%	Q3-median	Q4-75%	Q5-max	Overhead
Clojure	JVM	fastestfibo	701	702	704	708	717	-
		fastfibo	839	845	848	851	855	-
		classicfibo	3969	3978	3984	4010	4093	-
Clojure	JVM + JooFlux Agent	fastestfibo	695	699	702	704	713	-0,3%
		fastfibo	856	861	864	868	891	+1,9%
		classicfibo	3957	3969	3980	4002	4069	-0,1%

Table 6: Fibo(40) for Clojure over a JVM+JooFlux execution plate-form (in ms)

## 5 Related work

**Software dynamic updates** The ability to dynamically update software at runtime is anything but a new idea, and it has been investigated by many research works [7, 9]. In fact, the techniques to be leveraged vastly dependent on target the programming language semantics and runtime. Dynamic software updates has been investigated as low as the operating system kernel level in works such as [2] and [16]. Closer to the context of applications running on a Java virtual machines, the work in [24] studies applying modifications of Java code at runtime. The approach supports modification of classes, including adding and removing fields and methods. However, this requires the usage of a modified virtual machine. The same authors continued their approach based on virtual machine modifications in [23], and applied it to dynamic aspect-oriented programming. It should also be mentioned that class definitions may be reloaded in the JVM. This is what debuggers do, and JVM agents such as the one provided

<sup>14</sup>Fibo(40) in Clojure does not have exactly the same performance in Table 6 and Table 3. To allow a good comparison, we prefer presenting Clojure+JVM results tested at the same time as Clojure+JVM+JooFlux, in the same test run.

<sup>15</sup>Why Clojure doesn't need `invokedynamic`, but it might be nice:  
<http://blog.fogus.me/2011/10/14/why-clojure-doesnt-need-invokedynamic-but-it-might-be-nice/>

by JBoss ByteMan leverages this mechanism. This works only with stable class definitions though, and accumulated JIT optimizations get lost on reload.

**Aspect-oriented programming** Many works focus on the design and implementation of aspect-oriented programming. AspectJ is a well-known tool comprising a compiler and language [10]. While static, it allows for efficient bytecode transformation when weaving aspects, and offers a fine-grained language to define pointcuts. Javassist is another proposal to do AOP in Java [5]. Some works focus on the ability to weave aspects in a dynamic fashion [19, 18]. JBoss ByteMan is particularly interesting as it features a *event-condition-action* rule language [6]. Also, it can attach its JVM agent to already running applications as well as apply or remove rules dynamically. The work in [4] further illustrates that providing modified virtual machines is a sound solution to support a complete range of modification operations to running Java applications. This concurs with the results in [23].

**Virtual machines** The field of virtual machines is no short of challenges. The JVM is an attractive platform for research given its open specification [13]. Many works focus on improving the JVMs performance [15, 12, 8]. The design of the JVM has an initial bias towards favouring statically typed languages. The rise of interest in dynamic languages running on the JVM yield to the `invokedynamic` instruction and the `java.lang.invoke` API [20]. This facilitates the implementation of such languages, and it also gives efficient plugs into the JVM internals so as to benefit from adaptive optimizations [21]. New usages of `invokedynamic` are starting to appear [1], and it is poised to serve as an implementation technique for the support of *lambdas* in Java 8<sup>16</sup>.

## 6 Conclusions and perspectives

We conclude this paper with some perspectives envisioned at this early stage of the project, and give pointers to obtain a copy of our prototype.

### 6.1 Conclusion

This paper introduced JooFlux, a JVM agent that allows both the dynamic replacement of method implementations and the application of aspect advices. Compared to existing approaches, JooFlux takes a novel route by taking advantage of the new `invokedynamic` instruction added in Java SE 7. To do so, it performs bytecode rewriting by introducing an `invokedynamic`-based indirection on method invocations, which is later used to dynamically modify the intercepted call sites. As shown in the earlier micro-benchmarks, the runtime overhead of JooFlux is marginal for method invocations, and fairly limited when aspects are being injected. In any case, JooFlux shows interesting performance compared to related approaches such as AOP tools or dynamic languages that rely on dynamic dispatch. More interestingly, JooFlux does not involve reloading whole classes on either method replacement or advice injection, which keeps a large range of just-in-time compilation optimizations valid. It does not require a specifically-tailored virtual machine. Also, the thin indirection layer introduced by JooFlux does not require lookups or guard checks for call site invalidation, greatly helping virtual machines in adaptive optimization work. Finally, JooFlux works in a transparent fashion by directly operating at the method call site level. As such, it does not require a dedicated language to specify where and what changes to apply.

### 6.2 Future work

JooFlux is currently a research prototype that demonstrates how `invokedynamic` can be cleverly used for other purposes than implementing dynamic languages on the JVM. While it demonstrates and validates our approach with small testing code bases, it still needs to be tested on a wider range of applications running on the JVM. The impact of the bytecode manipulation that we perform is significant and we are well aware that it can break some bytecode constructions. As more testing is being made, we hope to iron out the large share of possible corner cases. By doing so, we intend to turn JooFlux into a tool that can be adopted at a larger scale than what our initial version allows.

---

<sup>16</sup>See <http://openjdk.java.net/projects/lambda/>

Given the promising performance figures exhibited by JooFlux, we intend to study its application in specialized contexts such as resource control, multi-tenant architectures or dynamic modular applications for the *Internet of Things* without the need to rely on dedicated middleware platforms such as OSGi<sup>17</sup>.

In terms of general purpose features, we intend to introduce modification transactions and rollback. The ability to statically verify the applicability of modifications is appealing. Also, we need to improve the tooling to make it more convenient than a JMX agent. A secure remote shell interface through SSH would be especially convenient, as it would also support new code remote transport over tunnelling.

### 6.3 Availability

JooFlux is available as an open source project at <https://github.com/dynamid/jooflux>. It is published under the terms of the *Mozilla Public License Version 2.0*<sup>18</sup>. The prototype version of JooFlux used while writing this paper corresponds to the `r0` annotated tag of the corresponding Git repository. The source code contains build and testing instructions, and it also contains some instructions for reproducing demonstrations such as the one related in Figure 3. It also contains the scripts or instructions to reproduce the experiments made in Section 4.

As further developments are being made, the range of features described in this paper may have changed, possibly also impacting performance figures. We encourage the wider researchers and practitioners community to report any issue with JooFlux, and contribute bug fixes and improvements.

## References

- [1] Malte Appeltauer, Michael Haupt, and Robert Hirschfeld. Layered method dispatch with IN-VOKEDYNAMIC: an implementation study. In *Proceedings of the 2nd International Workshop on Context-Oriented Programming*, COP'10, pages 4:1–4:6, New York, NY, USA, 2010. ACM.
- [2] Jeff Arnold and M. Frans Kaashoek. Ksplice: automatic rebootless kernel updates. In *Proceedings of the 4th ACM European conference on Computer systems*, EuroSys '09, pages 187–198, New York, NY, USA, 2009. ACM.
- [3] Michael Batchelder and Laurie Hendren. Obfuscating Java: the most pain for the least gain. In *Proceedings of the 16th International Conference on Compiler Construction*, CC'07, pages 96–110, Berlin, Heidelberg, 2007. Springer-Verlag.
- [4] Christoph Bockisch, Michael Haupt, Mira Mezini, and Klaus Ostermann. Virtual machine support for dynamic join points. In *Proceedings of the 3rd International Conference on Aspect-Oriented Software Development*, AOSD '04, pages 83–92, New York, NY, USA, 2004. ACM.
- [5] Shigeru Chiba and Muga Nishizawa. An easy-to-use toolkit for efficient Java bytecode translators. In *Proceedings of the 2nd International Conference on Generative Programming and Component Engineering*, GPCE'03, pages 364–376, New York, NY, USA, 2003. Springer-Verlag New York, Inc.
- [6] Andrew E. Dinn. Flexible, dynamic injection of structured advice using byteman. In *Proceedings of the 10th International Conference on Aspect-Oriented Software Development companion*, AOSD'11, pages 41–50, New York, NY, USA, 2011. ACM.
- [7] Ophir Frieder and Mark E. Segal. On dynamically updating a computer program: from concept to prototype. *J. Syst. Softw.*, 14(2):111–128, February 1991.
- [8] Christian Häubl and Hanspeter Mössenböck. Trace-based compilation for the Java HotSpot virtual machine. In *Proceedings of the 9th International Conference on Principles and Practice of Programming in Java*, PPPJ'11, pages 129–138, New York, NY, USA, 2011. ACM.
- [9] Michael Hicks and Scott Nettles. Dynamic software updating. *ACM Trans. Program. Lang. Syst.*, 27(6):1049–1096, November 2005.

---

<sup>17</sup>See <http://www.osgi.org/>

<sup>18</sup>See <http://www.mozilla.org/MPL/2.0/>

- [10] Gregor Kiczales, Erik Hilsdale, Jim Hugunin, Mik Kersten, Jeffrey Palm, and William G. Griswold. Getting started with ASPECTJ. *Commun. ACM*, 44(10):59–65, 2001.
- [11] Gregor Kiczales, John Lamping, Anurag Mendhekar, Chris Maeda, Cristina Lopes, Jean marc Loingtier, and John Irwin. Aspect-oriented programming. In *Proceedings of the 11th European Conference on Object-Oriented Programming*, ECOOP’97, pages 220–242, Jyväskylä, Finland, 1997. SpringerVerlag.
- [12] Thomas Kotzmann, Christian Wimmer, Hanspeter Mössenböck, Thomas Rodriguez, Kenneth Russell, and David Cox. Design of the Java HotSpot client compiler for Java 6. *ACM Trans. Archit. Code Optim.*, 5(1):7:1–7:32, May 2008.
- [13] Tim Lindholm and Frank Yellin. *Java Virtual Machine Specification*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2nd edition, 1999.
- [14] Martin Odersky and Matthias Zenger. Independently Extensible Solutions to the Expression Problem. In *Proceedings of 12th International Workshop on Foundations of Object-Oriented Languages*, January 2005. <http://homepages.inf.ed.ac.uk/wadler/fool>.
- [15] Michael Paleczny, Christopher Vick, and Cliff Click. The java hotspot™ server compiler. In *Proceedings of the 2001 Symposium on Java™ Virtual Machine Research and Technology Symposium - Volume 1*, JVM’01, pages 1–1, Berkeley, CA, USA, 2001. USENIX Association.
- [16] Nicolas Palix, Gaël Thomas, Suman Saha, Christophe Calvès, Julia Lawall, and Gilles Muller. Faults in linux: ten years later. In *Proceedings of the 16th International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS ’11, pages 305–318, New York, NY, USA, 2011. ACM.
- [17] Julien Ponge. Fork and Join: Java Can Excel at Painless Parallel Programming Too!, July 2011. Oracle Technology Network, Oracle Corporation.
- [18] Andrei Popovici, Gustavo Alonso, and Thomas Gross. Just-in-time aspects: efficient dynamic weaving for Java. In *Proceedings of the 2nd International Conference on Aspect-Oriented Software Development*, AOSD ’03, pages 100–109, New York, NY, USA, 2003. ACM.
- [19] Andrei Popovici, Thomas Gross, and Gustavo Alonso. Dynamic weaving for aspect-oriented programming. In *Proceedings of the 1st International Conference on Aspect-Oriented Software Development*, AOSD’02, pages 141–147, New York, NY, USA, 2002. ACM.
- [20] John R. Rose. Bytecodes meet combinators: invokedynamic on the JVM. In *Proceedings of the 3rd Workshop on Virtual Machines and Intermediate Languages*, VMIL’09, pages 2:1–2:11, New York, NY, USA, 2009. ACM.
- [21] Christian Thalinger and John Rose. Optimizing invokedynamic. In *Proceedings of the 8th International Conference on the Principles and Practice of Programming in Java*, PPPJ’10, pages 1–9, New York, NY, USA, 2010. ACM.
- [22] John Viega and Jeffrey Voas. Can Aspect-Oriented Programming Lead to More Reliable Software? *IEEE Softw.*, 17(6):19–21, November 2000.
- [23] Thomas Würthinger, Danilo Ansaloni, Walter Binder, Christian Wimmer, and Hanspeter Mössenböck. Safe and atomic run-time code evolution for Java and its application to dynamic AOP. In *Proceedings of the 2011 ACM International Conference on Object Oriented Programming Systems Languages and Applications*, OOPSLA ’11, pages 825–844, New York, NY, USA, 2011. ACM.
- [24] Thomas Würthinger, Christian Wimmer, and Lukas Stadler. Dynamic code evolution for Java. In *Proceedings of the 8th International Conference on the Principles and Practice of Programming in Java*, PPPJ ’10, pages 10–19, New York, NY, USA, 2010. ACM.