

Termination in Language-based Systems

Algis Rudys
arudys@cs.rice.edu

John Clements
clements@cs.rice.edu

Dan S. Wallach
dwallach@cs.rice.edu

Department of Computer Science, Rice University

Abstract

Language runtime systems are increasingly being embedded in systems to support runtime extensibility via mobile code. Such systems raise a number of concerns when the code running in such systems is potentially buggy or untrusted. While sophisticated access controls have been designed for mobile code and are shipping as part of commercial systems such as Java, there is no support for terminating mobile code short of terminating the entire language runtime. This paper presents a concept called “soft termination” which can be applied to virtually any mobile code system. Soft termination allows mobile code threads to be safely terminated while preserving the stability of the language runtime. In addition, function bodies can be permanently disabled, thwarting attacks predicated on system threads eventually calling untrusted functions. We present a formal design for soft termination and an implementation of it for Java, built using Java bytecode rewriting, and demonstrating reasonable performance (3-25% slowdowns on benchmarks).

1. Introduction

In recent years, many systems have turned to language runtime systems for enforcement of security. Language-based enforcement of security was popularized by Java and the Java Virtual Machine (JVM), which were adopted by Netscape for its Navigator 2.0 browser in 1995. Java promised an environment where untrusted and malicious code (hereafter called a “codelet”¹) could run safely inside the Web browser, enhancing the user’s Web experience without jeopardizing the user’s security. Rather than using kernel-based protection, the JVM would run inside the same address space as the browser, providing protection

and separation as a side-effect of enforcing its type system. A combination of static and dynamic type checking would serve to prevent a malicious codelet from forging a reference to an arbitrary memory location and subverting the system. In addition to its applications within Web browsers, codelets have also been touted for OS kernel extensions, active networking, extensible databases, agent-based negotiating systems, and presumably other problem domains.

The *promise* of Java may be attractive, but a large number of security flaws have been discovered in Java since its release [10, 24]. Significant strides have been made at understanding the type system [1, 27, 11, 12, 8, 9] and protecting the Java system classes from being manipulated into violating security [30, 31, 18, 13, 14], but efforts to control resource exhaustion have lagged behind. A simple infinite loop will still freeze the latest Web browsers. The most successful systems to date either run the JVMs in separate processes or machines [23, 26], surrendering any performance benefits from running the JVM together with its host application, or create a process-like abstraction inside the JVM [4, 28, 20, 5]. These process abstractions either complicate memory sharing or make it completely impossible.

This paper describes a new language runtime-based mechanism called *soft termination*. While it is not specific to Java, soft termination can be deployed on Java, and we present a Java-based implementation. Soft termination is intended to be invoked either by an administrator or by a system resource monitor which has concluded that a codelet is exceeding its allotted resources and may no longer be allowed to run. Soft termination provides semantics similar to the UNIX `ps` and `kill` commands while preserving system integrity after termination, yet requires neither process-like structures nor limits on memory sharing. Our implementation of soft termination is defined as a code-to-code transformation, and is thus more easily portable across languages and implementations of the same language. Soft termination supports two kinds of program termination: safe thread termination and safe

¹The term “codelet” is also used in artificial intelligence, numerical processing, XML tag processing, and PDA software, all with slightly different meanings. When we say “codelet,” we refer to a small program meant to be executed in conjunction with or as an internal component of a larger program.

code disabling. Safe thread termination must deal with cases where the target thread is currently executing critical system code that may not necessarily be designed to respond to an asynchronous signal. Safe codelet disabling must deal with cases where future threads may invoke functions of the disabled codelet, yet the codelet should not be able to “hijack” the thread and continue execution.

In the following sections, we discuss the concept of soft termination, and present our implementation. Section 2 discusses how prior work has addressed these concerns. Section 3 formalizes and describes what we mean by termination. Section 4 describes our Java-based implementation of soft termination, and mentions a number of Java-specific issues that we encountered. We present performance measurements in section 5. Finally, section 6 describes some experience with using soft termination in real-world situations.

2. Related work

Systems such as Smalltalk, Pilot, Cedar, Lisp Machines, and Oberon have taken advantage of language-based mechanisms to provide OS-like services. Perhaps as early as the Burroughs B5000 series computers, language based mechanisms were used for security purposes. More recently, language-based enforcement of security was popularized by Java, originally deployed by Netscape for its Navigator 2.0 browser in 1995 to run untrusted applets.

However, these systems provide little or no support for resource management on the programs they run. A number of projects have been developed to address this. A recent Scheme system called MrEd [17] supports thread termination and management of resources like open files but has no way of disabling code from running in future threads. Some systems, such as PLAN [21], restrict the language to guarantee that programs will terminate. In general, many language systems support interactive debugging, which includes the ability to interrupt a running program and inspect its state. This can be performed with operating-system services or by generating inline code to respond to an external debugger.

Much of the recent research in this area has been focused on the Java programming language. PERC [25], for instance, is a Java-derived language that extends Java to support asynchronous exceptions. A programmer may specify blocks with provable limits on their runtime and asynchronous exceptions are deferred until the block completes execution.

J-Kernel [20] is a system for managing multiple Java codelets running in the same JVM. It is written entirely in Java, giving it the advantage of working with multiple JVMs with minimal adjustment. It is implemented as a transformation on Java bytecode as the bytecodes are

loaded by the system. J-Kernel isolates threads to run within specific codelets; cross-domain calls are supported via message passing from one codelet thread to another or to the system. By isolating threads to their codelets, it becomes safe to arbitrarily deschedule a thread.

JRes [7] is a resource management system for Java. Bytecode rewriting is used to instrument memory allocation and object finalization in order to maintain a detailed accounting of memory usage. Again, termination is mentioned, but no significant details are provided.

KaffeOS [2, 3] provides an explicit process-like abstraction for Java codelets. KaffeOS is implemented as a heavily customized JVM with significant changes to the underlying language runtime system and system libraries. Code termination is supported in the same manner as a traditional operating system: user codelets are strongly separated from the kernel by running in separate heaps. Memory references across heaps are heavily restricted. Bernadat *et. al.* [5] and van Doorn [29] describe similar systems that customize a JVM in order to support better memory account and security. van Doorn takes advantage of lightweight mechanisms provided by an underlying micro-kernel. KaffeOS provides a style of termination we call hard termination (see section 3.2). It’s not clear how termination is supported in the other systems.

3. System design

A large space of possible designs exist for supporting termination in language runtimes. We first consider the naïve solutions and explain the hard problems raised by their failings. We then discuss how operating systems perform termination and finally, describe our own system.

3.1. Naïve termination

One naïve solution to termination would be to identify undesired threads and simply remove them from the thread scheduler. This technique is used by Java’s deprecated `Thread.destroy()` operation². Unfortunately, there are numerous reasons this cannot work in practice.

Critical sections A thread may be in a critical section of system code, holding a lock, and updating a system data structure. Descheduling the thread would either leave the system in a deadlock situation (if the lock is not released) or leave the system data structures in an undefined state, potentially destabilizing the entire system (if the lock is forcibly released).

Boundary-crossing threads In an object-oriented system, a program wishing to inspect or manipulate an

²see <http://java.sun.com/products/jdk/1.2/docs/guide/misc/threadPrimitiveDeprecation.html>

object invokes methods on that object. When memory sharing is unrestricted between the system and its codelets or among the codelets, these method invocations could allow a malicious codelet to “hijack” the thread from its caller and perhaps never release it. This is especially a concern if the thread in question is performing system functions, such as finalizing dead objects prior to garbage collection.

Blocking calls The language runtime system has support for making native OS system calls. A thread should not be descheduled while it is blocking on a system call such as an I/O call.

Another naïve solution is to force an asynchronous exception, as done by Java’s deprecated `Thread.stop()` operation. While this exception will wait for blocking calls to complete, it may still occur inside a critical section of system code. In addition, blocking calls could potentially never return, resulting in a non-terminable thread. Finally, a workaround is needed to prevent user-level code from catching the exception.

3.2. Hard termination

Operating systems like Unix support termination by carefully separating the kernel from the user program. When a process is executing in user space, the kernel is free to immediately deschedule all user threads and reclaim the resources in use.

If the process is executing in the kernel, termination is normally delayed; a flag is checked when the kernel is about to return control to the user process. In cases where the kernel may perform an operation that could potentially block forever (e.g., reading from the network), the kernel may implement additional logic to interrupt the system call. Other system calls, those that complete in a guaranteed finite time, need not check whether their user process has been terminated, as the kernel will handle the termination signal on the way out. We call such a mechanism a *hard termination system* because once termination is signaled, user-level code may be terminated immediately with no harmful side-effects.

3.3. Soft termination

Unlike a traditional operating system, the boundary between user and system code in a language runtime is harder to define. While all code within the system is generally tagged with its protection domain, there is nothing analogous to a system call boundary where termination signals can be enforced. Furthermore, because a thread could easily cross from user to system code and back many times, there may never be a correct time at which it becomes safe to terminate a thread.

This section introduces a design we call *soft termination* and describes the properties we would find desirable. We present an implementation of soft termination based on code rewriting for a simplified language and prove that all programs will terminate when signaled to do so.

3.3.1. Key ideas: Soft termination is based on the idea that a codelet may be instrumented to check for a termination condition during the normal course of its operation. Our goal is to perform these checks as infrequently as necessary — only enough to ensure that a codelet may not execute an infinite loop. Furthermore, as with the Unix kernel, we would like the termination of a codelet to not disturb any system code it may be using at the time it is terminated, thus preserving system correctness.

The soft termination checks are analogous to *safe points*, which are used in language environments to insert checks for implementing stack overflow detection, preemptive multitasking, inter-process and inter-task communication, barrier synchronization, garbage collection systems, and debugging functions. A more complete discussion is provided in Feeley [15]. In Feeley’s terminology, our implementation uses “minimal polling.”

3.3.2. Formal design: For our analysis, we begin with a simple programming language having methods or functions, conditional expressions, and simple exceptions (see figures 1 and 2). In our language, a program is a collection of function definitions followed by an expression to be evaluated. A function body contains function/method calls as well as primitive operations, conditionals, and exceptions. A **(CheckTermination)** expression is assumed to return a boolean value, indicating whether termination for the current codelet has been externally (and asynchronously) requested. We write the semantics of our language using the same style as Felleisen and Hieb [16]. We include a delta (δ) function which maps primitive names and valid arguments to values, and we assume the syntactic property that no name collisions occur; that is, no function is defined more than once, and no primitive function is redefined. The unusual definition of evaluation contexts is designed to permit single-step reduction of **(throw)** instructions. We state without proof that our language is consistent and well-formed.

The soft termination transformation, \mathcal{X} , is described in figure 3. Rule 4 of this transformation inserts the check for termination before every function invocation. The other rules describe how the transformation continues recursively on expressions.

Proving that a transformed program terminates in a finite time, given that **(CheckTermination)** returns true, is a straightforward exercise. We start with all possible expressions M from figure 1. After applying \mathcal{X} to M , and assuming **CheckTermination** is true, we wish to show

$P = \Gamma M$
 $\Gamma = D \dots D$
 $D = (\mathbf{define} (f x \dots) M)$
 $M = (f M \dots) \mid (\mathbf{if} M M M) \mid (\mathbf{let} (x M) \dots M) \mid (\mathbf{try} M M) \mid (\mathbf{throw}) \mid V \mid (\mathbf{CheckTermination})$
 $V = x \mid \mathbf{true} \mid \mathbf{false} \mid c$
 $f, x = \text{identifiers}$
 $E_a = (f V \dots E_a M \dots) \mid (\mathbf{if} E_a M M) \mid (\mathbf{try} E_a M) \mid (\mathbf{try} E_b M) \mid (\mathbf{let} (x V) \dots (x E_a) (x M) \dots M)$
 $E_b = (f V \dots E_c M \dots) \mid (\mathbf{if} E_c M M) \mid (\mathbf{let} (x V) \dots (x E_c) (x M) \dots M)$
 $E_c = [] \mid (f V \dots E_c M \dots) \mid (\mathbf{if} E_c M M) \mid (\mathbf{let} (x V) \dots (x E_c) (x M) \dots M)$
 $E = E_a \mid E_b \mid []$
 final states = $V \mid \mathbf{error}$

Figure 1. Simple language used for our analysis.

$$\begin{aligned}
 E[(f V \dots)] &\mapsto \begin{cases} E[V_0] & \text{if } \delta(f, V, \dots) = V_0 \\ E[[V \dots / x \dots] M] & \text{if } (\mathbf{define} (f x \dots) M) \in \Gamma \\ \mathbf{error} & \text{otherwise} \end{cases} \\
 E[(\mathbf{CheckTermination})] &\mapsto \text{true or false (depending on external conditions)} \\
 E[(\mathbf{if} \text{ true } M_1 M_2)] &\mapsto E[M_1] \\
 E[(\mathbf{if} \text{ false } M_1 M_2)] &\mapsto E[M_2] \\
 E[(\mathbf{if} V M_1 M_2)] &\mapsto \mathbf{error} \text{ if } V \neq \text{true, false} \\
 E[(\mathbf{try} V M)] &\mapsto E[V] \\
 E[(\mathbf{try} (\mathbf{throw}) M)] &\mapsto E[M] \\
 E[(\mathbf{throw})] &\mapsto \begin{cases} \mathbf{error} \text{ if } E = [] \\ E_a[(\mathbf{throw})] \text{ where } E = E_a[E_b] & \text{otherwise} \end{cases} \\
 E[(\mathbf{let} (x V) \dots M)] &\mapsto E[[V \dots / x \dots] M]
 \end{aligned}$$

Figure 2. An operational semantics for our language.

- (1) $\mathcal{X}[[\Gamma M]] = \mathcal{X}[[\Gamma]] \mathcal{X}[[M]]$
- (2) $\mathcal{X}[[D \dots D]] = \mathcal{X}[[D]] \dots \mathcal{X}[[D]]$
- (3) $\mathcal{X}[[\mathbf{define} (f x \dots) M]] = (\mathbf{define} (f x \dots) \mathcal{X}[[M]])$
- (4) $\mathcal{X}[[f M \dots]] = (\mathbf{let} (t \mathcal{X}[[M]]) \dots (\mathbf{if} (\mathbf{CheckTermination}) (\mathbf{throw}) (f t \dots)))$
- (5) $\mathcal{X}[[\mathbf{if} M M M]] = (\mathbf{if} \mathcal{X}[[M]] \mathcal{X}[[M]] \mathcal{X}[[M]])$
- (6) $\mathcal{X}[[\mathbf{try} M M]] = (\mathbf{try} \mathcal{X}[[M]] \mathcal{X}[[M]])$
- (7) $\mathcal{X}[[\mathbf{throw}]] = (\mathbf{throw})$
- (8) $\mathcal{X}[[\mathbf{let} (x M) \dots M]] = (\mathbf{let} (x \mathcal{X}[[M]]) \dots \mathcal{X}[[M]])$
- (9) $\mathcal{X}[[V]] = V$

Figure 3. The soft termination transformation.

$P = \Gamma M$
 $\Gamma = D \dots D$
 $L = \mathbf{system} \mid \mathbf{codelet}$
 $D = (\mathbf{define} L (f x \dots) M) \mid$
 $\quad (\mathbf{define} \mathbf{blocking} (f_{\mathbf{blocking}} x \dots) f_{\mathbf{nonblocking}} M)$
 $M = (f M \dots) \mid (\mathbf{if} M M M) \mid (\mathbf{let} (x M) \dots M) \mid (\mathbf{try} M M) \mid (\mathbf{throw}) \mid V \mid (\mathbf{CheckTermination})$
 $V = x \mid \mathbf{true} \mid \mathbf{false} \mid c$
 $f, x = \text{identifiers}$
 $E_a = (f V \dots E_a M \dots) \mid (\mathbf{if} E_a M M) \mid (\mathbf{try} E_a M) \mid (\mathbf{try} E_b M) \mid (\mathbf{let} (x V) \dots (x E_a) (x M) \dots M)$
 $E_b = (f V \dots E_c M \dots) \mid (\mathbf{if} E_c M M) \mid (\mathbf{let} (x V) \dots (x E_c) (x M) \dots M)$
 $E_c = [] \mid (f V \dots E_c M \dots) \mid (\mathbf{if} E_c M M) \mid (\mathbf{let} (x V) \dots (x E_c) (x M) \dots M)$
 $E = E_a \mid E_b \mid []$
 $\text{final states} = V \mid \mathbf{error}$

Figure 4. An extended language for analysis distinguishing codelets from system code.

(1) $\mathcal{X}[\Gamma M] = \mathcal{X}[\Gamma] \mathcal{X}_{\mathbf{codelet}}[[M]]$
(2) $\mathcal{X}[D \dots D] = \mathcal{X}[[D]] \dots \mathcal{X}[[D]]$
(3a) $\mathcal{X}[(\mathbf{define} \mathbf{codelet} (f x \dots) M)] = (\mathbf{define} (f x \dots) \mathcal{X}_{\mathbf{codelet}}[[M]])$
(3b) $\mathcal{X}[(\mathbf{define} \mathbf{system} (f x \dots) M)] = (\mathbf{define} (f x \dots) \mathcal{X}_{\mathbf{system}}[[M]])$
(3c) $\mathcal{X}[(\mathbf{define} \mathbf{blocking} (f_{\mathbf{blocking}} x \dots) f_{\mathbf{nonblocking}} M)] = (\mathbf{define} (f_{\mathbf{wrapper}} x \dots) (\mathbf{let} (b, t (f_{\mathbf{nonblocking}} x \dots)) (\mathbf{if} b t (\mathbf{if} (\mathbf{CheckTermination}) (\mathbf{throw}) (f_{\mathbf{blocking}} x \dots))))))$
(4a) $\mathcal{X}_{\mathbf{codelet}}[[f M \dots]] = (\mathbf{let} (t \mathcal{X}_{\mathbf{codelet}}[[M]]) \dots (\mathbf{if} (\mathbf{CheckTermination}) (\mathbf{throw}) (\mathcal{X}[[f]] t \dots)))$
(4b) $\mathcal{X}_{\mathbf{system}}[[f M \dots]] = (\mathcal{X}[[f]] \mathcal{X}_{\mathbf{system}}[[M]] \dots)$
(5a) $\mathcal{X}[[f_{\mathbf{blocking}}]] = f_{\mathbf{wrapper}} \text{ if } (\mathbf{define} \mathbf{blocking} (f_{\mathbf{blocking}} \dots) f_{\mathbf{nonblocking}} \dots) \in \Gamma$
(5b) $\mathcal{X}[[f]] = f \text{ otherwise}$
(6a) $\mathcal{X}_{\mathbf{codelet}}[[\mathbf{if} M M M]] = (\mathbf{if} \mathcal{X}_{\mathbf{codelet}}[[M]] \mathcal{X}_{\mathbf{codelet}}[[M]] \mathcal{X}_{\mathbf{codelet}}[[M]])$
(6b) $\mathcal{X}_{\mathbf{system}}[[\mathbf{if} M M M]] = (\mathbf{if} \mathcal{X}_{\mathbf{system}}[[M]] \mathcal{X}_{\mathbf{system}}[[M]] \mathcal{X}_{\mathbf{system}}[[M]])$
...

Figure 5. The soft termination transformation with codelets and blocking calls.

Note the domain of this transformation is in language of figure 1 with a small extension to allow multiple return values. The semantics of figure 2 apply.

the program enters a “locked” state, where termination is guaranteed. We call this set of expressions M_{lock} .

$$\begin{aligned}
M_{lock} = & \text{(if (CheckTermination) (a)} \\
& \text{(throw) (f V ...)) |} \\
& \text{(if true (throw) (f V ...)) | (b)} \\
& \text{(if } M_{lock} M_{lock} M_{lock}) | (c)} \\
& \text{(let (x } M_{lock}) \dots M_{lock}) | (d)} \\
& \text{(try } M_{lock} M_{lock}) | (e)} \\
& \text{(throw) | (f)} \\
& V (g)
\end{aligned}$$

To prove termination, we must show that, once **(CheckTermination)** returns *true*, M_{lock} is closed under program stepping, and that the syntactic length of the program will be strictly decreasing.

Closure may be stated as follows: if $M \mapsto M'$, and $M \in M_{lock}$, then $M' \in M_{lock}$. By inspection, for all possible expressions in M_{lock} , we observe that our semantics preserves closure.

The syntactic length property may be stated as follows: If **(CheckTermination)** always returns *true*, then for $M, M' \in M_{lock}$ such that $M \mapsto M'$, $|M'| < |M|$. That is, at each step, a program in M_{lock} must get smaller. To prove this, it suffices to observe that all reductions other than function application make the program smaller, and that for an expression in M_{lock} , all function applications are “fenced out” by a call to **(CheckTermination)**.

When a program is being evaluated and **(CheckTermination)** is not guaranteed to return *true*, a program may not be in M_{lock} . The case where this matters is when termination is requested just before an application; in this case, the next step taken will be the application itself, which will increase the length of the program. However, this resulting state must necessarily be in M_{lock} , and at this point the program length will be strictly decreasing. In other words, an annotated program which is not in M_{lock} is guaranteed to enter M_{lock} in one step.

Since program lengths are finite, a program annotated by \mathcal{X} is guaranteed to finish in a finite number of steps, once the **CheckTermination** condition is raised.

Codelet, system, and blocking code Termination becomes trickier when blocking calls are introduced and when we distinguish between **codelet** and **system** code. Figure 4 introduces an extension to our little language where functions are labeled as to the origin of their code, and whether or not they block. The **blocking** label would normally be applied to system functions known to block, such as I/O operations. By convention, every blocking function $f_{blocking}$ is defined to have some roughly equivalent $f_{nonblocking}$ which returns a tuple containing a boolean and a value. The boolean indicates whether the

call succeeded and if so, the value is the result. Polling $f_{nonblocking}$ is intended to be semantically equivalent to calling $f_{blocking}$.

The transformation in figure 5 is more complex than that of figure 3 so it bears some explanation.

Rules 3a through 3c describe the transformation on function definitions. Rules 3a and 3b say that system code and codelets have separate transformations, \mathcal{X}_{system} and $\mathcal{X}_{codelet}$, respectively. Rule 3c says a non-blocking wrapper is created for every blocking function. In this case, the wrapper invokes the non-blocking call equivalent to the original. If the call succeeds, the value is returned. Otherwise, the wrapper polls to see if termination has been indicated and throws an exception if it has. Otherwise, the non-blocking call is recursively polled again. When soft termination is implemented in a concrete language, other mechanisms may be available to create these wrappers.

Rules 4a and 4b describe how function calls are handled for \mathcal{X}_{system} and $\mathcal{X}_{codelet}$. For codelets, the termination checks are added, just as in figure 3. For system code, no termination checks are added. The purpose of this is to avoid destabilizing system code. When termination is desired, it will only change the control flow of the codelet, not the system. If system code makes an up-call to a codelet, it becomes possible for a terminated codelet to throw an exception. System code would be responsible for catching this exception and proceeding appropriately. In a mobile code environment, where codelets are untrusted, system code must already be prepared for such up-calls to throw exceptions, but system code now need not be concerned with asynchronous termination.

Rules 5a and 5b describe how blocking function calls are replaced with calls to their non-blocking wrappers. The remainder of the rules describe how the transformation continues recursively on expressions.

Despite the domain of the transformation in figure 5 being within the language of figures 1 and 2, we can no longer prove termination; system code may have infinite loops. It is now the programmer’s responsibility to label all code that may potentially have such infinite loops and provide some kind of non-blocking wrappers. It’s important to point out that a program consisting strictly of functions labeled $\mathcal{X}_{codelet}$ will be transformed to precisely the same result as if it were written in the simpler language, and thus can be terminated. When control flow is executing in a system function, the termination checks will be deferred. When control flow is executing in a codelet function, termination will happen normally. This also addresses the the boundary-crossing thread concern (see section 3.1): if system code calls a function within a terminated codelet, that function will be guaranteed to terminate in a finite time, and thus a system thread cannot be “hijacked.”

4. Java implementation

In an effort to understand the practical issues involved with soft termination, we implemented it for Java as a transformation on Java bytecodes. Our implementation relies on a number of Java-specific features. We also address a number of Java-specific quirks that would hopefully not pose a problem in other language systems.

4.1. Termination check insertion

Java compilers normally output Java bytecode. Every Java source file is translated to one or more class files, later loaded dynamically by the JVM as the classes are referenced by a running program. JVMs know how to load class files directly from the disk or indirectly through “class loaders,” invoked as part of Java’s dynamic linking mechanism. A class loader, among other things, embodies Java’s notion of a name space. Every class is tagged with the class loader that installed it, such that a class with unresolved references is linked against other classes from the same source. A class loader provides an ideal location to rewrite Java bytecode, implementing the soft termination transformation. A codelet appears in Java as a set of classes loaded by the same class loader. System code is naturally loaded by a different class loader than codelets, allowing us to naturally implement the codelet transform ($X_{codelet}$) separately from the system code transform (X_{system}).

Our implementation uses the CFParse³ and JOIE [6] packages, which provide interfaces for parsing and manipulating Java class files.

The basic structure of our bytecode modification is exactly as described in section 3.3.2. A static boolean field is added to every Java class, initially set to false. The Check-Termination operation, implemented in-line, tests if this field is true, and if so, calls a handler method that decides whether to throw an exception. As an extension to the semantics of figure 5, we allow threads and thread groups to be terminated as well as specific codelets, regardless of the running thread. The termination handler, when invoked, compares its caller and its current thread against a list of known termination targets. Note that, if the boolean field is set to false, the runtime overhead is only the cost of loading and checking the value, and then branching forward to the remainder of the method body.

Figure 6 shows how the codelet soft termination transform ($X_{codelet}$) would be applied to a Java method invocation.

4.2. Control flow

Java has a much richer control flow than the little language introduced earlier. First and foremost, Java byte-

³<http://www.alphaworks.ibm.com/tech/cfparse>

```
getstatic termination_signal Z
ifeq SKIP
invokestatic termination_handler()V
SKIP: original_invoke_instruction
```

Figure 6. Bytecodes inserted into Java for the soft termination check.

code has a general-purpose branch instruction. We do nothing special for forward branches, but we treat backward branches as if they were method invocations and perform the appropriate code transformation. An additional special case we must handle is a branch instruction which targets itself.

Java bytecode also supports many constructions that have no equivalent Java source code representation. In particular, it is possible to arrange for the `catch` portion of an exception handler to be equal to the `try` portion. That means an exception handler can be defined to handle *its own exceptions*. Such a construction allows for infinite loops without any method invocation or backward branching. While such code should most likely be rejected by the Java bytecode verifier, as it is not allowed in the JVM specification [22], the bytecode verifier currently treats such constructions as valid. We specifically check for and reject programs with overlapping `try` and `catch` blocks.

Lastly, Java bytecode supports a notion of subroutines within a Java method using the `jsr` and `ret` instructions. `jsr` pushes a return address on the stack which `ret` consumes before returning. The Java bytecode verifier imposes a number of restrictions on how these instructions may be used. In particular, a return address is an opaque type which may be consumed at most once. The verifier’s intent is to insure that these instructions may be used only to create subroutines, not general-purpose branching. As such, we instrument `jsr` instructions the same way we would instrument a method invocation and we do nothing for `ret` instructions.

4.3. Blocking calls

To address blocking calls, we follow the transformation outlined in section 3.3.2. Luckily, all blocking method calls in the Java system libraries are `native` methods (implemented in C) and can be easily enumerated and studied by examining the source code of the Java class libraries.

Java provides a mechanism for interrupting blocking calls: `Thread.interrupt()`. This method causes the blocking method to throw a `java.lang.InterruptedException` or `java.io.InterruptedIOException` exception.

While the X transformation is defined to poll a non-blocking version, we would prefer to take advantage of the interruption support already inside the JVM.

To accomplish this, we must track which threads are currently blocking and the codelets on behalf of whom they are blocking. The wrapper classes now get the current thread and save it in a global table for later reference. In order to learn the codelet on whose behalf we are about to block, we take advantage of the stack inspection primitives built into modern Java systems [31, 18].

Stack inspection provides two primitives that we use: `java.security.AccessController.doPrivileged()` and `getContext().getContext()` returns an array of `ProtectionDomains` that map one-to-one with codelets. The `ProtectionDomain` identities are then saved alongside the current thread before the blocking call is performed.

When we wish to terminate a codelet, we look up whether it is currently waiting on a blocked call, and only then do we interrupt the thread.

Taking advantage of another property of Java stack inspection, we can distinguish between blocking calls being performed on behalf of system code and those being performed indirectly by a codelet. Generally, we would rather not interrupt a blocking call if system code is depending on its result and system state could become corrupted if the call were interrupted. On the other hand, we have no problem interrupting a blocking call if only a codelet is depending on its result. Java system code already uses `doPrivileged()` to mark regions of privileged system code and `getContext()` to get dynamic traces for making access control checks. These regions are exactly the same regions where preserving system integrity upon termination is important; if system code is using its own security privileges, it wants the operation to succeed regardless of its caller's privileges. Thus, we overload the semantics of these existing security primitives to include whether blocking calls should be interrupted.

4.4. Invoking termination

Our system supports three kinds of termination: termination of individual threads, termination of thread groups, and termination of codelets.

To terminate a thread or thread group, we must map the threads we wish to terminate to the set of codelets potentially running those threads and set the termination signal on all classes belonging to the target codelet. Furthermore, we must check if any of these threads are currently blocking and interrupt them, as appropriate (see section 4.3). At this point, the thread requesting termination performs a `Thread.join()` on the target thread(s), waiting until

they complete execution. Once the target thread has completed, the termination signals are cleared and execution returns to its normal performance.

If multiple threads are executing concurrently over the same set of classes and only one is terminated, the termination handler will be invoked for threads not targeted, only to return shortly thereafter. These threads will experience degraded performance while the target thread is still running.

In the case where we wish to terminate a specific codelet, disabling all its classes forever, we simply set the termination signal on all classes in the codelet and immediately return. Any code that invokes a method on a disabled class will receive an exception indicating the class has been terminated.

Once a codelet has been signalled to terminated, if a codelet's thread is executing in a system class at the time, execution continues until the thread returns to a user class. If the codelet is currently making a blocking call, the call is interrupted and the thread resumes execution. Once the thread has resumed executing in the user's class, it becomes subject to the soft termination system.

For all codelet threads which are executing within the codelet, if they try to call a method within the codelet, the method fails with an exception. If they try to make a backward branch, the soft termination code added will throw an exception. In all cases, each thread of control unwinds, preventing the codelet from performing any meaningful work. Finally, if any other codelet or the system makes a call into this codelet, it will fail immediately, preventing the codelet from performing any meaningful work. As shown in section 3.3.2, the codelet is guaranteed to terminate.

Note that termination requests can be handled concurrently. A potential for deadlock occurs where a thread could request its own termination, or where a cycle of threads may request each others' termination. In production, where a user is manually terminating threads or codelets, this would not be an issue. The termination operation itself should not be provided to untrusted codelets. Instead, it is protected using the same security mechanisms as other Java privileged calls.

4.5. Optimizations

If a Java method contains a large number of method invocations, the transformed method may be significantly larger than the original, causing performance problems. We address this concern by observing that we get similar termination semantics by adding the soft termination check at the entry points to methods rather than at the call sites. So long as a soft termination check occurs either before or immediately after a method invocation, the resulting program will execute the same.

Additionally, we implemented an optimization to statically determine if a method has no outgoing method calls (a leaf method). For leaf methods, a termination check at the beginning of the method is unnecessary. If the method has loops, they will have their own termination checks. If not, the method is guaranteed to complete in a finite time. Regardless, removing the initial termination check from leaf methods preserves the semantics of soft termination and should offer a significant performance improvement, particularly for short methods such as “getter” and “setter” methods.

A more aggressive optimization, which we have not yet performed, would be an inter-procedural analysis of statically terminating methods. A method which only calls other terminating methods and has no backward branches will always terminate. Likewise, we have not attempted to distinguish loops that can be statically determined to terminate in a finite time (*i.e.*, loops that can be completely unrolled). Such analyses could offer significant performance benefits to a production implementation of soft termination.

4.6. Synchronization

A particularly tricky aspect of supporting soft termination in a Java system is supporting Java’s synchronization primitives.

The Java language and virtual machine specifications are not clear on how the system behaves when a deadlock is encountered [19, 22]. With Sun’s JDK 1.2, the only way to recover from a deadlock is to terminate the JVM. Obviously, this is an unsatisfactory solution. Ideally, we would like to see a modification to the JVM where locking primitives such as the `monitorenter` bytecode are interruptible as are other blocking calls in Java. We could then apply standard deadlock detection techniques and chose the threads to interrupt.

Additionally, it is possible to construct Java classes where the `monitorenter` and `monitorexit` operations are not properly balanced. Despite the fact that there exist no equivalent Java source programs, current JVM bytecode verifiers accept such programs. As such, it is possible for a malicious program to acquire a series of locks and terminate without those locks being released until the JVM terminates.

Our current system makes no attempt to address these issues.

4.7. Thread scheduling

Our work fundamentally assumes the Java thread system is preemptive. This was not the case in many early Java implementations. Without a preemptive scheduler, a malicious codelet could enter an infinite loop and no other thread would have the opportunity to run and request the

termination of the malicious thread. This would defeat soft termination.

4.8. System code interruptibility

Our work fundamentally assumes that all system methods that may be invoked by a codelet will either return in a finite time or will reach a blocking native method call which can be interrupted. It may be possible to construct an input to system code that will cause the system code itself to have an infinite loop. Addressing this concern would require a lengthy audit of the system code to guarantee there exist no possible inputs to system functions that may infinitely loop.

4.9. Memory consistency models

The Java language defines a relaxed consistency model where updates need not be propagated without the use of locking primitives. In our current prototype, we use no synchronization primitives when accessing this variable. Since external updates to the termination signal could potentially be ignored by the running method, this could defeat the soft termination system.

Instead, we take advantage of Java’s `volatile` modifier. This modifier is provided to guarantee that changes must be propagated immediately [19]. On the benchmark platform we used, the performance impact of using `volatile` versus not using it is negligible. However, on other platforms, especially multiprocessing systems, this may not be the case.

4.10. Defensive security

Our prototype implementation makes no attempt to protect itself from bytecode designed specifically to attack the termination system (*e.g.*, setting the termination flag to *false*, either directly or through Java’s reflection interface). Such protection could be added as a verification step before the bytecode rewriting.

5. Performance

We measured the performance of our soft termination system using Sun Microsystems Ultra 10 workstations (440 MHz UltraSparc II CPU’s, 128 MB memory, Solaris 2.6), and Sun’s Java 2, version 1.2.1 build 4, which includes a JIT. Our benchmarks were compiled with the corresponding version of `javac` with optimization turned on.

We used two classes of benchmark programs: microbenchmarks that test the impact of soft termination on various Java language constructs (and also measuring worst case performance), and macrobenchmarks which represent a number of real-world applications. We mea-

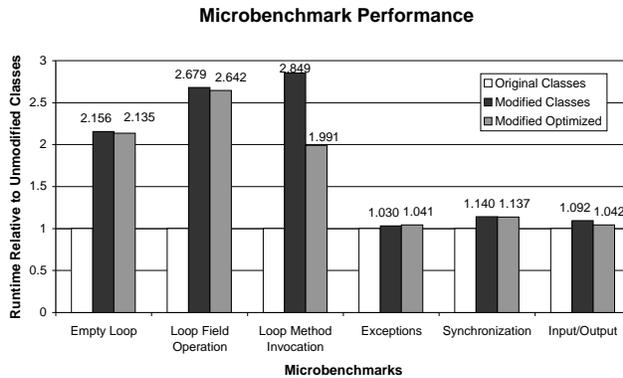


Figure 7. Performance of rewritten microbenchmark class files relative to the performance of the corresponding original class files.

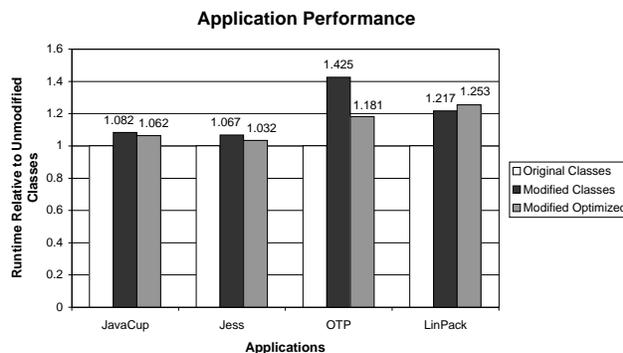


Figure 8. Performance of rewritten application class files relative to the performance of the corresponding original class files.

sured the performance of these systems in three configurations: their original unmodified state, their state after being rewritten, and their state after being rewritten with the leaf method optimization discussed in section 4.5. Generally, when we discuss results in this section, we refer to the optimized numbers.

5.1. Microbenchmarks

We first measured a series of microbenchmarks to stress-test the JVM with certain language constructs: looping, method and field accesses, exception handling, synchronization, and I/O. We used a microbenchmark package developed at University of California, San Diego, and modified at University of Arizona for the Sumatra Project⁴. The results are shown in figure 7.

⁴The original web site is <http://www-cse.ucsd.edu/users/wgg/JavaProf/javaprof.html>. The source we used was distributed from <http://www.cs.arizona.edu/sumatra/>

As one would expect, tight loops suffered the worst slow-downs: roughly a factor of two. When we rewrite the loop, the termination check costs roughly the same as the original loop termination check, so it's sensible to see a factor of two performance degradation.

For other microbenchmarks, we saw much smaller overheads. The overhead of handling exceptions, performing synchronization, or doing I/O operations dominates the cost of checking for termination. The largest overhead of these was 14% for the synchronization microbenchmark. The additional overhead can be attributed to performing the termination check once for each iteration of the loop.

For the I/O and exception-handling microbenchmarks, the performance figures are much better. Since I/O and exception handling are relatively costly operations, modifications don't have as significant an impact on performance.

We observe that the leaf method optimization generally has a small performance benefit. The loop method invocation microbenchmark shows the most dramatic improvement; the optimized benchmark runs 30% faster than the unoptimized benchmark. In one case, the exception handling benchmark, the optimized program ran roughly 1% slower than the unoptimized program. Similar behavior occurred in the Linpack macrobenchmark. The optimized programs are genuinely performing fewer termination checks, but still have a longer running time. The culprit appears to be Sun's JIT compiler (sunwjit). When the benchmarks are run with the JIT disabled, the optimized programs are strictly faster than the non-optimized programs. We have observed similar deviant behavior with Sun's HotSpot JIT running on Sparc/Solaris and Linux/x86. We have sent an appropriate bug report to Sun.

5.2. Application benchmarks

We benchmarked the real-world applications JavaCup⁵, Linpack⁶, Jess⁷, and JOTP⁸. These programs were chosen to provide sufficiently broad insight into our system's performance.

JavaCup, a parser-generator, was chosen to demonstrate how the rewritten classes perform in handling text processing. Jess, an expert system, was chosen to demonstrate the performance of the rewritten classes in handling symbolic data and solving logic problems. Linpack is a loop-intensive floating-point benchmark. JOTP is a one-time password generator which uses a cryptographic hash function. The results are shown in figure 8.

<ftp://benchmarks/Benchmark.java>

⁵<http://www.cs.princeton.edu/~appel/modern/java/CUP/>

⁶<http://netlib2.cs.utk.edu/benchmark/linpackjava/>

⁷<http://herzberg1.ca.sandia.gov/jess/>

⁸<http://www.cs.umd.edu/~harry/jotp/>

Microbenchmarks			Application Benchmarks		
Microbenchmark	Checks per Second		Application Benchmark	Checks per Second	
	Unoptimized	Optimized		Unoptimized	Optimized
Empty Loop	4.723×10^7	4.843×10^7	JavaCup	2.771×10^7	2.152×10^7
Loop Field Operation	3.887×10^7	3.914×10^7	Jess	1.013×10^7	1.501×10^7
Loop Method Invocation	3.995×10^7	3.515×10^7	JOTP	4.102×10^7	2.905×10^7
Exceptions	3.729×10^7	7.366×10^7	Linpack	3.128×10^7	2.598×10^7
Synchronization	5.427×10^7	4.082×10^7			
Input/Output	5.978×10^5	9.713×10^5			

Figure 9. Average number of termination checks performed per one second of increase in runtime for the micro and application benchmarks.

For the JavaCup test, we generated a parser for the Java 1.1 grammar. There was a modest 6% increase in execution time. For the Jess test, we ran several of the sample problems included with Jess through the system, and calculated the cumulative runtimes. There was a 3% increase in execution time. Both JavaCup and Jess represent applications which do not make extensive use of tight loops. Instead, these applications spend more of their time performing I/O and symbolic computations. Their performance closely tracks the performance of the I/O microbenchmark.

For the Java OTP generator, we generated a one-time password from a randomly-chosen seed and password, using 200,000 iterations. There was a 18% increase in runtime. For the Linpack benchmark, there was a 25% increase in runtime. Linpack is a loop-intensive program, while JOTP makes extensive use of method calls as well as loops. As a result, we would expect their performance to more closely track the loop-based microbenchmarks. Note in particular the benefit JOTP got from the leaf method optimization.

5.3. Termination check overhead

To gauge the actual impact of our class file modifications, we counted the number of times we checked the termination flag for each benchmark. Using this, we calculated the number of termination checks being performed per every second increase in runtime overhead. The results for all benchmarks are listed in figure 9.

For the three looping microbenchmarks, We found that roughly 40,000,000 checks are performed for every second of runtime overhead. This evaluates to around 10 CPU cycles for each check performed. Likewise, for the exception and synchronization microbenchmarks, while the cost of performing these original operations far outweighs the cost of the termination checks, still we see the same number of, and sometimes more, termination checks

per second of overhead.

For the input/output microbenchmark, however, only around 970,000 termination checks are performed per second of overhead. This can be attributed to the other operations being performed by the input/output benchmark. In particular, it needs to contend with the additional overhead of the blocking call management code. It is also important to keep in mind that the cost of performing I/O far outweighs the cost of termination checks.

The application benchmarks reflect the results of the microbenchmarks. All of these results fell within the same range, between 15,000,000 and 29,000,000 checks per second of overhead. Since none of the benchmarks perform any significant amount of I/O, these figures are in line with the microbenchmarks results.

These performance figures seem to indicate that for real-world applications, the showdown will be roughly proportional to how much the application's performance is dependent on tight loops. Applications which have tight loops may experience at worst a factor of two slowdown and more commonly 15 – 25%. Applications without tight loops can expect more modest slowdowns, most likely below 7%. The number of termination checks the system can perform per second seems not to be a limiting factor in system performance.

6. Soft termination in practice

In order to demonstrate our soft termination system, we integrated it into the Jigsaw web server. Jigsaw is a free, open Java-based web server which supports servlets⁹. We integrated our bytecode-rewriting system into the servlet loader for Jigsaw. Thus, every servlet that is loaded into Jigsaw is first rewritten to support soft termination.

We also wrote an administrative screen similar to the `top` utility in UNIX. It provides a list of all active servlets

⁹<http://www.w3.org/Jigsaw/>

on the system. It also gives an option to terminate a servlet. Selecting this option activates the terminate signal in the specified servlet. We observed that servlets would take at most ten seconds to terminate.

A number of attacks against Java focusing on resource exhaustion have been proposed [24]. Several of these focus on flaws in Java's access control framework. For example, the standard recipe for designing such attack includes setting a threads priority to `tt MAX_PRIORITY` to help ensure the program does its job. Java specifies an access control privilege for changing thread priority. The more serious *Business Assassin* applet relied on Java allowing unprivileged threads to stop one-another. Our soft termination system does not try to stop these attacks.

A number of other resource exhaustion attacks do not take advantage of flaws in Java's security system. These include creating threads which loop infinitely, overriding the `Applet.stop()` method, or catching a `ThreadDeath` exception and recreating the thread. As mentioned in section 4.4, soft termination successfully stops such applets.

7. Conclusion

While Java and other general-purpose language-based systems for codelets have good support for memory protection, authorization, and access controls, there is little support for termination. Without termination, a system can be vulnerable to denial-of-service attacks or even to bugs where a faulty codelet has an infinite loop.

We have introduced a concept we call soft termination, along with a formal design and an implementation for Java that allows for asynchronous and safe termination of misbehaving or malicious codelets. Soft termination can be implemented without making any changes to the underlying language or runtime system. Our Java implementation relies solely on class-file bytecode rewriting, making it portable across Java systems and easier to consider applying to non-Java systems. In real-world benchmarks, our system shows slow-downs of 3 – 25%. This could possibly be further reduced if we could leverage a safe point mechanism already implemented within the JVM.

A larger research area remains: building language runtimes that support the general process-management semantics of operating systems. Because language runtimes allow and take advantage of threads and memory references that easily cross protection boundaries, traditional operating system processes may not be the appropriate in this new setting. Opportunities exist to design new mechanisms to add these semantics to programming language runtimes.

8. Acknowledgments

Jiangchun “Frank” Luo and Liwei Peng helped implement an early prototype of this system. Matthias Felleisen and Shriram Krishnamurthy provided many helpful comments. This work is supported by NSF Grant CCR-9985332.

References

- [1] J. Alves-Foss, editor. *Formal Syntax and Semantics of Java*. Number 1523 in Lecture Notes in Computer Science. Springer-Verlag, July 1999.
- [2] G. Back and W. Hsieh. Drawing the Red Line in Java. In *Proceedings of the Seventh IEEE Workshop on Hot Topics in Operating Systems*, Rio Rico, Arizona, Mar. 1999. <http://www.cs.utah.edu/flux/papers/redline-hotos7.ps>.
- [3] G. Back, W. C. Hsieh, and J. Lepreau. Processes in KaffeOS: Isolation, resource management, and sharing in Java. In *Proceedings of the Fourth Symposium on Operating Systems Design and Implementation (OSDI 2000)*, San Diego, California, Oct. 2000. <http://www.cs.utah.edu/flux/papers/kaffeos-osdi00-base.html>.
- [4] G. Back, P. Tullmann, L. Stoller, W. C. Hsieh, and J. Lepreau. Techniques for the design of Java Operating System. In *Proceedings of the 2000 Usenix Annual Technical Conference*, San Diego, California, June 2000. <http://www.cs.utah.edu/flux/papers/javaos-usenix00-base.html>.
- [5] P. Bernadat, D. Lambright, and F. Travostino. Towards a resource-safe Java for service guarantees in uncooperative environments. In *IEEE Workshop on Programming Languages for Real-Time Industrial Applications*, Madrid, Spain, Dec. 1998.
- [6] G. Cohen, J. Chase, and D. Kaminsky. Automatic program transformation with JOIE. In *Proceedings of the 1998 Usenix Annual Technical Symposium*, pages 167–178, New Orleans, Louisiana, June 1998.
- [7] G. Czajkowski and T. von Eicken. JRes: A resource accounting interface for Java. In *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 21–35, Vancouver, British Columbia, Oct. 1998.
- [8] D. Dean. The security of static typing with dynamic linking. In *Fourth ACM Conference on Computer and Communications Security*, Zurich, Switzerland, Apr. 1997.
- [9] D. Dean. *Formal Aspects of Mobile Code Security*. PhD thesis, Princeton University, Princeton, New Jersey, Nov. 1998.
- [10] D. Dean, E. W. Felten, D. S. Wallach, and D. Balfanz. Java security: Web browsers and beyond. In D. E. Denning and P. J. Denning, editors, *Internet Besieged: Countering Cyberspace Scofflaws*, pages 241–269. ACM Press, New York, New York, Oct. 1997.
- [11] S. Drossopoulou and S. Eisenbach. Java is type safe — probably. In *Proceedings of the European Conference on*

- Object-Oriented Programming (ECOOP '97)*, Jyväskylä, Finland, June 1997.
- [12] S. Drossopoulou, D. Wragg, and S. Eisenbach. What is Java binary compatibility? In *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 341–358, Vancouver, British Columbia, Oct. 1998.
- [13] G. Edjlali, A. Acharya, and V. Chaudhary. History-based access control for mobile code. In *Proceedings of the 5th ACM Conference on Computer and Communications Security (CCS '98)*, pages 38–48, San Francisco, California, Nov. 1998. ACM Press.
- [14] U. Erlingsson and F. B. Schneider. SASI enforcement of security policies: A retrospective. In *Proceedings of the 1999 New Security Paradigms Workshop*, Caledon Hills, Ontario, Canada, Sept. 1999. ACM Press.
- [15] M. Feeley. Polling efficiently on stock hardware. In *Proceedings of the conference on Functional programming languages and computer architecture*, Copenhagen, Denmark, June 1993.
- [16] M. Felleisen and R. Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 102:235–271, 1992.
- [17] M. Flatt, R. B. Findler, S. Krishnamurthy, and M. Felleisen. Programming languages as operating systems (or revenge of the son of the Lisp machine). In *Proceedings of the 1999 ACM International Conference on Functional Programming (ICFP '99)*, Paris, France, Sept. 1999. <http://www.cs.rice.edu/CS/PLT/Publications/icfp99-ffkf.ps.gz>.
- [18] L. Gong. *Inside Java 2 Platform Security: Architecture, API Design, and Implementation*. Addison-Wesley, Reading, Massachusetts, June 1999.
- [19] J. Gosling, B. Joy, and G. Steele. *The Java Language Specification*. Addison-Wesley, Reading, Massachusetts, 1996.
- [20] C. Hawblitzel, C.-C. Chang, G. Czajkowski, D. Hu, and T. von Eicken. Implementing multiple protection domains in Java. In *USENIX Annual Technical Conference*, New Orleans, Louisiana, June 1998. USENIX.
- [21] M. Hicks, P. Kakkar, J. T. Moore, C. A. Gunter, and S. Nettles. PLAN: A Packet Language for Active Networks. In *Proceedings of the Third ACM SIGPLAN International Conference on Functional Programming Languages*, pages 86–93. ACM, 1998. <http://www.cis.upenn.edu/~switchware/papers/plan.ps>.
- [22] T. Lindholm and F. Yellin. *The Java Virtual Machine Specification*. Addison-Wesley, Reading, Massachusetts, 1996.
- [23] D. Malkhi, M. Reiter, and A. Rubin. Secure execution of Java applets using a remote playground. In *Proceedings of the 1998 IEEE Symposium on Security and Privacy*, pages 40–51, Oakland, California, May 1998.
- [24] G. McGraw and E. W. Felten. *Securing Java: Getting Down to Business with Mobile Code*. John Wiley and Sons, New York, New York, 1999.
- [25] K. Nilsen, S. Mitra, S. Sankaranarayanan, and V. Thanuvan. Asynchronous Java exception handling in a real-time context. In *IEEE Workshop on Programming Languages for Real-Time Industrial Applications*, Madrid, Spain, Dec. 1998. NewMonics, Inc. <http://www.newmonics.com/news/conferences/rtss/plrtia/async.exception.pdf>.
- [26] E. G. Sirer, R. Grimm, A. J. Gregory, and B. N. Bershad. Design and implementation of a distributed virtual machine for networked computers. In *Proceedings of the Seventeenth ACM Symposium on Operating System Principles*, pages 202–216, Kiawah Island Resort, South Carolina, Dec. 1999. ACM.
- [27] R. Stata and M. Abadi. A type system for Java bytecode subroutines. In *Proceedings of the 25th ACM Symposium on Principles of Programming Languages*, pages 149–160. ACM, Jan. 1998.
- [28] P. Tullman and J. Lepreau. Nested Java processes: OS structure for mobile code. In *Eighth ACM SIGOPS European Workshop*, Sept. 1998.
- [29] L. van Doorn. A secure Java virtual machine. In *Ninth USENIX Security Symposium Proceedings*, Denver, Colorado, Aug. 2000.
- [30] D. S. Wallach, D. Balfanz, D. Dean, and E. W. Felten. Extensible security architectures for Java. In *Proceedings of the Sixteenth ACM Symposium on Operating System Principles*, pages 116–128, Saint-Malo, France, Oct. 1997.
- [31] D. S. Wallach and E. W. Felten. Understanding Java stack inspection. In *Proceedings of the 1998 IEEE Symposium on Security and Privacy*, pages 52–63, Oakland, California, May 1998.