Java bytecode verification: algorithms and formalizations

Xavier Leroy

Java bytecode verification: algorithms and formalizations

Xavier Leroy

INRIA Rocquencourt and Trusted Logic S.A.

`Xavier.Leroy@inria.fr`

Abstract

Bytecode verification is a crucial security component for Java applets, on the Web and on embedded devices such as smart cards. This paper reviews the various bytecode verification algorithms that have been proposed, recasts them in a common framework of dataflow analysis, and surveys the use of proof assistants to specify bytecode verification and prove its correctness.

1 Introduction

Web applets have popularized the idea of downloading and executing untrusted compiled code on the personal computer running the Web browser, without user's approval or intervention. Obviously, this raises major security issues: without appropriate security measures, a malicious applet could mount a variety of attacks against the local computer, such as destroying data (e.g. reformatting the disk), modifying sensitive data (e.g. registering a bank transfer via a home-banking software [5]), divulging personal information over the network, or modifying other programs (Trojan attacks).

Beyond Web services, the applet model is now being transferred to high-security embedded devices such as smart cards: the Java Card architecture [6] allows for post-issuance downloading of applets on smart cards. Smart cards are used as security tokens in sensitive application areas such as payment, mobile telephony, and authentication. This makes the security issues with applets even more acute.

The solution put forward by the Java programming environment is to execute the applets in a so-called "sandbox", which is an insulation layer preventing direct access to the hardware resources and implementing a suitable access control policy [18, 54, 34]. The security of the sandbox model relies on the following three components:

1. Applets are not compiled down to machine executable code, but rather to bytecode for a virtual machine. The virtual machine manipulates higher-level, more secure abstractions of data than the hardware processor, such as object references instead of memory addresses.

2. Applets are not given direct access to hardware resources such as the serial port, but only to a carefully designed set of API classes and methods that perform suitable access control before performing interactions with the outside world on behalf of the applet.

3. Upon downloading, the bytecode of the applet is subject to a static analysis called bytecode verification, whose purpose is to make sure that the code of the applet is well typed and does not attempt to bypass protections 1 and 2 above by performing ill-typed operations at run-time, such as forging object references from integers, illegal casting of an object reference from one class to another, calling directly private methods of the API, jumping in the middle of an API method, or jumping to data as if it were code [19, 58, 33].

Thus, bytecode verification is a crucial security component in the Java "sandbox" model: any bug in the verifier causing an ill-typed applet to be accepted can potentially enable a security attack. At the same time, bytecode verification is a complex process involving elaborate program analyses. Consequently, considerable research efforts have been expended to specify the goals of bytecode verification, formalize bytecode verification algorithms, and prove their correctness.

The purpose of the present paper is to survey this work on bytecode verification. We explain what bytecode verification is, describe the various algorithms that have been proposed, outline the main problems they are faced with, and give references to machine-assisted proofs of correctness.

The remainder of this paper is organized as follows. Section 2 gives a quick overview of the Java virtual machine and of bytecode verification. Section 3 presents the basic bytecode verification algorithm based on dataflow analysis. Sections 4 and 5 concentrate on two delicate verification issues: checking object initialization and dealing with JVM subroutines. Section 6 presents polyvariant verification algorithms that address the subroutine issue. Some issues specific to low-resources embedded systems are discussed in section 7, followed by conclusions and perspectives in section 8.

2   Overview of the JVM and of bytecode verification

The Java Virtual Machine (JVM) [33] is a conventional stack-based abstract machine. Most instructions pop their arguments off the stack, and push back their results on the stack. In addition, a set of registers (also called local variables) is provided; they can be accessed via "load" and "store" instructions that push the value of a given register on the stack or store the top of the stack in the given register, respectively. While the architecture does not mandate it, most Java compilers use registers to store the values of source-level local variables and method parameters, and the stack to hold temporary results during evaluation of expressions. Both the stack and the registers are part of the activation record for a method. Thus, they are preserved across method calls. The entry point for a method specifies the number of registers and stack slots used by the method, thus allowing an activation record of the right size to be allocated on method entry.

Control is handled by a variety of intra-method branch instructions: unconditional branch ("goto"), conditional branches ("branch if top of stack is 0"), multi-way branches (corresponding to the `switch` Java construct). Exception handlers can be specified as a table of (pc_1, pc_2, C, h) quadruples, meaning that if an exception of class C or a subclass of C is raised by any instruction between locations pc_1 and pc_2, control is transferred to the instruction at h (the exception handler).

About 200 instructions are supported, including arithmetic operations, comparisons, object creation, field accesses and method invocations. The example in Figure 1 should give the general flavor of JVM bytecode.

An important feature of the JVM is that most instructions are typed. For instance, the `iadd` instruction (integer addition) requires that the stack initially contains at least two elements, and that these two elements are of type `int`; it then pushes back a result of type `int`. Similarly, a

Source Java code:

```
static int factorial(int n)
{
  int res;
  for (res = 1; n > 0; n--) res = res * n;
  return res;
}
```

Corresponding JVM bytecode:

```
method static int factorial(int), 2 registers, 2 stack slots
    0:  iconst_1          // push the integer constant 1
    1:  istore_1          // store it in register 1 (the res variable)
    2:  iload_0           // push register 0 (the n parameter)
    3:  ifle 14           // if negative or null, go to PC 14
    6:  iload_1           // push register 1 (res)
    7:  iload_0           // push register 0 (n)
    8:  imul              // multiply the two integers at top of stack
    9:  istore_1          // pop result and store it in register 1
   10:  iinc 0, -1        // decrement register 0 (n) by 1
   11:  goto 2            // go to PC 2
   14:  iload_1           // load register 1 (res)
   15:  ireturn           // return its value to caller
```

Figure 1: An example of JVM bytecode

`getfield` $C.f.\tau$ instruction (access the instance field f of type τ declared in class C) requires that the top of the stack contains a reference to an instance of class C or one of its sub-classes (and not, for instance, an integer – this would correspond to an attempt to forge an object reference by an unsafe cast); it then pops it and pushes back a value of type τ (the value of the field f). More generally, proper operation of the JVM is not guaranteed unless the code meets at least the following conditions:

- Type correctness: the arguments of an instruction are always of the types expected by the instruction.

- No stack overflow or underflow: an instruction never pops an argument off an empty stack, nor pushes a result on a full stack (whose size is equal to the maximal stack size declared for the method).

- Code containment: the program counter must always point within the code for the method, to the beginning of a valid instruction encoding (no falling off the end of the method code; no branches into the middle of an instruction encoding).

- Register initialization: a load from a register must always follow at least one store in this register; in other terms, registers that do not correspond to method parameters are not initialized on method entrance, and it is an error to load from an uninitialized register.

- Object initialization: when an instance of a class C is created, one of the initialization methods for class C (corresponding to the constructors for this class) must be invoked before the class instance can be used.

One way to guarantee these conditions is to check them dynamically, while executing the bytecode. This is called the "defensive JVM approach" in the literature [11]. However, checking these conditions at run-time is expensive and slows down execution significantly. The purpose of bytecode verification is to check these conditions once and for all, by static analysis of the bytecode at loading-time. Bytecode that passes verification can then be executed faster, omitting the dynamic checks for the conditions above.

It must be emphasized that bytecode verification by itself does not guarantee secure execution of the code: many crucial properties of the code still need to be checked dynamically, for instance via array bounds checks and null pointer checks in the virtual machine, and access control checks in the API. The purpose of bytecode verification is to shift the verifications listed above from run-time to loading-time.

3   Basic verification by dataflow analysis

The first JVM bytecode verification algorithm is due to Gosling and Yellin at Sun [19, 58, 33]. Almost all existing bytecode verifiers implement this algorithm. It can be summarized as a dataflow analysis applied to a type-level abstract interpretation of the virtual machine. Some advanced aspects of the algorithm that go beyond standard dataflow analysis are described in sections 4 and 5. In this section, we describe the basic ingredients of this algorithm: the type-level abstract interpreter and the dataflow framework.

$$\text{iconst } n \; : \; (S, \; R) \to (\texttt{int}.S, \; R) \text{ if } |S| < M_{stack}$$
$$\texttt{ineg} \; : \; (\texttt{int}.S, \; R) \to (\texttt{int}.S, \; R)$$
$$\texttt{iadd} \; : \; (\texttt{int}.\texttt{int}.S, \; R) \to (\texttt{int}.S, \; R)$$
$$\texttt{iload } n \; : \; (S, \; R) \to (\texttt{int}.S, \; R)$$
$$\qquad \text{if } 0 \le n < M_{reg} \text{ and } R(n) = \texttt{int} \text{ and } |S| < M_{stack}$$
$$\texttt{istore } n \; : \; (\texttt{int}.S, \; R) \to (S, \; R\{n \leftarrow \texttt{int}\}) \text{ if } 0 \le n < M_{reg}$$
$$\texttt{aconst\_null} \; : \; (S, \; R) \to (\texttt{null}.S, \; R) \text{ if } |S| < M_{stack}$$
$$\texttt{aload } n \; : \; (S, \; R) \to (R(n).S, \; R)$$
$$\qquad \text{if } 0 \le n < M_{reg} \text{ and } R(n) <: \texttt{Object} \text{ and } |S| < M_{stack}$$
$$\texttt{astore } n \; : \; (\tau.S, \; R) \to (S, \; R\{n \leftarrow \tau\})$$
$$\qquad \text{if } 0 \le n < M_{reg} \text{ and } \tau <: \texttt{Object}$$
$$\texttt{getfield } C.f.\tau \; : \; (\tau'.S, \; R) \to (\tau.S, \; R) \text{ if } \tau' <: C$$
$$\texttt{putfield } C.f.\tau \; : \; (\tau_1.\tau_2.S, \; R) \to (S, \; R) \text{ if } \tau_1 <: \tau \text{ and } \tau_2 <: C$$
$$\texttt{invokestatic } C.m.\sigma \; : \; (\tau'_n \ldots \tau'_1.S, \; R) \to (\tau.S, \; R)$$
$$\qquad \text{if } \sigma = \tau(\tau_1, \ldots, \tau_n), \; \tau'_i <: \tau_i \text{ for } i = 1 \ldots n, \text{ and } |\tau.S| \le M_{stack}$$
$$\texttt{invokevirtual } C.m.\sigma \; : \; (\tau'_n \ldots \tau'_1.\tau'.S, \; R) \to (\tau.S, \; R)$$
$$\qquad \text{if } \sigma = \tau(\tau_1, \ldots, \tau_n), \; \tau' <: C, \; \tau'_i <: \tau_i \text{ for } i = 1 \ldots n, \; |\tau.S| \le M_{stack}$$

Figure 2: Selected rules for the type-level abstract interpreter. M_{stack} is the maximal stack size and M_{reg} the maximal number of registers.

3.1 The type-level abstract interpreter

At the heart of all bytecode verification algorithms described in this paper is an abstract interpreter for the JVM instruction set that executes JVM instructions like a defensive JVM (including type tests, stack underflow and overflow tests, etc), but operates over types instead of values. That is, the abstract interpreter manipulates a stack of types (a sequence of types) and a register type (a tuple of types associating a type to each register number). It simulates the execution of instructions at the level of types. For instance, for the `iadd` instruction (integer addition), it checks that the stack of types contains at least two elements, and that the top two elements are the type `int`. It then pops the top two elements and pushes back the type `int` corresponding to the result of the addition.

Figure 2 defines more formally the abstract interpreter on a number of representative JVM instructions. The abstract interpreter is presented as a transition relation $i : (S, R) \to (S', R')$, where i is the instruction, S and R the stack type and register type before executing the instruction, and S' and R' the stack type and register type after executing the instruction. Errors such as type mismatches on the arguments, stack underflow, or stack overflow, are denoted by the absence of a transition. For instance, there is no transition on `iadd` from an empty stack.

Notice that method invocations (such as the `invokestatic` and `invokevirtual` instructions in Figure 2) are not treated by branching to the code of the invoked method, like the concrete JVM does, but simply assume that the effect of the method invocation on the stack is as described by the method signature given in the "invoke" instruction. All bytecode verification algorithms described in this paper proceed method per method, assuming that all other methods are well-typed when verifying the code of a method. A simple coinductive argument shows that if this is the case, the
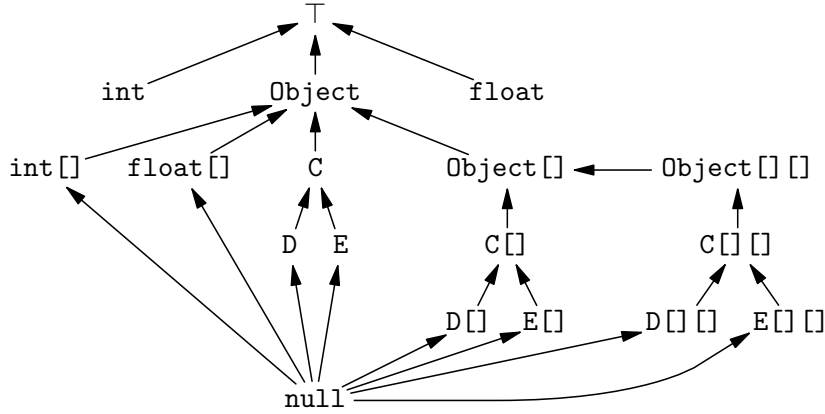
Figure 3: Some type expressions used by the verifier, with their subtyping relation. `C`, `D`, `E` are user-defined classes, with `D` and `E` extending `C`. Not all types are shown.

program as a whole (the collection of all methods) is well typed.

The types manipulated by the abstract interpreter are similar to the source-level types of the Java language. They include primitive types (`int`, `long`, `float`, `double`), array types, and object reference types represented by the fully qualified names of the corresponding classes. The `boolean`, `byte`, `short` and `char` types of Java are identified with `int`. Three extra types are introduced: `null` to represent the type of the null reference, $\bot$ to represent the absence of any value (at unreachable instructions), and $\top$ to represent the contents of uninitialized registers, that is, any value. ("Load" instructions explicitly check that the accessed register does not have type $\top$, thus detecting accesses to uninitialized registers.)

These types are equipped with a subtyping relation, written $<:$, which is essentially identical to the subtyping relation of the Java language (the "assignment compatibility" predicate). Figure 3 illustrates the subtyping relation. Precise definitions of $<:$ can be found in [46, 38, 29, 27], but are omitted here. For the purposes of this article, all we require from the subtyping relation is the following property:

Property 1 *(Well-founded semi-lattice.) The set of types ordered by the $<:$ relation is a semi-lattice: any pair of types has a least upper bound. Moreover, the $<:$ ordering is well founded: there does not exist infinite strictly increasing sequences of types.*

The $<:$ relation between types is extended pointwise to register types and stack types. Two stack types are in the $<:$ relation only if they have the same size. It is easy to see that the $<:$ relation on register types and on stack types is well founded.

The type-level abstract interpreter must satisfy several formal properties. First and foremost is correctness with respect to the dynamic semantics of the defensive JVM: if the abstract interpreter can do a transition $i : (S, R) \to (S', R')$, then for all concrete states (s, r) matching (S, R), the defensive JVM, started in state (s, r), will not stop on a run-time type violation: it will either loop or transition to a state (s', r') that matches (S', R'). This correctness property was formalized and proved by several authors, including Pusch [43] (using the Isabelle/HOL prover), Coglio et al [10] (using SpecWare), Goldberg [17] and Qian [44] (using standard mathematics).

$$r_0 : D \qquad r_0 : E$$
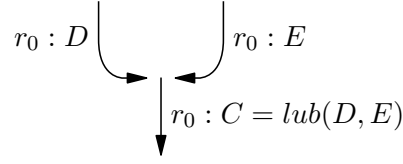
$$r_0 : C = lub(D, E)$$

Figure 4: Handling joins in the control flow

Two other properties of the abstract interpreter are essential to ensure correctness and termination of the bytecode verification algorithms reviewed in this paper: determinacy of transitions, and monotonicity with respect to the subtyping relation.

Property 2 *(Determinacy.)* *The transitions of the abstract interpreter define a partial function: if $i : (S, R) \to (S_1, R_1)$ and $i : (S, R) \to (S_2, R_2)$, then $S_1 = S_2$ and $R_1 = R_2$.*

Property 3 *(Monotonicity.)* *If $i : (S, R) \to (S', R')$, then for all stack types $S_1 <: S$ and register types $R_1 <: R$, there exists a stack type S_1' and a register type R_1' such that $i : (S_1, R_1) \to (S_1', R_1')$, and moreover $S_1' <: S'$ and $R_1' <: R$.*

3.2 The dataflow analysis

Verifying a method whose body is a straight-line piece of code (no branches) is easy: we simply iterate the transition function of the abstract interpreter over the instructions, taking the stack type and register type "after" the preceding instruction as the stack type and register type "before" the next instruction. The initial stack and register types reflect the state of the JVM on method entrance: the stack type is empty; the types of the registers $0 \ldots n - 1$ corresponding to the n method parameters are set to the types of the corresponding parameters in the method signature; the other registers $n \ldots M_{reg} - 1$ corresponding to uninitialized local variables are given the type $\top$.

If the abstract interpreter gets "stuck", i.e. cannot make a transition from one of the intermediate states, then verification fails and the code is rejected. Otherwise, verification succeeds, and since the abstract interpreter is a correct approximation of a defensive JVM, we are certain that a defensive JVM will not get stuck either executing the code. Thus, the code is correct and can be executed safely by a regular, non-defensive JVM.

Branches and exception handlers introduce forks and joins in the control flow of the method. Thus, an instruction can have several predecessors, with different stack and register types "after" these predecessor instructions. Sun's bytecode verifier deals with this situation in the manner customary for data flow analysis: the state (stack type and register type) "before" an instruction is taken to be the least upper bound of the states "after" all predecessors of this instruction. For instance, assume classes D and E extend C, and we analyze a conditional construct that stores a value of type D in register 0 in one arm, and a value of type E in the other arm. (See Figure 4.) When the two arms meet, register 0 is assumed to have type C, which is the least upper bound (the smallest common supertype) of D and E.

More precisely, write $instr(p)$ for the instruction at program point p, $in(p)$ for the state "before" the instruction at p, and $out(p)$ for the state "after" the instruction at p. The verification algorithm

sets up the following forward dataflow equations:

$$instr(p) : in(p) \rightarrow out(p)$$
$$in(p) = lub\{out(q) \mid q \text{ predecessor of } p\}$$

for every valid program point p, plus

$$in(0) = (\varepsilon, (P_0, \ldots, P_{n-1}, \top, \ldots, \top))$$

for the entry point (the P_k are the types of the method parameters). These equations are then solved by standard fixpoint iteration using Kildall's worklist algorithm [35, section 8.4]: a program point p is taken from the worklist and its state "after" $out(p)$ is determined from its state "before" $in(p)$ using the abstract interpreter; then, we replace $in(q)$ by $lub(in(q), out(p))$ for each successor q of p, and enter those successors q for which $in(q)$ changed in the worklist. The fixpoint is reached when the worklist is empty, in which case verification succeeds. Verification fails if a state with no transition is encountered, or one of the least upper bounds is undefined.

As a trivial optimization of the algorithm above, the dataflow equations can be set up at the level of extended basic blocks rather than individual instructions. In other terms, it suffices to keep in working memory the states $in(p)$ where p is the beginning of an extended basic block (i.e. a branch target); the other states can be recomputed on the fly as needed.

The least upper bound of two states is taken pointwise, both on the stack types and the register types. It is undefined if the stack types have different heights, which causes verification to fail. This situation corresponds to a program point where the run-time stack can have different heights depending on the path by which the point is reached; such code cannot be proved correct in the framework described in this section, and must be rejected. (See section 6.2 for an alternate verification algorithm that can handle this situation.)

The least upper bound of two types for a register (or stack slot) can be $\top$, causing this register to have type $\top$ in the merged state. This corresponds to the situation where a register holds values of incompatible types in two arms of a conditional (e.g. `int` in one arm and an object reference in the other), and therefore is treated as uninitialized (no further loads from this register) after the merge point.

Several formalizations and proofs (on paper or on machine) of the bytecode verification algorithm described above have been published. Nipkow and Klein's development in Isabell/HOL [38, 29, 27] is the closest to the dataflow presentation that we gave above. Other formalizations and correctness proofs of the dataflow approach include those of Qian [44], Coglio et al [10], and Stärk et al [49]. Bytecode verification can also be specified and proved sound using type systems [51, 14, 13]; in this framework, the forward dataflow analysis algorithm that we described can be viewed as a type inference algorithm for these type systems. Hartel and Moreau [21] survey other approaches.

3.3 Interfaces and least upper bounds

The dataflow framework presented above requires that the type algebra, ordered by the subtyping relation, constitutes a semi-lattice. That is, every pair of types possesses a smallest common supertype (least upper bound).

Unfortunately, this property does not hold if we take the verifier type algebra to be the Java source-level type algebra (extended with $\top$ and `null`) and the subtyping relation to be the Java

source-level assignment compatibility relation. The problem is that interfaces are types, just like classes, and a class can implement several interfaces. Consider the following classes:

```
interface I  { ... }
interface J  { ... }
class C1 implements I, J   { ... }
class C2 implements I, J   { ... }
```

The subtyping relation induced by these declarations is:

```
                    Object
           I                   J

              C1            C2
```
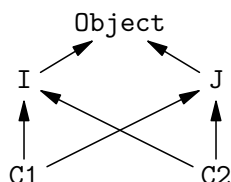
This is obviously not a semi-lattice, since the two types `C1` and `C2` have two common super-types `I` and `J` that are not comparable (neither is subtype of the other).

There are several ways to address this issue. One approach is to manipulate sets of types during verification instead of single types as we described earlier. These sets of types are to be interpreted as conjunctive types, i.e. the set $\{I, J\}$, like the conjunctive type $I \wedge J$, represents values that have both types `I` and `J`, and therefore is a suitable least upper bound for the types $\{C1\}$ and $\{C2\}$ in the example above. This is the approach followed in [17, 44, 43, 49].

Another approach is to complete the class and interface hierarchy of the program into a lattice before performing verification [30]. This is an instance of a general mathematical construction known as the Dedekind-MacNeille completion of a poset. In the example above, the completion would add a point `IandJ` to the lattice of types, which is a subtype of both `I` and `J`, and a supertype of `C1` and `C2`. We then obtain the following semi-lattice:

```
                    Object
           I                   J
                    IandJ
         C1                       C2
```
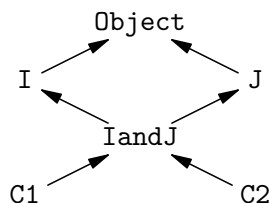
The additional type `IandJ` plays the same role as the type set $\{I, J\}$ in the first approach described above. The difference is that the completion of the class/interface hierarchy is performed once and for all, and verification manipulates only simple types rather than sets of types. This keeps verification simple and fast.

The simplest solution to the interface problem is to be found in Sun's implementation of the JDK bytecode verifier. (This approach is documented nowhere, but can easily be inferred by experimentation.) Namely, bytecode verification ignores interfaces, treating all interface types as the class type `Object`. Thus, the type algebra used by the verifier contains only proper classes and no interfaces, and subtyping between proper classes is simply the inheritance relation between them. Since Java has single inheritance (a class can implement several interfaces, but inherit from one class only), the subtyping relation is tree-shaped and trivially forms a semi-lattice: the least upper bound of two classes is simply their closest common ancestor in the inheritance tree.

The downside of Sun's approach, compared with the set-based or completion-based approach, is that the verifier cannot guarantee statically that an object reference implements a given interface. In particular, the `invokeinterface` $I.m$ instruction, which invokes method m of interface I on an object, is not guaranteed to receive at run-time an object that actually implements I: the only guarantee provided by Sun's verifier is that it receives an argument of type `Object`, that is, any object reference. The `invokeinterface` $I.m$ instruction must therefore check dynamically that the object actually implements I, and raise an exception if it does not.

4   Verifying object initialization

Object creation in the Java virtual machine is a two-step process: first, the instruction `new` C creates a new object, instance of the class C, with all instance fields filled with default values (0 for numerical fields and `null` for reference fields); second, one of the initialization methods for class C (methods named C.`<init>` resulting from the compilation of the constructors of class C) must be invoked on the newly created object. Initialization methods, just like their source-level counterpart (constructors), are typically used to initialize instance fields to non-default values, although they can also perform nearly arbitrary computations.

The JVM specification requires that this two-step object initialization protocol be respected. That is, the object instance created by the `new` instruction is considered uninitialized, and none of the regular object operations (i.e. store the object in a data structure, return it as method result, access one of its fields, invoke one of its methods) is allowed on this uninitialized object. Only when one of the initialization methods for its class is invoked on the new object and returns normally is the new object considered fully initialized and usable like any other object. (Additional restrictions that we will not discuss here are imposed on initialization methods themselves; see [33, 15, 13].)

Unlike the register initialization property, this object initialization property is not crucial to ensure type safety at run-time: since the `new` instruction initializes the instance fields of the new object with correct values for their types, type safety is not broken if the resulting default-initialized object is used right away without having called an initializer method. However, the object initialization property is important to ensure that some invariants between instance fields that is established by the constructor of a class actually hold for all objects of this class.

Static verification of object initialization is made more complex by the fact that initialization methods operate by side-effect: instead of taking an uninitialized object and returning an initialized object, they simply take an uninitialized object, update its fields, and return nothing. Hence, the code generated by Java compilers for the source-level statement `x = new C(arg)` is generally of the following form:

```
new C                      // create uninitialized instance of C
dup                        // duplicate the reference to this instance
code to compute arg
invokespecial C.<init>     // call the initializer
astore 3                   // store initialized object in x
```

That is, two references to the uninitialized instance of C are held on the stack. The topmost reference is "consumed" by the invocation of `C.<init>`. When this initializer returns, the second reference is now at the top of the stack and now references a properly initialized object, which is then stored in the register allocated to `x`. The tricky point is that the initializer method is applied

to one object reference on the stack, but it is another object reference contained in the stack (which happens to reference the same object) whose status goes from "uninitialized" to "fully initialized" in the process.

As demonstrated above, static verification of object initialization requires a form of alias analysis (more precisely a must-alias analysis) to determine which object references in the current state are guaranteed to refer to the same uninitialized object that is passed as argument to an initializer method. While any must-alias analysis can be used, Sun's verifier uses a fairly simple analysis, whereas an uninitialized object is identified by the position (program counter value) of the `new` instruction that created it. More precisely, the type algebra is enriched by the types $\overline{C}_p$ denoting an uninitialized instance of class C created by a `new` instruction at PC p. An invocation of an initializer method $C.$`<init>` checks that the first argument of the method is of type $\overline{C}_p$ for some p, then pops the arguments off the stack type as usual, and finally finds all other occurrences of the type $\overline{C}_p$ in the abstract interpreter state (stack type and register types) and replaces them by C. The following example shows how this works for a nested initialization corresponding to the Java expression `new C(new C(null))`:

```
 0: new C      // stack type after: C̄₀
 3: dup                    //     C̄₀, C̄₀
 4: new C                  //     C̄₀, C̄₀, C̄₄
 7: dup                    //     C̄₀, C̄₀, C̄₄, C̄₄
 8: aconst_null            //     C̄₀, C̄₀, C̄₄, C̄₄, null
 9: invokespecial C.<init> //     C̄₀, C̄₀, C
12: invokespecial C.<init> //     C
15: ...
```

In particular, the first `invokespecial` initializes only the instance created at PC 4, but not the one created at PC 0.

This approach is correct only if at any given time, the machine state contains at most one uninitialized object created at a given PC. Thus, the verifier must prevent situations where several distinct objects created by the same instruction `new` C at p can be "in flight": these would be given the same type $\overline{C}_p$, and initializing one would cause the verifier to assume incorrectly that the others are also initialized. This could potentially happen if a `new` instruction is executed repeatedly as part of a loop; another example involving subroutines is given in [15].

To avoid this problem, Sun's verifier requires that no uninitialized object type appear in the machine state when a backward branch is taken. Freund and Mitchell [15] formalize a simpler, equally effective restriction: they propose that when verifying a `new` C instruction at location p, there must be no occurrence of the type $\overline{C}_p$ in the stack type and register type at p. Bertot [3] proves the correctness of this approach using the Coq theorem prover, and extracts a verification algorithm from the proof. (Bertot's algorithm is not a standard dataflow analysis, since it features an additional pass of constraint solving.)

Both Sun's restriction and Freund and Mitchell's restriction are not monotone, in the sense that raising the type of a register or stack location from $\overline{C}_p$ to $\top$ can transform an abstract state from which no transitions are possible into an abstract state from which a transition is possible. In other terms, property 3 does not hold, and this causes difficulties with dataflow analysis. To address this issue, a modified formulation of Freund and Mitchell's restriction is introduced in [49] and used in [29, 27, 13]: a `new` C instruction at location p does not fail if there are occurrences of the type $\overline{C}_p$ in the stack and register types before p; instead, it sets these entries to $\top$ in the resulting state.

This suffices to enforce non-aliasing (the unitialized objects whose types were set to $\top$ cannot be initialized nor used later on), while respecting the monotonicity property.

Subroutines, as described in section 5, complicate the verification of object initialization. As discovered by Freund and Mitchell [15], a `new` instruction inside a subroutine can result in distinct uninitialized objects having the same static type $\overline{C}_p$, thus fooling Sun's verifier into believing that all of them become initialized after invoking an initialization method on one of them. The solution is to prohibit or set to $\top$ all registers and stack locations that have type $\overline{C}_p$ across a subroutine call.

Coglio [9] observes that Sun's restriction on backward branches as well as Freund and Mitchell's restriction on `new` are unnecessary for a bytecode verifier based on monovariant dataflow analysis. More precisely, [9, section 5.8.2] shows that, in the absence of subroutines, a register or stack location cannot have the type $\overline{C}_p$ just before a program point p containing a `new` C instruction. Thus, the only program points where uninitialized object types in stack types or register types must be prohibited (or turned into $\top$) are subroutine calls.

5 Subroutines

Subroutines in the JVM are code fragments that can be called from several points inside the code of a method. To this end, the JVM provides two instructions: `jsr` branches to a given label in the method code and pushes a return address to the following instruction; `ret` recovers a return address (from a register) and branches to the corresponding instruction. Subroutines are used to compile certain exception handling constructs, and can also be used as a general code-sharing device. The difference between a subroutine call and a method invocation is that the body of the subroutine executes in the same activation record as its caller, and therefore can access and modify the registers of the caller.

5.1 The verification problem with subroutines

Subroutines complicate significantly bytecode verification by dataflow analysis. First, it is not obvious to determine the successors of a `ret` instruction, since the return address is a first-class value. As a first approximation, we can say that a `ret` instruction can branch to any instruction that follows a `jsr` in the method code. (This approximation is too coarse in practice; we will describe better approximations later.) Second, the subroutine entry point acts as a merge point in the control-flow graph, causing the register types at the points of call to this subroutine to be merged. This can lead to excessive loss of precision in the register types inferred, as the example in Figure 5 shows.

The two `jsr 100` at 0 and 52 have 100 as successor. At 0, register 0 has type $\top$; at 52, it has type `int`. Thus, at 100, register 0 has type $\top$ (the least upper bound of $\top$ and `int`). The subroutine body (between 101 and 110) does not modify register 0, hence its type at 110 is still $\top$. The `ret 1` at 110 has 3 and 55 as successors (the two instructions following the two `jsr 100`). Thus, at 55, register 0 has type $\top$ and cannot be used as an integer by instructions 55 and 56. This code is therefore rejected.

This behavior is counter-intuitive. Calling a subroutine that does not use a given register does not modify the run-time value of this register, so one could expect that it does not modify the

```
                       // register 0 uninitialized here
   0: jsr 100          // call subroutine at 100
   3: ...

  50: iconst_0
  51: istore_0         // register 0 has type "int" here
  52: jsr 100          // call subroutine at 100
  55: iload_0          // load integer from register 0
  56: ireturn          // and return to caller
      ...
                       // subroutine at 100:
 100: astore_1         // store return address in register 1
 101: ...              // execute some code that does not use register 0
 110: ret 1            // return to caller
```

Figure 5: An example of subroutine

verification-time type of this register either. Indeed, if the subroutine body was expanded inline at the two `jsr` sites, bytecode verification would succeed as expected.

The subroutine-based compilation scheme for the `try...finally` construct produces code very much like the above, with a register being uninitialized at one call site of the subroutine and holding a value preserved by the subroutine at another call site. Hence it is crucial that similar code passes bytecode verification. In the remainder of this section and in section 6, we will present refinements of the dataflow-based verification algorithm that achieve this goal.

5.2 Sun's solution

We first describe the approach implemented in Sun's JDK verifier. It is described informally in [33, section 4.9.6], and formalized (with various degrees of completeness and faithfulness to Sun's implementation) in [51, 44, 49, 13, 50]. This approach implements the intuition that a call to a subroutine should not change the types of registers that are not used in the subroutine body.

First, we need to make precise what a "subroutine body" is: since JVM bytecode is unstructured, subroutines are not syntactically delimited in the code; subroutine entry points are easily detected (as targets of `jsr` instructions), but it is not immediately apparent which instructions can be reached from a subroutine entry point. Thus, a dataflow analysis is performed, either before or in parallel with the main type analysis. The outcome of this analysis is a consistent labeling of every instruction by the entry point(s) for the subroutine(s) it logically belongs to. From this labeling, we can then determine, for each subroutine entry point ℓ, the return instruction $Ret(\ell)$ for the subroutine, and the set of registers $Used(\ell)$ that are read or written by instructions belonging to that subroutine.

The dataflow equation for subroutine calls is then as follows. Let i be an instruction `jsr` ℓ, and j be the instruction immediately following i. Let $(S_{jsr}, R_{jsr}) = out(i)$ be the state "after" the `jsr`, and $(S_{ret}, R_{ret}) = out(Ret(\ell))$ be the state "after" the `ret` that terminates the subroutine. Then:

$$in(j) = \left(S_{ret}, \{ r \mapsto \begin{cases} R_{ret}(r) & \text{if } r \in Used(\ell) \\ R_{jsr}(r) & \text{if } r \notin Used(\ell) \end{cases} \} \right)$$

13

In other terms, the state "before" the instruction j following the $\mathtt{jsr}$ is identical to the state "after" the $\mathtt{ret}$, except for the types of the registers that are not used by the subroutine, which are taken from the state "after" the $\mathtt{jsr}$.

In the example above, we have $Ret(100) = 110$ and register 0 is not in $Used(100)$. Hence the type of register 0 before instruction 55 (the instruction following the $\mathtt{jsr}$) is equal to the type after instruction 52 (the $\mathtt{jsr}$ itself), that is $\mathtt{int}$, instead of $\top$ (the type of register 0 after the $\mathtt{ret\ 1}$ at 110).

While effective in practice, Sun's approach to subroutine verification raises a challenging issue: determining the subroutine structure is difficult. Not only are subroutines not syntactically delimited, but return addresses are stored in general-purpose registers rather than on a subroutine-specific stack, which makes tracking return addresses and matching $\mathtt{ret}/\mathtt{jsr}$ pairs more difficult. To facilitate the determination of the subroutine structure, the JVM specification states a number of restrictions on correct JVM code, such as "two different subroutines cannot 'merge' their execution to a single $\mathtt{ret}$ instruction" [33, section 4.9.6]. These restrictions seem rather ad-hoc and specific to the particular subroutine labeling algorithm that Sun's verifier uses. Moreover, the description of subroutine labeling given in the JVM specification is very informal and incomplete.

Several rational reconstructions and formalizations of this part of Sun's verifier have been published. The presentations closest to Sun's implementation are due to Qian [44] and Stärk et al [49]. A characteristic feature of Sun's implementation, correctly captured by these presentations, is that the subroutine structure and the $Used(\ell)$ sets are not determined prior to setting up and solving the dataflow equations, as we suggested above; instead, the types and the $Used(\ell)$ sets are inferred simultaneously during the dataflow analysis. This simultaneous computation of types and $Used(\ell)$ sets complicates the analysis: as shown by Qian [45], the transfer function of the dataflow analysis is no longer monotonous, and special iteration strategies are required to reach the fixpoint.

5.3 Other approaches to the verification of subroutines

Many alternatives to Sun's verification of subroutine have been proposed in the literature, often in the context of small subsets of the JVM, leading to algorithms that are simpler and more elegant, but do not scale to the whole JVM.

The first formal work on subroutine verification is due to Abadi and Stata [51]: it relies on a separate analysis, performed before type verification proper, that labels bytecode instructions with the names of the subroutines they belong to, thus reconstructing the subroutine structure. The $Used(\ell)$ sets can then be computed for each subroutine ℓ, and injected in the dataflow equations as described in section 5.2.

Later work by Hagiya and Tozawa [20] also relies on prior determination of the subroutine structure, but expresses the flow of types through subroutines in a different way, using special types $\mathtt{last}(n)$ to refer to the type of register n in the caller of the subroutine. The $\mathtt{last}(n)$ types behave very much like type variables in a type system featuring parametric polymorphism. In other terms, the subroutine is given a type that is polymorphic over the types of the local variables that it does not use.

While these works shed considerable light on the issue of subroutines, they are carried in the context of a small subset of the JVM that excludes exceptions and object initialization in particular. The delicate interactions between subroutines and object initialization were discussed in section 4. As for exceptions, exception handling complicates significantly the determination of the subroutine

structure. Examination of bytecode produced by a Java compiler shows two possible situations: either an exception handler covers a range of instructions entirely contained in a subroutine, in which case the code of the exception handler should be considered as part of the same subroutine (e.g. it can branch back to the `ret` instruction that terminates the subroutine); or, an exception handler covers both instructions belonging to a subroutine and non-subroutine instructions, in which case the code of the handler should be considered as outside the subroutine. The problem is that in the second case, we have a branch (via the exception handler) from a subroutine instruction to a non-subroutine instruction, and this branch is not a `ret` instruction; this situation is not allowed in Abadi and Stata's subroutine labeling system.

Thus, it is desirable to develop subroutine verification strategies that do not rely on prior determination of the subroutine structure, but instead "discover" this structure during verification. The polyvariant verification algorithms that we discuss next in section 6 satisfy this requirement. O'Callahan [39] proposes a different approach, based on continuation types: the type assigned to a return address contains the complete stack and register type expected at the program point following the `jsr` instruction. However, O'Callahan gives only type checking rules, but no effective type inference algorithm.

6 Polyvariant bytecode verification

In this paper so far, we have presented bytecode verification as a *monovariant* flow analysis: at each program point, only one abstract state (register type and stack type) is considered. *Polyvariant* flow analyses, also called context-sensitive analyses [37, section 3.6] lift this restriction: several states are allowed per program point. As we show in this section, polyvariant bytecode verification provides an alternate solution to the subroutine problem of section 5: polyvariant analysis allows instructions inside subroutine bodies to be analyzed several times, e.g. once per call site for the subroutine, without merging the corresponding states like the monovariant analysis of section 3.2 does.

6.1 Polyvariant verification based on contours

The first polyvariant bytecode verification analysis that we describe is based on *contours*, and is used in the Java Card off-card verifier [53]. In contour-based polyvariant flow analysis, the distinct states maintained at each program point are indexed by contours that usually approximate the control-flow path that led to each state.

In the case of bytecode verification, contours are subroutine call stacks: lists of return addresses for the sequence of `jsr` instructions that led to the corresponding state. In the absence of subroutines, all the bytecode for a method is analyzed in the empty contour. Thus, only one state is associated to each instruction and the analysis degenerates into the monovariant dataflow analysis of section 3.2. However, when a `jsr` ℓ instruction is encountered in the current contour c, it is treated as a branch to the instruction at ℓ in the augmented contour $\ell.c$. Similarly, a `ret` r instruction is treated as a branch that restricts the current context c by popping one or several return addresses from c (as determined by the type of the register r).

In the example of Figure 5, the two `jsr 100` instructions are analyzed in the empty context ε. This causes two "in" states to be associated with the instruction at 100; one has contour $3.\varepsilon$, assigns

type $\top$ to register 0, and contains `retaddr(3)` at the top of the stack[1]; the other state has contour $55.\varepsilon$, assigns type `int` to register 0, and contains `retaddr(55)` at the top of the stack. Then, the instructions at $101\ldots110$ are analyzed twice, in the two contours $3.\varepsilon$ and $55.\varepsilon$. In the contour $3.\varepsilon$, the `ret 1` at 110 is treated as a branch to 3, where register 0 still has type $\top$. In the contour $55.\varepsilon$, the `ret 1` is treated as a branch to 55 with register 0 still having type `int`. By analyzing the subroutine body in a polyvariant way, under two different contours, we avoided merging the types $\top$ and `int` of register 0 at the subroutine entry point, and thus obtained the desired type propagation behavior for register 0: $\top$ before and after the `jsr 100` at 3, but `int` before and after the `jsr 100` at 52.

More formally, the polyvariant dataflow equation for a `jsr` ℓ instruction at i followed by an instruction at j is

$$in(\ell, j.c) = (\mathtt{retaddr}(j).S, T) \text{ where } (S, T) = out(i, c)$$

For a `ret` r instruction at i, the equation is

$$in(j, c') = out(i, c)$$

where the type of register r in the state $out(i, c)$ is `retaddr`(j) and the context c' is obtained from c by popping return addresses until j is found, that is, $c = c''.j.c'$. Finally, for instructions i other than `jsr` and `ret`, the equation is simply

$$i : in(i, c) \rightarrow out(i, c)$$

that is, the instruction triggers no change of contour.

Another way to view polyvariant verification is that it is exactly equivalent to performing monovariant verification on an expanded version of the bytecode where every subroutine call has been replaced by a distinct copy of the subroutine body. Instead of actually taking N copies of the subroutine body, we analyze them N times in N different contours. Of course, duplicating subroutine bodies before the monovariant verification is not practical, because it requires prior knowledge of the subroutine structure (to determine which instructions are part of which subroutine body), and as shown in section 5.3, the subroutine structure is hard to determine exactly. The beauty of the polyvariant analysis is that it determines the subroutine structure along the way, via the computations on contours performed during the dataflow analysis. Moreover, this determination takes advantage of typing information such as the `retaddr`(ra) types to determine with certainty the point to which a `ret` instruction branches in case of an early return from nested subroutines.

Another advantage of polyvariant verification is that it handles code that is reachable both from subroutine bodies and from the main program, such as the exception handlers mentioned in section 5.3: rather than deciding whether such exception handlers are part of a subroutine or not, the polyvariant analysis simply analyzes them several times, once in the empty contour and once or several times in subroutine contours.

One downside of polyvariant verification is that it is more computationally expensive than Sun's approach. In particular, if subroutines are nested to depth N, and each subroutine is called k times, the instructions from the innermost subroutine are analyzed k^N times instead of only once in Sun's algorithm. However, typical Java code has low nesting of subroutines: most methods have $N \leq 1$, very few have $N = 2$, and $N > 2$ is unheard of. Hence, the extra cost of polyvariant verification is entirely acceptable in practice.

---

[1]The type `retaddr`(i) represents a return address to the instruction at i. It is subtype of $\top$ and itself only.

```
while (true) {                    0:   invokestatic m
  try {                           3:   jsr 15
    m();                          6:   goto 0
  } finally {                     9:   astore 0     // exception handler
    continue;                     10:  jsr 15
  }                               13:  aload 0
}                                 14:  athrow
                                  15:  astore 1     // subroutine
                                  16:  goto 0
```

Exception table: $[0, 2] \rightarrow 9$

Figure 6: The problem with contour-based polyvariant verification (left: Java source code; right: corresponding JVM bytecode)

A more serious drawback of this approach is that contour-based polyvariant verification can fail to accept valid JVM code because the subroutine structure that it infers through the construction of contours can be infinite. Consider the Java method shown in Figure 6, along with the corresponding JVM code. The subroutine at 15 does not terminate by a `ret` instruction, but by a `goto 0` that branches back to the non-subroutine part of the method code. (This branch corresponds to the `continue` statement in the source code.)

Polyvariant verification starts by analyzing instructions at 0, 3 and 6 in the empty contour ε. The `jsr 15` at 3 causes instructions at 15 and 16 to be analyzed in the contour $15.\varepsilon$. The `goto 0` at 16 does not affect the contour (unlike a `ret` instruction), thus causing instructions at 0, 3 and 6 to be re-analyzed in the contour $15.\varepsilon$. Iterating this process, we end up analyzing instructions 0, 3 and 6 in an infinite number of contexts: ε, $15.\varepsilon$, $15.15.\varepsilon$, ... Thus, verification does not terminate.

Alternatively, termination can be ensured by requiring that contours never contain twice the same subroutine label: a `jsr ℓ` in a contour containing ℓ is rejected, in accordance with the JVM specification that states that subroutines cannot be recursive. (This is what the Java Card verifier [53] does.) But in this case we end up rejecting a valid piece of JVM bytecode, generated by compilation of a valid (albeit artificial) Java program.

As demonstrated by this example, using data-independent contours to direct polyvariant verification can cause too many – or even infinitely many – different states to be kept for a given program point, even if these states are exactly identical. (In the example, the states at point 0 in contours $15.\varepsilon$, $15.15.\varepsilon$, etc, are all identical.) We now describe an alternate approach to polyvariant verification, not using contours, that avoids these issues.

6.2   Model checking of abstract interpretations

It is folk lore that dataflow analyses can be viewed as model checking of abstract interpretations [48]. Since a large part of bytecode verification is obviously an abstract interpretation (of a defensive JVM at the type level), it is natural to look at the remaining parts from a model-checking perspective.

Posegga and Vogt [40] were the first to do so. They outline an algorithm that takes the bytecode for a method and generates a temporal logic formula that holds if and only if the bytecode is safe. They then use an off-the-shelf model checker to determine the validity of the formula. While this

application uses only a small part of the power and generality of temporal logic and of the model checker, the approach sounds interesting for establishing finer properties of the bytecode that go beyond the basic safety properties of bytecode verification (see section 8). The article by Basin *et al.* in this volume [2] explores the model-checking approach to bytecode verification in great details.

Brisset [4] and independently Coglio [8] extract the essence of the "model checking" approach: the idea of exploring all reachable states of the abstract interpreter. They consider the transition relation obtained by combining the transition relation of the type-level abstract interpreter (Figure 2) with the "successor" relation between instructions. This relation is of the form $(p, S, R) \to (p', S', R')$, meaning that the abstract interpreter, started at PC p with stack type S and register type R, can abstractly execute the instruction at p and arrive at PC p' with stack type S' and register type R'. Additional transitions $(p, S, R) \to \texttt{err}$ are introduce to reflect states (p, S, R) in which the abstract interpreter is "stuck" (cannot make a transition because some check failed).

$$(p, S, R) \quad \to \quad (p', S', R') \quad \text{if } instr(p) : (S, R) \to (S', R')$$
$$\text{and } p' \text{ is a successor of } p$$
$$(p, S, R) \quad \to \quad \texttt{err} \quad \text{if } instr(p) : (S, R) \not\to$$

The BC (Brisset-Coglio) verification algorithm simply explores all states reachable by repeated applications of the extended transition relation starting with the initial state $\sigma_0 = (0, \varepsilon, (P_0, \dots, P_{n-1}, \top, \dots, \top))$ corresponding to the method entry. In other terms, writing $C(\Sigma) = \Sigma \cup \{\sigma' \mid \exists \sigma \in \Sigma. \; \sigma \to \sigma'\}$ for the one-step closure of a set of states Σ, the BC algorithm computes by fixpoint iteration the smallest set Σ_c containing σ_0 and closed under C. If $\texttt{err} \in \Sigma_c$, the error state is reachable and the bytecode is rejected. Otherwise, the bytecode passes verification. In the latter case ($\texttt{err}$ is not reachable), the correctness of the abstract interpretation (as proved in [43]) guarantees that the concrete, defensive JVM interpreter will never fail a safety check during the execution of the method code, hence the bytecode is safe.

This algorithm always terminates because the number of distinct states is finite (albeit large), since there is a finite number of distinct types used in the program, and the height of the stack is bounded, and the number of registers is fixed.

The problem with subroutines described in section 5.1 completely disappears in this approach. It suffices to use the following transitions for the $\texttt{jsr}$ and $\texttt{ret}$ instructions:

$$(p, S, R) \quad \to \quad (\ell, \texttt{retaddr}(p+3).S, R) \quad \text{if } instr(p) = \texttt{jsr } \ell$$
$$(p, S, R) \quad \to \quad (q, S, R) \quad \text{if } instr(p) = \texttt{ret } r \text{ and } R(r) = \texttt{retaddr}(q)$$
$$(p, S, R) \quad \to \quad \texttt{err} \quad \text{if } instr(p) = \texttt{ret } r \text{ and } R(r) \neq \texttt{retaddr}(q)$$

The fact that the BC algorithm never merges the types inferred along two execution paths leading to the same instruction guarantees that subroutine bodies are analyzed as many times as necessary to propagate type information correctly across subroutine calls. However, we will never consider twice the same stack and register types at a given point, thus guaranteeing termination and avoiding the problem with contour-based polyvariance. For instance, the example of Figure 6 that cannot be accepted by contour-based polyvariant verification is now correctly accepted: instructions at 0, 3 and 6 are verified exactly twice, once under the assumption $r_0 : \top$, the other under the assumption $r_0 : \texttt{retaddr}(3)$.

Another interest of this approach is that it allows us to reconsider some of the design decisions explained in sections 3.2, 3.3 and 4. In particular, the BC algorithm never computes least upper

$$r_0 : D \qquad r_0 : E$$

$$r_0 : D$$
or
$$r_0 : E$$

$$r_1 : D \qquad r_1 : E$$
$$r_2 : D \qquad r_2 : E$$
$$\vdots$$
$$r_n : D \qquad r_n : E$$

2^n register types
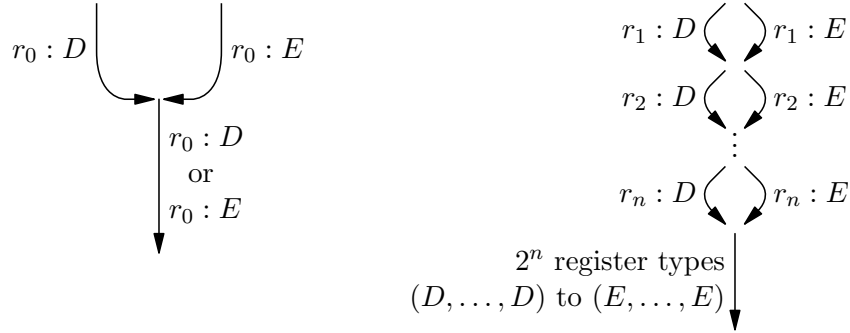$(D, \ldots, D)$ to $(E, \ldots, E)$

Figure 7: Control-flow joins in the model-checking approach

bounds of types, but simply checks subtyping relations between types. Thus, it can be applied to any well-founded subtyping relation, not just relations that form a semi-lattice. Indeed, it can keep track of interface types and verify `invokeinterface` instructions accurately, without having to deal with sets of types or lattice completion. Similarly, it is possible to verify code where the stack size is not the same on all paths leading to an instruction.

Brisset [4] formalized and proved the correctness of this approach in the Coq proof assistant, and extracted the ML code of a bytecode verifier from the proof. Klein and Wildmoser [26] also prove the correctness of this approach using Isabelle/HOL. Their proof builds on a generic dataflow analysis framework, and thus seems reusable for the variants of the BC algorithm discussed in section 6.3. Coglio [8] argues that the BC verification algorithm is the most precise of all "reasonable" verification algorithms, in the sense that it accepts all bytecode that does not crash a defensive virtual machine that would follow all execution paths across conditional jumps, regardless of the value of the condition. Thus, the only correct bytecode that could be rejected by the BC algorithm is bytecode whose correctness depends on additional knowledge on the values (and not just the types) of booleans, integers and object references manipulated by the code.

6.3 Widening in the model checking approach

The BC verification algorithm based on model checking, described in section 6.2, can be impractical: it runs in time exponential in the number of conditional or N-way branches in the method. Consider the control-flow joint depicted in Figure 7, left part. While the dataflow-based algorithms verify the instructions following the join point only once under the assumption $r : C$ where $C = lub(D, E)$, the BC algorithm verifies them twice, once under the assumption $r : D$, once under the assumption $r : E$. Consider now the control-flow graph shown in the right part of Figure 7. It comprises N such conditional constructs in sequence, each assigning a different type to registers $r_1 \ldots r_N$. This causes the instructions following the last conditional to be verified 2^N times under 2^N different register types.

One way to improve the efficiency of the BC algorithm is to reduce the number of states that need be explored via judicious application of *widening steps*: some transitions $(pc, S, R) \rightarrow (pc', S', R')$ can be replaced by $(pc, S, R) \rightarrow (pc', S'', R'')$ where $R' <: R''$ and $S' <: S''$. If the error state `err` is still not reachable, the bytecode remains safe. A widening step such as the above reduces the total number of states to be explored if we judiciously choose R'' and S'', e.g. if the widened state

(pc', S'', R'') is already reachable.

More formally, define a subtyping relation between states by $\sigma <: \mathtt{err}$ for all states σ, and $(p, S, R) <: (p', S', R')$ if $p = p'$, $S <: S'$ and $R <: R'$. A widening scheme W is a function from sets of states to sets of states, such that for all state $\sigma \in \Sigma$, there exists a state $\sigma' \in W(\Sigma)$ such that $\sigma <: \sigma'$. Note that this property implies that if $\mathtt{err} \in \Sigma$, then $\mathtt{err} \in W(\Sigma)$.

Bytecode verification by "widened model checking", then, computes by fixpoint iteration of $W \circ C$ the smallest state set Σ_w containing the initial state σ_0 and such that $W(C(\Sigma_w)) = \Sigma_w$. The bytecode passes verification if and only if $\mathtt{err} \notin \Sigma_w$.

It is easy to show that all bytecode accepted by a widened model checking algorithm is also accepted by the BC model checking algorithm, and is therefore type-safe at runtime. The proof is given in appendix A.

As proposed in [23], we can construct interesting examples of widening functions from an equivalence relation $\approx$ determining which states should be merged in a single state, and a merging function M from sets of equivalent states to states such that $\sigma <: M(\Sigma)$ for all $\sigma \in \Sigma$. The widening function W is then defined by

$$W(\Sigma) = \{M(\Sigma_1), \ldots, M(\Sigma_n)\}$$

where $\Sigma_1, \ldots, \Sigma_n$ is a partition of Σ in equivalence classes for the $\approx$ relation. The W function thus defined is a widening function, since for every $\sigma \in \Sigma$, writing Σ_i for the equivalence class of σ in Σ, we have $\sigma <: M(\Sigma_i) \in W(\Sigma)$.

The monovariant dataflow analysis of section 3.2 is a trivial instance of widened model checking where the widening function W merges all stack and register types associated with the same program point by taking their least upper bound. More precisely, the widening function is constructed as described above from the following equivalence relations and merging functions. The equivalence relation $\approx$ is defined by $\mathtt{err} \approx \mathtt{err}$ and $(p_1, S_1, R_1) \approx (p_2, S_2, R_2)$ if and only if $p_1 = p_2$. The merging function is defined by

$$
\begin{aligned}
M(\{\mathtt{err}\}) &= \mathtt{err} \\
M(\{(p, S_1, R_1), \ldots, (p, S_n, R_n)\}) &= (p, lub(S_1, \ldots, S_n), lub(R_1, \ldots, R_n)) \\
&\quad \text{if } |S_1| = \ldots = |S_n| \\
M(\{(p, S_1, R_1), \ldots, (p, S_n, R_n)\}) &= \mathtt{err}, \text{ otherwise}
\end{aligned}
$$

A more interesting widening function is described by Henrio and Serpette [23]. Similar ideas were implemented by Frey [16] and by Coglio [8]. They propose to merge stack and register types at a given program point if and only if these types agree on return addresses, that is, contain the same $\mathtt{retaddr}(p)$ types in the same registers or stack slots, but may differ on non-$\mathtt{retaddr}$ types.

More precisely, we say that two types agree on return addresses if they are equal, or none of them is a $\mathtt{retaddr}$ type. Two register types agree on return addresses if the types they assign to every register agree. Two stack types agree on return addresses if they have the same size, and the types they contain agree pointwise. Then, define the equivalence relation $\approx$ by $\mathtt{err} \approx \mathtt{err}$ and $(p_1, S_1, R_1) \approx (p_2, S_2, R_2)$ if and only if $p_1 = p_2$ and S_1, S_2 agree on return addresses, and R_1, R_2 agree on return addresses. The merging function is defined by

$$
\begin{aligned}
M(\{\mathtt{err}\}) &= \mathtt{err} \\
M(\{(p, S_1, R_1), \ldots, (p, S_n, R_n)\}) &= (p, lub(S_1, \ldots, S_n), lub(R_1, \ldots, R_n))
\end{aligned}
$$

(By construction of $\approx$, M is never presented with stack types of different heights.) The CFHS (Coglio-Frey-Henrio-Serpette) verification algorithm, then, corresponds to widened model checking using the widening function induced by the $\approx$ relation and M merging function.

In bytecode that contains no subroutines, stack and register types never contain return address types, thus always agree on return addresses; in this case, the CFHS algorithm essentially reduces to standard, efficient monovariant dataflow analysis[2]. In the presence of subroutines, however, sufficient polyvariance is ensured because return address types prevent type merging inside subroutines. In the example of Figure 5, analysis of the two `jsr` at 0 and 52 causes two states with $pc = 100$ to be explored: one has `retaddr(3)` at the top of the stack type, and $\top$ as the type of register 0; the other has `retaddr(55)` at the top of the stack type, and `int` as the type of register 0. These two states disagree on return addresses (namely, on the top element of the stack type), thus are not merged. This causes the `astore_1` at 100 to be analyzed twice, resulting in two states at 101 that are not mergeable either because one has `retaddr(3)` for register 1 and the other has `retaddr(55)` for register 1. Thus, the instructions in the subroutine body (from 100 to 110) are given two independent sets of stack and register types, as needed to verify the two subroutine calls with sufficient precision.

In appendix B, we prove that the CFHS algorithm accepts exactly the same bytecode as the BC algorithm. In other terms, it turns out that the additional merging of type information performed in the CFHS algorithm does not degrade the precision of the analysis. This is a consequence of the following "continuity" property of the type-level abstract interpreter with respect to least upper bounds.

Property 4 *(Continuity.) Assume that the stack types S_1 and S_2 agree on return addresses, and the register types R_1 and R_2 agree on return addresses. Further assume that $i : (S_1, R_1) \rightarrow (S'_1, R'_1)$ and $i : (S_2, R_2) \rightarrow (S'_2, R'_2)$. Then, $i : (lub(S_1, S_2), lub(R_1, R_2)) \rightarrow (lub(S'_1, S'_2), lub(R'_1, R'_2))$.*

7 Bytecode verification on small computers

Java virtual machines run not only in personal computers and workstations, but also in a variety of embedded computers, such as personal digital assistants, mobile phones, and smart cards. Extending the Java model of safe post-issuance code downloading to these devices requires that bytecode verification be performed on the embedded system itself. However, bytecode verification is an expensive process that can exceed the resources (processing power and working memory space) of small embedded systems. For instance, a typical Java card (Java-enabled smart card) has 1 or 2 *kilo*-bytes of working memory, 16 to 32 kilo-bytes of rewritable persistent memory[3], and an 8-bit microprocessor that is approximately 1000 times slower than a personal computer.

On such small computing devices, a conventional bytecode verification algorithm (such as Sun's or any of the polyvariant verifiers described in this paper) cannot run in working memory: the

---

[2]One minor difference with monovariant dataflow analysis is that the latter fails if the stack height differs on several paths that reach a given program point, while the CFHS algorithm does not fail immediately in this case, but analyzes the join point and its successors in a polyvariant manner, once per different stack height. At any rate, subroutine-free bytecode that passes monovariant verification also passes the CFHS algorithm, with only one state being considered for each program point.

[3]Rewritable persistent memory (EEPROM or Flash) is distinct from working memory (RAM) in that writes to the former take 1000–10000 times longer than writes to the latter. Moreover, rewritable persistent memory allows a limited number of writes to the same memory location.

amount of RAM available is too small to store the stack and register types inferred during bytecode verification. One avenue is to use persistent rewritable memory instead of RAM to hold the data structures of the verifier. It was long believed that this solution was impractical, because these data structures change rapidly during bytecode verification, resulting in excessive writing times and "wear" on the persistent memory. Recent work by Deville and Grimaud [12] develops specially-designed encodings of types and memory representations of the verification data that alleviate this issue, and suggests that this approach is indeed feasible.

Another avenue for fitting a bytecode verifier into a small computing device is to design new verification algorithms that can run in tiny amounts of working memory. We now discuss two such algorithms.

7.1   Lightweight bytecode verification using certificates

Inspired by Necula and Lee's proof-carrying code [36], Rose and Rose [47] propose to split bytecode verification into two phases: the code producer computes the stack and register types at branch targets and transmit these so-called certificates along with the bytecode; the embedded system, then, simply checks that the code is well-typed with respect to the types given in the certificates, rather than inferring these types itself. In other terms, the embedded system no longer solves iteratively the dataflow equations characterizing correct bytecode, but simply checks (in a single, linear pass over the bytecode) that the types provided in the code certificates are indeed a solution of these equations.

The benefits of this approach are twofold. First, checking a solution is faster than inferring one, since we avoid the cost of the fixpoint iteration. This speeds up verification to some extent[4]. Second, certificates are only read, but never modified during verification. Hence, they can be stored in persistent rewritable memory without risking to "wear" this memory space by repeated rewriting of the data during verification.

A practical limitation of this approach is that certificates are relatively large: about 50% of the size of the code they annotate [31]. Even if certificates are stored in persistent memory, they can still exceed the available memory space.

Later work by Rose [46] addresses the issue of the certificate size by noticing that type information for certain branch targets can be omitted from the certificate, as long as this type information is correctly computed by the verification algorithm during its linear scan of the bytecode. This hybrid verification strategy thus reduces the size of the certificate, at the cost of increased working memory requirements to store the inferred type information for branch targets not described in the certificate.

Lightweight bytecode verification is used in the KVM, one of Sun's embedded variants of the JVM [52]. It was formalized and proved sound and complete by Klein and Nipkow [28, 27], using the Isabelle/HOL prover. These presentations differ in their treatment of subroutines. Rose's presentation [47, 46] deals only with monovariant verification and does not handle subroutines. The KVM implementation follows this approach, and relies on subroutine expansion before verification, as part of the certificate generation phase. Klein and Nipkow's formalization, however, is applicable not only to monovariant verification, but to any verification algorithm that can be expressed using dataflow equations, including the polyvariant algorithms of section 6.

---

[4]The speedup is not as important as one might expect, since experiments show that the fixpoint is usually reached after examining every instruction at most twice [31].

7.2 On-card verification with off-card code transformation

The Java Card bytecode verifier described in [31] attacks the memory problem from another angle. Like the standard bytecode verifiers, it solves dataflow equations using fixpoint iteration. To reduce memory requirements, however, it has only one global register type that is shared between all control points in the method. In other terms, the solution it infers is such that a given register has the same type throughout the method. For similar reasons, it also requires that the expression stack be empty at each branch instruction and at each branch target instruction.

With these extra restrictions, bytecode verification can be done in space $O(M_{stack} + M_{reg})$, instead of $O(N_{branch} \times (M_{stack} + M_{reg}))$ for Sun's algorithm, where N_{branch} is the number of branch targets. In practice, the memory requirements are small enough that all data structures fit comfortably in RAM on a smart card.

One drawback of this approach is that register initialization can no longer be checked statically, and must be replaced by run-time initialization of registers to a safe value (typically, the value of `null`) on method entry. Another drawback is that the extra restrictions imposed by the on-card verifier cause perfectly legal bytecode (that passes Sun's verifier) to be rejected.

To address the latter issue, this verifier relies on an off-card transformation, performed on the bytecode of the applet, that transforms any legal bytecode (that passes Sun's verifier) into equivalent bytecode that passes the on-card verifier. The off-card transformations include stack normalizations around branches and register reallocation by graph coloring, and are described in [31].

These transformations can increase the size of the code, as well as the number of registers used by a method. However, experience shows that these increases are minimal: on typical Java Card code, the increase in code size is less than 2%, and the increase in registers is negligible [31].

This off-card code transformation phase plays a role similar to that of adding certificates in Rose and Rose's approach: both are off-card processing that adds enough information to the applet to facilitate its on-card verification. However, the off-card code transformer embeds directly its information inside the code, via code transformations, instead of storing it into separate certificates. Moreover, the size of the extra information is much smaller (2% vs. 50%).

8 Conclusions and perspectives

Java bytecode verification is now a well-researched technique. The considerable body of formal work reviewed in this paper led not only to a precise understanding of what bytecode verification is and what it guarantees, but also to a number of new verification algorithms (besides Sun's original implementation) that cover a large part of the precision/cost spectrum.

A largely open question is whether bytecode verification can go beyond basic type safety and initialization properties, and statically establish more advanced properties of applets, such as resource usage (bounding the amount of memory allocated) and reactiveness (bounding the running time of an applet between two interactions with the outside world). Controlling resource usage is especially important for Java Card applets: since Java Card does not guarantee the presence of a garbage collector, applets are supposed to allocate all the objects they need at installation time, then run in constant space.

Other properties of interest include access control and information flow. Currently, the Java security manager performs all access control checks dynamically, using stack inspection [18]. Various

static analyses and program transformations have been proposed to perform some of these checks statically [25, 57, 42]. In the Java Card world, inter-applet communications via shared objects raises delicate security issues; [7] applies model-checking technology to this problem. As for information flow (an applet does not "leak" confidential information that it can access), this property is essentially impossible to check dynamically; several type systems have been proposed to enforce it statically [56, 55, 22, 1, 41].

Finally, the security of the sandbox model relies not only on bytecode verification, but also on the proper implementation of the API given to the applet. The majority of known applet-based attacks exploit bugs in the API in a type-safe way, rather than breaking type safety through bugs in the verifier. Verification of the API is a promising area of application for formal methods [32, 24].

Acknowledgements

A   Correctness of widened model checking

In this appendix, we formally prove that all bytecode accepted by a widened model checking algorithm (section 6.3) is also accepted by the BC model checking algorithm (section 6.2, and is therefore type-safe at runtime. In the following, W stands for an arbitrary widening function.

Lemma 1 Let $\sigma_1, \sigma_2, \sigma_1'$ be three states such that $\sigma_1 <: \sigma_2$ and $\sigma_1 \rightarrow \sigma_1'$. Further assume that $\sigma_2 \neq \mathtt{err}$ and $\sigma_2 \not\rightarrow \mathtt{err}$. Then, there exists σ_2' such that $\sigma_2 \rightarrow \sigma_2'$ and $\sigma_1' <: \sigma_2'$.

Proof: We first show that $\sigma_1' \neq \mathtt{err}$. Assume, by way of contradiction, that $\sigma_1 \rightarrow \mathtt{err}$. This means that $\sigma_1 = (pc, S_1, R_1)$ and the type-level abstract interpreter cannot make a transition from (S_1, R_1) on instruction $instr(pc)$. By hypotheses $\sigma_1 <: \sigma_2$ and $\sigma_2 \neq \mathtt{err}$, we have $\sigma_2 = (pc, S_2, R_2)$ with $S_1 <: S_2$ and $R_1 <: R_2$. The monotonicity property of the abstract interpreter guarantees that it cannot make a transition from (S_2, R_2) on $instr(pc)$. Thus, $\sigma_2 \rightarrow \mathtt{err}$, which is a contradiction.

Therefore, $\sigma_1' \neq \mathtt{err}$, and by definition of the $\rightarrow$ and $<:$ relations, this implies $\sigma_1 \neq \mathtt{err}$ and $\sigma_2 \neq \mathtt{err}$. We can therefore write

$$\sigma_1 = (p, S_1, R_1) \qquad \sigma_1' = (p', S_1', R_1') \qquad \sigma_2 = (p, S_2, R_2)$$

and the following properties hold:

$$S_1 <: S_2 \qquad R_1 <: R_2 \qquad instr(p) : (S_1, R_1) \rightarrow (S_1', R_1')$$

By hypothesis $\sigma_2 \not\rightarrow \mathtt{err}$, the abstract interpreter can make a transition on $instr(p)$ from state (S_2, R_2). Let (S_2', R_2') be the result of this transition. The monotonicity and determinacy properties of the type-level abstract interpreter guarantee that $S_1' <: S_2'$ and $R_1' <: R_2'$. Define $\sigma_2' = (p', S_2', R_2')$. We thus have $\sigma_1' <: \sigma_2'$.

To conclude the proof, it remains to show that $\sigma_2 \rightarrow \sigma_2'$. Since $instr(p) : (S_2, R_2) \rightarrow (S_2', R_2')$, it suffices to show that p' is a valid successor of the instruction at p in state (S_2, R_2). By hypothesis $\sigma_1 \rightarrow \sigma_1'$, we know that p' is a valid successor of the instruction at p in state (S_1, R_1). If the

instruction at p is not a `ret` instruction, its successors are independent of the stack and register types "before" the instruction; thus, $\sigma_2 \to \sigma_2'$. If the instruction at p is `ret` r, we have $R_1(r) = $ `retaddr`(p'). Since $\sigma_2 \not\to$ `err`, it must be the case that $R_2(r)$ is a `retaddr` type. Moreover, $R_1 <: R_2$. This entails $R_2(r) = R_1(r) = $ `retaddr`(p'). Hence, p' is the successor of the instruction at p in state (S_2, R_2). The expected result $\sigma_2 \to \sigma_2'$ follows. $\qquad\square$

In the following, we write $\Sigma \sqsubseteq \Sigma'$ to mean that for all $\sigma \in \Sigma$, there exists $\sigma' \in \Sigma'$ such that $\sigma <: \sigma'$.

Lemma 2 *Let Σ and Σ' be two sets of states such that $\Sigma \sqsubseteq \Sigma'$ and* `err` $\notin C(\Sigma')$*. Then, $C(\Sigma) \sqsubseteq W(C(\Sigma'))$.*

Proof: The $\sqsubseteq$ relation is transitive, and by hypothesis on the widening function W, we have $\Sigma \sqsubseteq W(\Sigma)$ for all Σ. It therefore suffices to show that $C(\Sigma) \sqsubseteq C(\Sigma')$.

Choose σ in $C(\Sigma)$. By definition of C, one of the following two cases holds:

- $\sigma \in \Sigma$. Since $\Sigma \sqsubseteq \Sigma'$, there exists $\sigma' \in \Sigma'$ such that $\sigma <: \sigma'$. Since $\Sigma' \subseteq C(\Sigma')$, it follows that σ is a subtype of an element of $C(\Sigma')$.

- $\rho \to \sigma$ for some $\rho \in \Sigma$. Let ρ' be an element of Σ' such that $\rho <: \rho'$. By hypothesis `err` $\notin C(\Sigma')$, we have $\rho' \neq$ `err` and $\rho' \not\to$ `err`. Applying lemma 1, we obtain σ' such that $\rho' \to \sigma'$ and $\sigma <: \sigma'$. Moreover, $\sigma' \in C(\Sigma')$ since σ' is obtained by taking one transition from an element of Σ'.

Since the result above holds for all $\sigma \in C(\Sigma)$, the expected result follows. $\qquad\square$

Theorem 1 *Let Σ_c be the smallest set closed under C containing σ_0, and Σ_w be the smallest set closed under $W \circ C$ containing σ_0. Then,* `err` $\notin \Sigma_w$ *implies* `err` $\notin \Sigma_c$*. In other terms, if the bytecode passes a widened model checking algorithm, it also passes the BC model-checking algorithm.*

Proof: Write $\Sigma_c = \bigcup_{n \in \mathbf{N}} \Sigma_c^n$ and $\Sigma_w = \bigcup_{n \in \mathbf{N}} \Sigma_w^n$, where $\Sigma_c^0 = \Sigma_w^0 = \{\sigma_0\}$ and $\Sigma_c^{n+1} = C(\Sigma_c^n)$ and $\Sigma_c^{n+1} = W(C(\Sigma_c^n))$.

We now show by induction on n that $\Sigma_c^n \sqsubseteq \Sigma_w^n$. The base case $n = 0$ is trivial. The inductive case follows from lemma 2, noticing that `err` $\notin \Sigma_w$ implies `err` $\notin C(\Sigma_w^n)$ for all n.

It follows that $\Sigma_c \sqsubseteq \Sigma_w$. Assume, by way of contradiction, that `err` $\in \Sigma_c$. Thus, there exists $\sigma \in \Sigma_w$ such that `err` $<: \sigma$. By definition of the $<:$ relation between states, this implies $\sigma =$ `err` and contradicts the hypothesis `err` $\notin \Sigma_w$. Hence `err` $\notin \Sigma_c$ as claimed. $\qquad\square$

B   Relative completeness of the CFHS algorithm

In this appendix, we prove that the CFHS verification algorithm (defined in section 6.3 as an instance of widened model checking) is complete with respect to the BC model-checking algorithm: all programs accepted by the latter algorithm are also accepted by the former. Combined with theorem 1, this result shows that the CFHS algorithm accepts exactly the same programs as the BC algorithm.

We recall the definitions that describe the CFHS algorithm: the equivalence relation $\approx$ is defined by `err` $\approx$ `err` and $(p_1, S_1, R_1) \approx (p_2, S_2, R_2)$ iff $p_1 = p_2$ and S_1, S_2 agree on return addresses

and R_1, R_2 agree on return addresses. The merging function M is defined by $M(\{\texttt{err}\}) = \texttt{err}$ and $M(\{(p, S_1, R_1), \ldots, (p, S_n, R_n)\}) = (p, lub(S_1, \ldots, S_n), lub(R_1, \ldots, R_n))$. Finally, the widening function W is defined by $W(\Sigma) = \{M(\Sigma_1), \ldots, M(\Sigma_n)\}$ where $\Sigma_1, \ldots, \Sigma_n$ is a partition of Σ into equivalence classes for $\approx$.

The following key lemma shows that the type-level abstract machine performs parallel transitions from a set of equivalent states and from the single state obtained by merging the equivalent states. In particular, if $\texttt{err}$ is not reachable from the equivalent states, it is not reachable from the merged state either. We write $T(\Sigma)$ for the states reachable by one transition from a state in Σ, that is, $T(\Sigma) = \{\sigma' \mid \exists \sigma \in \Sigma. \ \sigma \to \sigma'\}$.

Lemma 3 *Let Σ be a non-empty set of equivalent states such that $\texttt{err} \notin \Sigma$ and $\texttt{err} \notin T(\Sigma)$. We then have $W(T(\Sigma)) = T(W(\Sigma))$.*

Proof: Write $\Sigma = \{(p, S_1, R_1), \ldots, (p, S_n, R_n)\}$ where the S_i agree on return addresses, and the R_i also agree on return addresses. Since $\texttt{err} \notin T(\Sigma)$, for every i, there exists S_i' and R_i' such that $instr(p) : (S_i, R_i) \to (S_i', R_i')$. According to the continuity property 4, we thus have $instr(p) : (lub(S_1 \ldots S_n), lub(R_1 \ldots R_n)) \to (lub(S_1' \ldots S_n'), lub(R_1' \ldots R_n'))$. Moreover, by the determinacy property 2, there are no other transitions from these initial states.

Consider the successors of the instruction at p. If $\texttt{instr}(p)$ is not a $\texttt{ret}$ instruction, these successors $p_1, \ldots, p_k$ depend only on $\texttt{instr}(p)$ and are independent of the stack and register types. If $\texttt{instr}(p)$ is a $\texttt{ret} \ r$ instruction, the successor for the instruction is determined by the $\texttt{retaddr}$ type found in $R(r)$, where R is the register type before the instruction. However, the register types "before" considered in this proof, namely $R_1, \ldots, R_n$ and $lub(R_1, \ldots, R_n)$, all agree on return addresses. This means that they all assign the same $\texttt{retaddr}$ type to register r. Thus, the successor for the $\texttt{ret} \ r$ instruction at p is the same in all these register types, just as in the case of a non-$\texttt{ret}$ instruction at p.

It follows from this discussion of successors that the set $T(\Sigma)$ of states reachable by one transition from Σ is exactly the cartesian product of the set of possible successors $p_1, \ldots, p_k$ with the set of stack and register types "after" $(S_1', R_1'), \ldots, (S_n', R_n')$, and similarly for $T(W(\Sigma))$:

$$
\begin{aligned}
T(\Sigma) &= \{(p_i, S_j', R_j') \mid i \in \{1, \ldots, k\}, \ j \in \{1, \ldots, n\}\} \\
T(W(\Sigma)) &= \{(p_i, lub(S_1' \ldots S_n'), lub(R_1' \ldots R_n')) \mid i \in \{1, \ldots, k\}\}
\end{aligned}
$$

Finally, notice that the stack and register types "after" $S_1', \ldots, S_n'$ and $R_1', \ldots, R_n'$ agree on return addresses: if $instr(p)$ is not a $\texttt{jsr}$ instruction, any $\texttt{retaddr}$ type in the state after comes from a matching $\texttt{retaddr}$ type in the state before; if $instr(p)$ is a $\texttt{jsr}$ instruction, the state after contains an additional $\texttt{retaddr}(p + 3)$ type at the top of the stack, but this type is the same in all stack types after. Thus,

$$
W(T(\Sigma)) = \{(p_i, lub(S_1' \ldots S_n'), lub(R_1' \ldots R_n')) \mid i \in \{1, \ldots, k\}\}.
$$

It follows that $W(T(\Sigma)) = T(W(\Sigma))$ as expected. □

We now prove a number of algebraic properties of the M and W functions.

Lemma 4 *Let Σ be a non-empty set of equivalent states. Then, $M(\Sigma)$ is equivalent to every $\sigma \in \Sigma$.*

Proof: We first show that if two types τ_1 and τ_2 agree on return addresses, then $lub(\tau_1, \tau_2)$ agrees on return addresses with both τ_1 and τ_2. Indeed, either $\tau_1 = \tau_2$, in which case $lub(\tau_1, \tau_2) = \tau_1 = \tau_2$ and the result is obvious; or neither τ_1 nor τ_2 are `retaddr` types, in which case $lub(\tau_1, \tau_2)$ is not a `retaddr` type either, thus agrees with τ_1 and with τ_2.

The previous result extends trivially to the least upper bound of one or several stack types, or one or several register types.

Let Σ be a non-empty set of equivalent states. Either $\Sigma = \{\texttt{err}\}$, or $\Sigma = \{(p, S_1, R_1), \ldots, (p, S_n, R_n)\}$. In the former case, $M(\Sigma) = \texttt{err}$ and the expected result obviously holds. In the latter case, $M(\Sigma) = (p, lub(S_1, \ldots, S_n), lub(R_1, \ldots, R_n))$. For every i, $lub(S_1, \ldots, S_n)$ and S_i agree on return addresses, and $lub(R_1, \ldots, R_n)$ and R_i agree on return addresses. Thus, $M(\Sigma) \approx \sigma$ for any $\sigma \in \Sigma$. $\qquad\square$

Lemma 5 *Let $\Sigma_1, \ldots, \Sigma_n$ be non-empty sets of states such that all states in $\Sigma_1 \cup \ldots \cup \Sigma_n$ are equivalent. Then, $M(\Sigma_1 \cup \ldots \cup \Sigma_n) = M(\{M(\Sigma_1), \ldots, M(\Sigma_n)\})$.*

Proof: There are two cases to consider. In the first case, the Σ_i are all $\{\texttt{err}\}$. In this case, the expected equality trivially holds. In the second case, `err` does not belong to any Σ_i. Write

$$\Sigma_i = \{(p, S_{i,1}, R_{i,1}), \ldots, (p, S_{i,n_i}, R_{i,n_i})\}$$

We thus have

$$
\begin{aligned}
M(\Sigma_i) &= (p, lub_j(S_{i,j}), lub_j(R_{i,j})) \\
M(\Sigma_1 \cup \ldots \cup \Sigma_n) &= (p, lub_{i,j}(S_{i,j}), lub_{i,j}(R_{i,j}))
\end{aligned}
$$

By hypothesis, all elements of Σ_i are equivalent to all elements of Σ_j. By lemma 4, $M(\Sigma_i)$ is equivalent to all elements of Σ_i, and $M(\Sigma_j)$ is equivalent to all elements of Σ_j. By transitivity, it follows that $M(\Sigma_i) \approx M(\Sigma_j)$. This holds for all i, j. We thus have

$$M(\{M(\Sigma_1), \ldots, M(\Sigma_n)\}) = (p, lub_i(lub_j(S_{i,j})), lub_i(lub_j(R_{i,j}))).$$

The expected result follows from the associativity and commutativity of the *lub* operation. $\qquad\square$

Lemma 6 $W(\Sigma_1 \cup \ldots \cup \Sigma_n) = W(W(\Sigma_1) \cup \ldots \cup W(\Sigma_n))$ *for all sets of states $\Sigma_1, \ldots, \Sigma_n$.*

Proof: Let $A_1, \ldots, A_k$ be a partition of $\Sigma_1 \cup \ldots \cup \Sigma_n$ into equivalence classes. Thus, for every i, $\{A_j \cap \Sigma_i \mid j \in \{1 \ldots k\}, A_j \cap \Sigma_i \neq \emptyset\}$ is a partition of Σ_i into equivalence classes. We thus have

$$
\begin{aligned}
&W(\Sigma_1) \cup \ldots \cup W(\Sigma_n) \\
&= \{M(A_j \cap \Sigma_i) \mid j \in \{1 \ldots k\}, i \in \{1 \ldots n\}, A_j \cap \Sigma_i \neq \emptyset\}
\end{aligned}
$$

Consider two elements of this set, $M(A_j \cap \Sigma_i)$ and $M(A_{j'} \cap \Sigma_{i'})$. By lemma 4, the former is equivalent to every element of $A_j \cap \Sigma_i$, and the latter to every element of $A_{j'} \cap \Sigma_{i'}$. Moreover, the $A_1, \ldots, A_k$ are disjoint equivalence classes. Hence, $M(A_j \cap \Sigma_i) \approx M(A_{j'} \cap \Sigma_{i'})$ if and only if $j = j'$. Therefore, the equivalence classes of $W(\Sigma_1) \cup \ldots \cup W(\Sigma_n)$ for the $\approx$ relation are the sets

$$B_j = \{M(A_j \cap \Sigma_i) \mid i \in \{1 \ldots n\}, A_j \cap \Sigma_i \neq \emptyset\}$$

for $j = 1, \ldots, k$. By lemma 5,

$$M(B_j) = M(\cup\{A_j \cap \Sigma_i \mid i \in \{1 \ldots n\}, A_j \cap \Sigma_i \neq \emptyset\}) = M(A_j)$$

It follows that

$$
\begin{aligned}
W(W(\Sigma_1) \cup \ldots \cup W(\Sigma_n)) &= \{M(B_1), \ldots, M(B_k)\} \\
&= \{M(A_1), \ldots, M(A_k)\} \\
&= W(\Sigma_1 \cup \ldots \cup \Sigma_n)
\end{aligned}
$$

This is the expected result. □

Lemma 7 *Let Σ be a set of states such that $\mathtt{err} \notin C(\Sigma)$. Then, $W(C(\Sigma)) = W(C(W(\Sigma)))$.*

Proof: Let $\Sigma_1, \ldots, \Sigma_n$ be a partition of Σ into equivalence classes. By definition of C, we have

$$C(\Sigma) = \Sigma_1 \cup \ldots \cup \Sigma_n \cup T(\Sigma_1) \cup \ldots \cup T(\Sigma_n).$$

Applying lemma 6, we obtain

$$W(C(\Sigma)) = W(W(\Sigma_1) \cup \ldots \cup W(\Sigma_n) \cup W(T(\Sigma_1)) \cup \ldots \cup W(T(\Sigma_n))).$$

It follows from the hypothesis $\mathtt{err} \notin C(\Sigma)$ that $\mathtt{err} \notin \Sigma_i$ and $\mathtt{err} \notin T(\Sigma_i)$ for all $i = 1, \ldots, n$. By lemma 3, we thus have $W(T(\Sigma_i)) = T(W(\Sigma_i))$. It follows that

$$
\begin{aligned}
W(C(\Sigma)) &= W(W(\Sigma_1) \cup \ldots \cup W(\Sigma_n) \cup T(W(\Sigma_1)) \cup \ldots \cup T(W(\Sigma_n))) \\
&= W(C(W(\Sigma_1) \cup \ldots W(\Sigma_n)))
\end{aligned}
$$

Since $\Sigma_1, \ldots, \Sigma_n$ is a partition of Σ into equivalence classes, we have $W(\Sigma_i) = \{M(\Sigma_i)\}$ and $W(\Sigma) = \{M(\Sigma_1), \ldots, M(\Sigma_n)\} = W(\Sigma_1) \cup \ldots \cup W(\Sigma_n)$. The expected result follows. □

Theorem 2 *Let Σ_c be the smallest set closed under C containing σ_0, and Σ_w be the smallest set closed under $W \circ C$ containing σ_0. Then, $\mathtt{err} \notin \Sigma_c$ implies $\mathtt{err} \notin \Sigma_w$. In other terms, if the bytecode passes the BC algorithm, it also passes the CFHS algorithm.*

Proof: Write $\Sigma_c = \bigcup_{n \in \mathbf{N}} \Sigma_c^n$ and $\Sigma_w = \bigcup_{n \in \mathbf{N}} \Sigma_w^n$, where $\Sigma_c^0 = \Sigma_w^0 = \{\sigma_0\}$ and $\Sigma_c^{n+1} = C(\Sigma_c^n)$ and $\Sigma_c^{n+1} = W(C(\Sigma_c^n))$. Assume that $\mathtt{err} \notin \Sigma_c$. This implies $\mathtt{err} \notin \Sigma_c^n$ for any n.

We now show by induction on n that $\Sigma_w^n = W(\Sigma_c^n)$. The base case $n = 0$ is trivial, since W is the identity function on singleton sets. For the inductive case, assume $\Sigma_w^n = W(\Sigma_c^n)$. Then, $\Sigma_w^{n+1} = W(C(\Sigma_w^n)) = W(C(W(\Sigma_c^n)))$. Note that $\mathtt{err} \notin C(\Sigma_c^n) = \Sigma_c^{n+1}$. Applying lemma 7, it follows that $\Sigma_w^{n+1} = W(C(\Sigma_c^n)) = W(\Sigma_c^{n+1})$, as desired.

By construction of the widening function W, for any set of states Σ we have $\mathtt{err} \in W(\Sigma)$ if and only if $\mathtt{err} \in \Sigma$. Thus, for any n, we have $\mathtt{err} \notin \Sigma_w^n$ since $\Sigma_w^n = W(\Sigma_c^n)$ and $\mathtt{err} \notin \Sigma_c^n$. It follows that $\mathtt{err} \notin \Sigma_w$. □

References

[1] Martín Abadi, Anindya Banerjee, Nevin Heintze, and Jon G. Riecke. A core calculus of dependency. In *26th symposium Principles of Programming Languages*, pages 147–160. ACM Press, 1999.

[2] David Basin, Stefan Friedrich, and Marek Gawkowski. Bytecode verification by model checking. *Journal of Automated Reasoning*. Special issue on bytecode verification (this issue).

[3] Yves Bertot. Formalizing a JVML verifier for initialization in a theorem prover. In *Proc. Computer Aided Verification (CAV'01)*, volume 2102 of *Lecture Notes in Computer Science*, pages 14–24. Springer-Verlag, 2001.

[4] Pascal Brisset. Vers un vérifieur de bytecode Java certifié. Seminar given at Ecole Normale Supérieure, Paris, October 2nd 1998, 1998.

[5] Klaus Brunnstein. Hostile ActiveX control demonstrated. *RISKS Forum*, 18(82), February 1997.

[6] Zhiqun Chen. *Java Card Technology for Smart Cards: Architecture and Programmer's Guide*. The Java Series. Addison-Wesley, 2000.

[7] Gennady Chugunov, Lars Åke Fredlund, and Dilian Gurov. Model checking multi-applet Java Card applications. In *Smart Card Research and Advanced Applications Conference (CARDIS'02)*, 2002.

[8] Alessandro Coglio. Simple verification technique for complex Java bytecode subroutines. In *4th ECOOP Workshop on Formal Techniques for Java-like Programs*, 2002. Extended version available as Kestrel Institute technical report.

[9] Alessandro Coglio. Improving the official specification of Java bytecode verification. *Concurrency and Computation: Practice and Experience*, 15(2):155–179, 2003.

[10] Alessandro Coglio, Allen Goldberg, and Zhenyu Qian. Towards a provably-correct implementation of the JVM bytecode verifier. In *OOPSLA Workshop on Formal Underpinnings of Java*, 1998.

[11] Richard Cohen. The defensive Java virtual machine specification. Technical report, Computational Logic Inc., 1997.

[12] Damien Deville and Gilles Grimaud. Building an "impossible" verifier on a Java Card. In *USENIX Workshop on Industrial Experiences with Systems Software (WIESS'02)*, 2002.

[13] Stephen N. Freund and John C. Mitchell. A type system for the Java bytecode language and verifier. *Journal of Automated Reasoning*. Special issue on bytecode verification (this issue).

[14] Stephen N. Freund and John C. Mitchell. A formal framework for the java bytecode language and verifier. In *Object-Oriented Programming Systems, Languages and Applications 1999*, pages 147–166. ACM Press, 1999.

[15] Stephen N. Freund and John C. Mitchell. A type system for object initialization in the Java bytecode language. *ACM Transactions on Programming Languages and Systems*, 21(6):1196–1250, 1999.

[16] Alexandre Frey. On-terminal verifier for JEFF files. Personal communication, 2001.

[17] Allen Goldberg. A specification of Java loading and bytecode verification. In *ACM Conference on Computer and Communications Security*, pages 49–58. ACM Press, 1998.

[18] Li Gong. *Inside Java 2 platform security: architecture, API design, and implementation*. The Java Series. Addison-Wesley, 1999.

[19] James A. Gosling. Java intermediate bytecodes. In *Proc. ACM SIGPLAN Workshop on Intermediate Representations*, pages 111–118. ACM, 1995.

[20] Masami Hagiya and Akihiko Tozawa. On a new method for dataflow analysis of Java virtual machine subroutines. In Giorgio Levi, editor, *SAS'98*, volume 1503 of *Lecture Notes in Computer Science*, pages 17–32. Springer-Verlag, 1998.

[21] Pieter H. Hartel and Luc A. V. Moreau. Formalizing the safety of Java, the Java virtual machine and Java Card. *ACM Computing Surveys*, 33(4):517–558, 2001.

[22] Nevin Heintze and Jon G. Riecke. The SLam calculus: programming with secrecy and integrity. In *25th symposium Principles of Programming Languages*, pages 365–377. ACM Press, 1998.

[23] Ludovic Henrio and Bernard Serpette. A framework for bytecode verifiers: Application to intra-procedural continuations. Research report, INRIA, 2001.

[24] Marieke Huisman, Bart Jacobs, and Joachim van den Berg. A case study in class library verification: Java's Vector class. *Software Tools for Technology Transfer*, 3(3):332–352, 2001.

[25] Thomas Jensen, Daniel Le Métayer, and Tommy Thorn. Verification of control flow based security properties. In *IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 1999.

[26] G. Klein and M. Wildmoser. Verified bytecode subroutines. *Journal of Automated Reasoning*. Special issue on bytecode verification (this issue).

[27] Gerwin Klein. *Verified Java bytecode verification*. PhD thesis, Technische Universität München, 2003.

[28] Gerwin Klein and Tobias Nipkow. Verified lightweight bytecode verification. *Concurrency and Computation: Practice and Experience*, 13:1133–1151, 2001.

[29] Gerwin Klein and Tobias Nipkow. Verified bytecode verifiers. *Theoretical Computer Science*, 2002. To appear.

[30] Todd Knoblock and Jakob Rehof. Type elaboration and subtype completion for Java bytecode. In *27th symposium Principles of Programming Languages*, pages 228–242. ACM Press, 2000.

[31] Xavier Leroy. Bytecode verification for Java smart card. *Software Practice & Experience*, 32:319–340, 2002.

[32] Xavier Leroy and François Rouaix. Security properties of typed applets. In J. Vitek and C. Jensen, editors, *Secure Internet Programming – Security issues for Mobile and Distributed Objects*, volume 1603 of *Lecture Notes in Computer Science*, pages 147–182. Springer-Verlag, 1999.

[33] Tim Lindholm and Frank Yellin. *The Java Virtual Machine Specification*. The Java Series. Addison-Wesley, 1999. Second edition.

[34] Gary McGraw and Ed Felten. *Securing Java*. John Wiley & Sons, 1999.

[35] Steven S. Muchnick. *Advanced compiler design and implementation*. Morgan Kaufmann, 1997.

[36] George C. Necula. Proof-carrying code. In *24th symposium Principles of Programming Languages*, pages 106–119. ACM Press, 1997.

[37] Flemming Nielson, Hanne Riis Nielson, and Chris Hankin. *Principles of program analysis*. Springer-Verlag, 1999.

[38] Tobias Nipkow. Verified bytecode verifiers. In *Foundations of Software Science and Computation Structures (FOSSACS'01)*, volume 2030 of *Lecture Notes in Computer Science*, pages 347–363. Springer-Verlag, 2001.

[39] Robert O'Callahan. A simple, comprehensive type system for Java bytecode subroutines. In *26th symposium Principles of Programming Languages*, pages 70–78. ACM Press, 1999.

[40] Joachim Posegga and Harald Vogt. Java bytecode verification using model checking. In *Workshop Fundamental Underpinnings of Java*, 1998.

[41] François Pottier and Vincent Simonet. Information flow inference for ML. In *29th symposium Principles of Programming Languages*, pages 319–330. ACM Press, 2002.

[42] François Pottier, Christian Skalka, and Scott Smith. A systematic approach to static access control. In David Sands, editor, *Proceedings of the 10th European Symposium on Programming (ESOP'01)*, volume 2028 of *Lecture Notes in Computer Science*, pages 30–45. Springer-Verlag, 2001.

[43] Cornelia Pusch. Proving the soundness of a Java bytecode verifier specification in Isabelle/HOL. In W. Rance Cleaveland, editor, *TACAS'99*, volume 1579 of *Lecture Notes in Computer Science*, pages 89–103. Springer-Verlag, 1999.

[44] Zhenyu Qian. A formal specification of Java virtual machine instructions for objects, methods and subroutines. In Jim Alves-Foss, editor, *Formal syntax and semantics of Java*, volume 1523 of *Lecture Notes in Computer Science*. Springer-Verlag, 1998.

[45] Zhenyu Qian. Standard fixpoint iteration for Java bytecode verification. *ACM Transactions on Programming Languages and Systems*, 22(4):638–672, 2000.

[46] Eva Rose. *Vérification de code d'octet de la machine virtuelle Java: formalisation et implantation*. PhD thesis, University Paris 7, September 2002.

[47] Eva Rose and Kristoffer Rose. Lightweight bytecode verification. In *OOPSLA Workshop on Formal Underpinnings of Java*, 1998.

[48] David A. Schmidt. Data flow analysis is model checking of abstract interpretations. In *25th symposium Principles of Programming Languages*, pages 38–48. ACM Press, 1998.

[49] Robert Stärk, Joachim Schmid, and Egon Börger. *Java and the Java Virtual Machine*. Springer-Verlag, 2001.

[50] Robert F. Stärk and Joachim Schmid. Completeness of a bytecode verifier and a certifying Java-to-JVM compiler. *Journal of Automated Reasoning*. Special issue on bytecode verification (this issue).

[51] Raymie Stata and Martín Abadi. A type system for Java bytecode subroutines. *ACM Transactions on Programming Languages and Systems*, 21(1):90–137, 1999.

[52] Sun Microsystems. Java 2 platform micro edition technology for creating mobile devices. White paper, http://java.sun.com/products/cldc/wp/KVMwp.pdf, 2000.

[53] Trusted Logic. Off-card bytecode verifier for Java Card. Distributed as part of Sun's Java Card Development Kit, 2001.

[54] G. Vigna, editor. *Mobile Agents and Security*, volume 1419 of *Lecture Notes in Computer Science*. Springer-Verlag, 1998.

[55] Dennis Volpano and Geoffrey Smith. A type-based approach to program security. In *Proceedings of TAPSOFT'97, Colloquium on Formal Approaches in Software Engineering*, volume 1214 of *Lecture Notes in Computer Science*, pages 607–621. Springer-Verlag, 1997.

[56] Dennis Volpano, Geoffrey Smith, and Cynthia Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3):1–21, 1996.

[57] David Walker. A type system for expressive security policies. In *27th symposium Principles of Programming Languages*, pages 254–267. ACM Press, 2000.

[58] Frank Yellin. Low level security in Java. In *Proceedings of the Fourth International World Wide Web Conference*, pages 369–379. O'Reilly, 1995.