

# Static Enforcement of Information Flow Policies for a Concurrent JVM-like Language\*

Gilles Barthe and Exequiel Rivas

IMDEA Software Institute, Madrid, Spain

**Abstract.** An essential security goal of mobile code platforms is to protect confidential data against untrusted third-party applications; yet, prevailing mechanisms for ensuring confidentiality of mobile code are limited to sequential programs, whereas existing applications are generally concurrent. To bridge this gap, we develop a sound information-flow type system for a JVM-like, low-level concurrent object-oriented language. The type system builds upon existing solutions for object-oriented languages and concurrency, solving a number of intricate issues in their combination. Moreover, we connect the type system for bytecode programs to a type system for Java programs, extending the results of type-preserving compilation developed in earlier works.

## 1 Introduction

Non-interference is a baseline information flow policy which provides strong end-to-end security guarantees about programs, ensuring that they do not reveal confidential data throughout execution. Because information flow policies focus on program behavior rather than program origin, they are suitable for mobile code security, where confidential data must be protected against applications originating from untrusted third parties. Hence, there have been efforts to enhance the Java bytecode verifier, one of the two main security mechanisms of the Java Virtual Machine, so that it provides a type-based enforcement mechanism for non-interference. For instance, Barthe, Pichardie and Rezk [4] give a sound information flow type system for a large fragment of the Java Virtual Machine, including objects and arrays, methods, and exceptions. One important limitation of [4] is that it is restricted to sequential programs. Indeed, mobile code applications are generally concurrent, and hence cannot be handled by the type system.

Developing flexible information flow type systems for concurrent languages is notoriously difficult. Indeed, the concurrent execution of two secure programs may be insecure [14], and hence information flow type systems impose stringent restrictions on programs. In [12,13], Russo and Sabelfeld introduce the notion of secure scheduler to provide a more flexible (and yet sound) approach to handle concurrency. Building on their idea, Barthe, Rezk, Russo and Sabelfeld [5] give a modular method for extending an information flow type system for a sequential low-level language into an information flow type system for a concurrent low-level language, and instantiate their method to

---

\* Partially funded by European Projects FP7-231620 HATS and FP7-256980 NESSoS, Spanish project TIN2009-14599 DESAFIOS 10, Madrid Regional project S2009TIC-1465 PROMETIDOS.

a minimal stack-based language that can serve as a target for compiling WHILE programs; the crux of the method is to make schedulers aware of the security environment that maps program points to a security level that lowers bound their effect, and to require that schedulers do not leak information through internal timing. Although the method is modular, its application to the type system of [4] raises a significant difficulty, because the soundness of the latter is based on a mix-step semantics, in which method calls are performed in one step. While the mix-step semantics makes soundness proofs considerably simpler, it is inappropriate for reasoning about concurrent programs, since the execution of separate threads might be interleaved. In this paper, we address this problem by enhancing the proof method of [4] so that all reasoning about program executions is performed directly with respect to a small-step semantics. This requires substantial generalizations in the notions of state equivalence, and substantial adaptations in the unwinding lemmas. Moreover, we provide a more realistic treatment of thread creation, so that it matches more closely the operational semantics of the JVM, and recast one technical hypothesis of [5] in terms of control dependence regions, making its verification simpler.

Moreover, we consider the issue of type-preserving compilation: we define an information flow type system for a concurrent fragment of Java and show that non-optimizing compilation preserves typability of programs, thereby generalizing [3,5]. Besides establishing that developers and consumers views on secure information flow (respectively embodied in the Java and JVM type systems) coincide, our result allows us to derive that the source type system is sound, i.e. enforces non-interference. To our best knowledge, this is the first proof of non-interference for a concurrent fragment of Java.

In summary, the main contributions of this article are:

- we provide the first sound information-flow type system for low-level, concurrent, object-oriented language;
- we define an information flow type system for a concurrent fragment of Java, prove type-preserving compilation, and derive soundness of the type system;
- we simplify one technical hypothesis in [5] that was extremely difficult to establish, even for a simple language.

## 2 A Concurrent JVM Language

We consider a concurrent JVM-like language, but omit some features most notably locks and exceptions.

*Programs.* Figure 2 gives the definition of programs. Informally, a program  $P$  is given by: i) a set  $\mathcal{C}$  of classes, including a **main** class; ii) a set  $\mathcal{F}$  of fields, and a function  $\text{fieldsof} : \mathcal{C} \rightarrow \mathcal{P}(\mathcal{F})$  that yields a partition of  $\mathcal{F}$ ; iii) a set  $\mathcal{M}_{id}$  of method identifiers; iv) a set  $\mathcal{M}_{vt}$  of virtual method identifiers, including an identifier `run` reserved for methods that run as new threads; v) a set  $\mathcal{M}$  of methods, including a method identifier, a virtual method identifier, an argument table that maps argument positions to local variable names, and a sequence of instructions from Figure 1; vi) a function `lookup` attached to each program that takes a method identifier and a class name and returns the

$$\text{Instr} ::= \text{binop } op \mid \text{push } n \mid \text{load } x \mid \text{store } x \mid \text{ifeq } j \mid \\ \text{goto } j \mid \text{invokevirtual } m_{ID} \mid \text{return} \mid \text{start}$$

where  $op$  is a binary operation,  $n$  is an integer,  $i$  and  $j$  are natural numbers, and  $x$  is an identifier.

**Fig. 1.** Instruction set

$\mathcal{C} \ni C$	(classes)
$\mathcal{F} \ni f$	(fields)
$\mathcal{C} \rightarrow \mathcal{O}(\mathcal{F}) \ni \text{fieldsof}$	(class description)
$\mathcal{M}_{id} \ni m_{id}$	(method ids.)
$\mathcal{M}_{vt} \ni m_{vt}$	(virt. method ids.)
$\text{Instr} \ni ins$	(instruction)
$\text{List} \ni lis ::= ins :: lis \mid ins$	(instruction listing)
$\mathbb{N} \rightarrow \mathcal{I} \ni \text{arg}$	(argument table)
$\mathcal{M} \ni m ::= \langle m_{vt}, m_{id}, \text{arg}, lis \rangle$	(methods)
$\mathcal{M}_{id} \times \mathcal{C} \rightarrow \mathcal{M} \ni \text{lookup}$	(class lookup)
$\mathcal{P} \ni i ::= \langle n, m_{id} \rangle$	(program points)
$\text{Program} \ni P ::= \langle \mathcal{C}, \mathcal{F}, \text{fieldsof}, \mathcal{M}, \mathcal{M}_{id}, \mathcal{M}_{vt}, \text{lookup} \rangle$	(programs)

**Fig. 2.** Programs

method to be executed. A program point consists of a method  $m$ , and an integer  $i$  that is smaller than the length of the list of instructions attached to  $m$ , and only  $i$  is used if the method is understood from context. We let  $P_m[i]$  denote the instruction at index  $i$  of method  $m$  and let  $\mathcal{P}$  denote the set of program points. We write  $\text{arg}_m$  for the argument table corresponding to method  $m$ , and assume a set  $\mathcal{I}$  of variable identifiers.

*States.* Figure 2 introduces the semantic domains used to interpret the program behavior. Informally, a state is a set of threads indexed by thread identifiers together with a heap. The local state of a thread is given by a stack of frames, where each frame contains some local state used to execute a method invocation. The heap is modeled as partial maps from locations to objects; objects are themselves modeled as pairs, consisting of a class and of a partial map from fields to values. We assume heaps satisfy standard well-formedness constraints, e.g the domain of  $p$  include the fields of  $C$ , and  $o$  maps fields to values of the correct type. Values can be either integer numbers or locations, i.e.  $\mathcal{V} = \mathbb{Z} \cup \mathcal{L} \cup \{\text{null}\}$ , where  $\mathcal{L}$  is an (infinite) set of locations used to store the address of an object in the heap and  $\text{null}$  denotes the null pointer. Following [4], we avoid pointer arithmetic by making a distinction between locations and integer values. A set of accessors that we use to extract information from frames and states is defined.

*Thread-local semantics.* The operational semantics of JVM programs is defined in two steps: we first define a thread-local semantics that characterize the effect of executing

	$\mathcal{L} \ni \ell$	(locations)
	$\mathcal{V} \ni v ::= \mathbb{Z} \cup \mathcal{L} \cup \{\text{null}\}$	(values)
	$\mathcal{V}^* \ni s ::= v :: s \mid \epsilon$	(operand stacks)
$\mathcal{F} \rightarrow$	$\mathcal{V} \ni p$	(properties)
	$\mathcal{O} \ni o ::= \langle C, p \rangle$	(objects)
$\mathbb{N} \rightarrow$	$\mathcal{V} \ni \rho$	(environments)
	Frame $\ni f ::= \langle m_{id}, i, \rho, s \rangle$	(frames)
	Thread $\ni t ::= f :: t \mid f$	(threads)
$\mathcal{T} \rightarrow$	Thread $\ni ts$	(thread set)
Mem = $\mathcal{L} \rightarrow$	$\mathcal{O} \ni \mu$	(heaps)
	State $\ni s ::= \langle ts, \mu \rangle$	(states)

  

$f.m \Rightarrow m_{id}$	$f.s \Rightarrow s$	$s.ts \Rightarrow ts$
$f.pc \Rightarrow i$	$t.s \Rightarrow \text{hd}(t).s$	$s.act \Rightarrow \text{dom}(ts)$
$f.env \Rightarrow \rho$	$t.pc \Rightarrow \text{hd}(t).pc$	$s.mem \Rightarrow \mu$

**Fig. 3.** Semantic domains and accessors

an instruction on a thread. We omit its definition, which is totally standard, and can be found in the full version of the paper. The thread-local semantics is defined as a relation  $\rightsquigarrow_s \subseteq \text{Thread} \times \text{Mem} \times \text{Thread}_\perp \times \text{Mem}$ : we write  $\langle t, \mu \rangle \rightsquigarrow_s \langle t', \mu' \rangle$  instead of  $\langle t, \mu, t', \mu' \rangle \in \rightsquigarrow_s$  as is standard. Note that termination of a thread corresponds to  $t' = \perp$ . The semantics is parametrized by a function  $\text{default} : \mathcal{C} \rightarrow \mathcal{O}$  that returns a default object of the corresponding class and it is used to create the new objects.

*Global semantics.* The semantics of programs is parametrized by a scheduler, that picks a thread to be executed, according to the current state, and a history  $h$ . A history is defined as a sequence of thread identifiers; we let  $\mathcal{H}$  denote the set of histories, and let  $\epsilon$  denote the initial history. We model a scheduler as a function  $\text{pick}t : \text{State} \times \mathcal{H} \rightarrow \mathcal{T}$  that picks the next instruction to be executed. Moreover, we assume given a function  $\text{updh} : \text{State} \times \mathcal{H} \rightarrow \mathcal{H}$  that updates the history after program execution. Formally, one-step execution is defined as a relation between pairs of states and histories. We write  $\langle s, h \rangle \rightsquigarrow \langle s', h' \rangle$  when executing  $s$  with history  $h$  leads to state  $s'$  with history  $h'$ ; the rules for the semantics are given in Figure 4. The semantics considers three cases: i) the instruction to be executed spawns a new thread; ii) it terminates the execution of a thread (in case the instruction is a return instruction and the thread has a single frame); iii) otherwise, the thread-local effect of the instruction is propagated to the active thread. We assume given a function  $\text{fresh}t$  that takes as input a set of thread identifiers and generates a new thread identifier. Finally, we define the evaluation semantics.

**Definition 1 (Evaluation semantics).** *The evaluation relation  $\Downarrow \subseteq (\text{State} \times \mathcal{H}) \times \text{Mem}$  is defined as:  $s, h \Downarrow \mu'$  iff there is a pair  $s', h'$  such that  $s, h \rightsquigarrow^* s', h'$  and  $s'.act = \emptyset$  and  $s'.mem = \mu$ . We write  $P, \mu \Downarrow \mu'$  as a shorthand for  $s_{init}(\mu), \epsilon \Downarrow \mu'$ , where  $s_{init}(\mu)$  stands for  $\langle \{\text{fresh}t(\emptyset)\} \mapsto \langle 1_{\text{main.run}}, \rho, \epsilon \rangle, \mu \rangle$  where  $\rho$  is an empty environment.*

$$\begin{array}{c}
 \frac{\text{pick}t(s, h) = ctid \quad \text{updh}(s, h) = h' \quad s_{ctid}.pc = i \quad P_m[i] = \text{start} \quad s_{ctid}.os = o :: os \\
 \text{fresh}t(s) = ntid \quad r = \text{lookup}(\text{run}, \text{class}(s.\text{mem}(o))) \quad s_{ctid}.pc := i + 1 = \sigma'}{s, h \rightsquigarrow s.[\text{lst}(ctid).\text{lst} := \sigma', \text{lst}(ntid) := \langle \langle 1, \{this \mapsto o\}, \epsilon \rangle \rangle], h'} \text{(Spawn)} \\
 \text{pick}t(s, h) = ctid \quad \text{updh}(s, h) = h' \quad s_{ctid}.pc = i \quad P_m[i] = \text{return} \quad \langle s_{ctid}.ts, s.\text{mem} \rangle \rightsquigarrow_s \langle \perp, \mu \rangle \\
 s, h \rightsquigarrow s.[\text{lst} := \text{lst} \setminus ctid, \text{mem} := \mu], h'} \text{(Kill)} \\
 \text{pick}t(s, h) = ctid \quad \text{updh}(s, h) = h' \quad s_{ctid}.pc = i \quad P_m[i] \neq \text{start} \quad \langle s_{ctid}.ts, s.\text{mem} \rangle \rightsquigarrow_s \langle \sigma, \mu \rangle \\
 s, h \rightsquigarrow s.[\text{lst}(ctid) := \sigma, \text{mem} := \mu], h'} \text{(Seq)}
 \end{array}$$

**Fig. 4.** Semantics of multithreaded programs

Next, we define non-interference using the evaluation semantics, and characterize the class of schedulers for which our type system enforces non-interference.

### 3 Policy, Schedulers and Type System

This section presents an information flow type system for the concurrent language described in the previous section. The type system, and the policy it enforces, rely on security annotations on programs. The annotations use security levels, drawn from the set  $\mathcal{S} = \{L, H\}$ ; the set  $\mathcal{S}$  is given a lattice structure by declaring  $L \sqsubseteq H$ . Figure 5 gives a summary of the security annotations for each program. They include: i) a function  $\text{ft}$  that assigns security levels to fields; ii) a security environment that assigns security levels to program points, and is used to track implicit flows, and; iii) a method signature table, that assigns to each virtual method a security signature; a signature for a virtual method identifier  $m$  consists of a function  $k_v$  that gives the security levels describing the security level for each of its local variables, a level  $k_h$  that gives the effect of the method on the heap (in the sequel we write  $\text{eff}(m)$  for  $k_h$ ), and a security level  $k_r$  for the value returned.

#### 3.1 Policy

The definition of non-interference relies on memory indistinguishability. Following [1,4], we allow the allocator to be non-deterministic and use a partial bijection  $\beta$  on locations to accommodate dynamic creation of low objects. More precisely, the partial bijection  $\beta$  is used to track the relation between low objects, since, given two identical states,

$$\begin{array}{ll}
 \mathcal{S} \ni s ::= L \mid H & \text{(security levels)} \\
 \mathcal{F} \rightarrow \mathcal{S} \ni \text{ft} & \text{(security level of fields)} \\
 \mathcal{I} \rightarrow \mathcal{S} \ni \Gamma & \text{(local variables security levels)} \\
 \text{SE} = \mathcal{P} \rightarrow \mathcal{S} \ni \text{se} & \text{(security environments)} \\
 \text{Sig} \ni \text{sig} ::= \Gamma \xrightarrow{s} s & \text{(signatures)} \\
 \mathcal{M}_{vt} \rightarrow \text{Sig} \ni \text{mt} & \text{(method table)}
 \end{array}$$

**Fig. 5.** Security setting

a non-deterministic allocator may allocate objects in different locations. High objects (i.e. objects created in a high state) are considered indistinguishable, and the partial bijection  $\beta$  needs not track them.

**Definition 2.** Let  $\beta$  be a partial bijection between locations.

– Value indistinguishability is defined by the clauses:

$$\frac{}{\text{null} \overset{v}{\sim}_{\beta,L} \text{null}} \quad \frac{}{v \overset{v}{\sim}_{\beta,H} v'} \quad \frac{v \in \mathbb{Z}}{v \overset{v}{\sim}_{\beta,L} v} \quad \frac{v, v' \in \mathcal{L} \quad \beta(v) = v'}{v \overset{v}{\sim}_{\beta,L} v'}$$

- Two objects  $o_1, o_2 \in \mathcal{O}$  are indistinguishable, written  $o_1 \overset{o}{\sim}_{\beta} o_2$ , iff  $o_1$  and  $o_2$  are of the same class, and  $o_1.f \overset{v}{\sim}_{\beta, \text{ft}(f)} o_2.f$  for all fields  $f$  in  $o_1$ .
- Two memories  $\mu_1$  and  $\mu_2$  are indistinguishable w.r.t.  $\beta$ , written  $\mu_1 \overset{\mu}{\sim}_{\beta} \mu_2$ , iff  $\text{dom}(\beta) \subseteq \text{dom}(\mu_1)$ ,  $\text{rng}(\beta) \subseteq \text{dom}(\mu_2)$  and for every  $l \in \text{dom}(\beta)$ , we have  $\mu_1(l) \overset{o}{\sim}_{\beta} \mu_2(\beta(l))$ .

We can now define our policy: following [4,5], we consider termination-insensitive non-interference. This policy is more permissive than bisimulation-based properties, which are overly conservative for most purposes.

**Definition 3 (Non-interfering program).** A program  $P$  is non-interfering if for all memories  $\mu_1, \mu_2, \mu'_1, \mu'_2$ , and partial bijection  $\beta$ , there exists a partial bijection  $\beta' \supseteq \beta$  s.t.

$$\mu_1 \overset{\mu}{\sim}_{\beta} \mu_2 \text{ and } P, \mu_1 \Downarrow \mu'_1 \text{ and } P, \mu_2 \Downarrow \mu'_2 \text{ implies } \mu'_1 \overset{\mu}{\sim}_{\beta'} \mu'_2$$

### 3.2 Schedulers

Execution of a concurrent JVM program proceeds by invoking the scheduler to select an active thread, and executing the current instruction for the selected thread. Following [12,13,5], we prevent internal timing leaks by instrumenting the state with security-relevant information, and making the scheduler aware of the security levels at which threads are executing. To this end, we separate thread identifiers in two sets  $\mathcal{T}_H$  of high threads and  $\mathcal{T}_L$  of low threads; moreover, we assume that  $\text{fresh}(\emptyset)$  is a low thread.

The critical property of a secure scheduler is that it always executes critical threads whenever they arise. A low thread becomes critical when it branches over a high value or it performs a method call on a high object. When it happens, we require that the scheduler executes the critical thread, until it terminates the execution of the conditional branch, or it returns from the high method. Formally, we say that a frame  $f$  is low, written  $L(f)$ , if  $se(f.pc) = L$  and  $eff(f.m) = L$ , and that a thread  $t$  is low, written  $L(t)$ , if all its frames are. The set of non-critical and critical threads of a state are defined as:

$$s.\text{lowT} = \{tid \in s.\text{act} \cap \mathcal{T}_L \mid L(s_{tid})\} \quad s.\text{critT} = (s.\text{act} \cap \mathcal{T}_L) \setminus s.\text{lowT}$$

Note that, in comparison to [5], we omit the definition of always high threads, since such an analysis breaks modularity for languages with method calls.

The other essential property of a secure scheduler is that, whenever it chooses a low thread, its choice only depends on the low part of the history. To formalize the latter

notion, we must taint in histories the instances of thread identifiers that correspond to the execution of critical threads. Then, we define the low part  $h_{|L}$  of a history  $h$  as the subsequence of  $h$  where high thread identifiers and tainted instances of thread identifiers are removed.

**Definition 4 (Secure scheduler).** *A function  $\text{pickT} : \text{State} \times \mathcal{H} \rightarrow \mathcal{T}$  is a secure scheduler iff for all states  $s$  and  $s'$  and for all histories  $h$  and  $h'$ :*

1.  $\text{pickT}(\langle s, h \rangle) \in s.\text{act}$
2. if  $s.\text{crit} \neq \emptyset$ , then  $\text{pickT}(\langle s, h \rangle) \in s.\text{crit}$ ;
3. if  $\text{pickT}(s, h) \in s.\text{lowT}$  and  $s'.\text{crit} = \emptyset$  and  $s'.\text{act} \neq \emptyset$ , and  $s.\text{lowT} = s'.\text{lowT}$ , and  $h_{|L} = h'_{|L}$ , then  $\text{pickT}(s, h) = \text{pickT}(s', h')$ .

Examples of secure schedulers include (mildly adapted) round robin schedulers [5].

### 3.3 Control Dependence Regions

The definition of the type system is parametrized by an approximation of control dependence regions. This approximation is used in the typing rule for branching statements to prevent implicit flows.

Control dependence regions are formulated in terms of the successor relation  $\mapsto$ , which is defined by the clauses: i)  $i \mapsto i + 1$  if  $P_m[i]$  is of the form `binop op`, `push n`, `load x`, `store x`, `invokevirtual mID`, `start`, `ifeq j`, `getfield f`, `putfield f`, or `new C`; ii)  $i \mapsto j$  if  $P_m[i]$  is of the form `goto j`, or `ifeq j`; iii)  $i \mapsto$ , i.e.  $i$  does not have a successor if  $P_m[i]$  is a return instruction. We let  $\mapsto^*$  denote the reflexive-transitive closure of  $\mapsto$ . Moreover, we say that  $i$  is a branching point iff there exists program points  $j$  and  $k$  such that  $j \neq k$ ,  $i \mapsto j$  and  $i \mapsto k$ . We let  $\mathcal{P}^\#$  be the set of branching points.

In the sequel, we consider that programs come equipped with a CDR structure  $\langle \text{region}, \text{jun} \rangle$ , where `region` is a function from branching points to sets of program points and `jun` is a partial function from branching points to program points, such that for all program points  $i, j$  and  $k$ :

- CDR1** if  $i \mapsto j$  then  $j \in \text{region}(i)$  or  $j = \text{jun}(i)$ ;
- CDR2** if  $j \in \text{region}(i)$  and  $j \mapsto k$ , then either  $k \in \text{region}(i)$  or  $k = \text{jun}(i)$ ;
- CDR3** if  $j \in \text{region}(i)$  and  $j \mapsto$  then  $\text{jun}(i)$  is undefined;
- CDR4** if  $j \in \text{region}(k)$  and  $i \mapsto j$  then  $i \in \text{region}(k)$  or  $k = i$ ;
- CDR5** if  $\text{jun}(i) \in \text{region}(j)$  then  $\text{region}(i) \subseteq \text{region}(j)$ ;
- CDR6** if  $\text{region}(i) \cap \text{region}(j) \neq \emptyset$  then  $\text{region}(i) \subseteq \text{region}(j)$  or  $\text{region}(j) \subseteq \text{region}(i)$ ;
- CDR7** if  $i \in \text{region}(j)$ , then  $j \mapsto^* i$ .

In comparison with earlier works, e.g. [4], the last four CDR assumptions are new. These additional assumptions have a limited impact in our development: we briefly comment on two key points. On the one hand, our infrastructure requires that a CDR checker is executed prior to type checking. In order to ensure soundness of our verification method, the CDR checker must therefore be enhanced to verify these additional properties; fortunately, the additional verifications can be performed without increasing the complexity of the algorithm. On the other hand, type-preserving compilation

also requires showing that source programs are compiled into JVM programs and CDR structures that are accepted by the checker. Thus, we must verify for each compiler that the CDR structures it outputs verify the additional assumptions. This is the purpose of Proposition 2.

On the other hand, these additional assumptions allow a direct proof of the existence of a function `next` that returns the next (relative to the successor relation) junction point that is not inside a high region. Before we proceed with the definition of the `next` function, we observe that extending CDR checkers so that they verify the above hypotheses, allows us to define a security environment that respects the CDR structure. It is assumed a function `ise` that each branching point is mapped to the security level of the value over which the branch is done.

**Proposition 1 (Next function).** *Let  $\langle \text{region}, \text{jun} \rangle$  be a CDR structure. Moreover, let  $\text{ise} : \mathcal{P}^\# \rightarrow \mathcal{S}$  and for all  $i \in \mathcal{P}$ , let  $\bar{i} = \{k \in \mathcal{P}^\# : i \in \text{region}(k) \wedge \text{ise}(k) = H\}$ . Set the security environment  $se$  as follows:*

$$se(i) = \begin{cases} H & \text{if } \bar{i} \neq \emptyset \\ L & \text{if } \bar{i} = \emptyset \end{cases}$$

Then there exists a function  $\text{next} : \mathcal{P} \rightarrow \mathcal{P}$  such that the following properties hold:

- NePd**  $\text{dom}(\text{next}) = \{i \in \mathcal{P} : se(i) = H\}$
- NeP1**  $i, j \in \text{dom}(\text{next}) \wedge i \mapsto j \Rightarrow \text{next}(i) = \text{next}(j)$
- NeP2**  $i \in \text{dom}(\text{next}) \wedge j \notin \text{dom}(\text{next}) \wedge i \mapsto j \Rightarrow \text{next}(i) = j$
- NeP3**  $j, k \in \text{dom}(\text{next}) \wedge i \notin \text{dom}(\text{next}) \wedge i \mapsto j \wedge i \mapsto k \wedge j \neq k \Rightarrow \text{next}(j) = \text{next}(k)$
- NeP4**  $i, j \in \text{dom}(\text{next}) \wedge k \notin \text{dom}(\text{next}) \wedge i \mapsto j \wedge i \mapsto k \wedge j \neq k \Rightarrow \text{next}(j) = k$

The proof of soundness makes use of the existence of the `next` function.

### 3.4 Type System

The type system adopts the principles of bytecode verification. It performs a modular (method-wise) data flow analysis that operates on stack of security levels, known as security stacks, and outputs a security type, i.e. a mapping from program points to security stacks:

$$\begin{array}{l} \mathcal{ST} \ni st ::= s :: st \mid \epsilon \text{ (security stacks)} \\ \text{Type} = \mathcal{P} \rightarrow \mathcal{ST} \ni S \quad \text{(security types)} \end{array}$$

More specifically, the type system is defined by a set of transfer rules that relate the stack type at a given program point to the stack type of its successors. The transfer rules are given in Figure 6. Their judgments are of the form  $\text{mt}, \text{region}, se, i \vdash st \Rightarrow st'$ , where  $st$  and  $st'$  are stack types; we simply write  $i \vdash st \Rightarrow st'$  when the other components are understood from the context.

**Definition 5 (Typable program).** *A program  $P$  is typable w. type  $S$ , written  $S \vdash P$ , if*

1.  $S(\langle 1, m_{id} \rangle) = \epsilon$ , for all methods  $m_{id}$ ; and
2.  $\text{mt}, \text{region}, i, se \vdash S(i) \Rightarrow st$  where  $st \sqsubseteq S(j)$ , for all program points  $i$  and  $j$  such that  $i \mapsto j$ ; and
3.  $\text{mt}, \text{region}, i, se \vdash S(i) \Rightarrow$ , for all program point  $i$  such that  $P_m[i] = \text{return}$ .

$$\begin{array}{c}
\frac{P_m[i] = \text{push } n}{\text{mt, region, } i, se \vdash st \Rightarrow se(i) :: st} \\
\frac{P_m[i] = \text{goto } j}{\text{mt, region, } i, se \vdash st \Rightarrow st} \\
\frac{P_m[i] = \text{store } x \quad \text{mt}(m) = k_v \xrightarrow{k_h} k_r}{\text{se}(i) \sqcup k \sqsubseteq k_v(x)} \\
\text{mt, region, } i, se \vdash k :: st \Rightarrow st \\
\frac{P_m[i] = \text{ifeq } j \quad \forall j' \in \text{region}(i), k \sqsubseteq se(j')}{\text{mt, region, } i, se \vdash k :: st \Rightarrow \text{lift}_k(st)} \\
\frac{P_m[i] = \text{start} \quad k \sqcup se(i) \sqcup \text{eff}(m) \sqsubseteq \text{eff}(\text{run})}{\text{mt, region, } i, se \vdash k :: st \Rightarrow st} \\
\frac{P_m[i] = \text{putfield } f \quad k_1 \sqcup se(i) \sqcup k_2 \sqsubseteq \text{ft}(f) \quad \text{eff}(m) \sqsubseteq \text{ft}(f)}{\text{mt, region, } i, se \vdash k_1 :: k_2 :: st \Rightarrow st} \\
\frac{P_m[i] = \text{load } x \quad \text{mt}(m) = k_v \xrightarrow{k_h} k_r}{\text{mt, region, } i, se \vdash st \Rightarrow (se(i) \sqcup k_v(x)) :: st} \\
\frac{P_m[i] = \text{binop } op}{\text{mt, region, } i, se \vdash k_1 :: k_2 :: st \Rightarrow (se(i) \sqcup k_1 \sqcup k_2) :: st} \\
\frac{P_m[i] = \text{invokevirtual } m_{vt} \quad \text{mt}(m_{ID}) = k_v \xrightarrow{k_h} k_r}{k :: \bar{k} \sqsubseteq k_v \cdot \text{arg}_{m_{vt}} \quad k \sqcup se(i) \sqcup \text{eff}(m) \sqsubseteq k_h} \\
\text{mt, region, } i, se \vdash \bar{k} :: k :: st \Rightarrow (k_r \sqcup k_h) :: st \\
\frac{P_m[i] = \text{return} \quad se(i) \sqcup k \sqsubseteq k_r \quad se(i) \sqsubseteq \text{eff}(m)}{\text{mt, region, } i, se \vdash k :: st \Rightarrow} \\
\frac{P_m[i] = \text{getfield } f}{\text{mt, region, } i, se \vdash k :: st \Rightarrow (k \sqcup \text{ft}(f) \sqcup se(i)) :: st}
\end{array}$$

Fig. 6. Transfer rules

## 4 Soundness

This section establishes the soundness property for the type system presented in the previous section: we prove that typable programs are non-interfering. The overall structure of the proof is similar to [4,5]: one first proves locally-respects and step-consistent unwinding lemmas, and then one derives a locally-respects unwinding lemma for a notion of visible execution; finally, one proves non-interference by an argument on traces. The former critically relies on an appropriate notion of state indistinguishability, whereas the latter relies on properties of the control dependence regions and the security environment. In both cases, significant extensions to [4,5] are required. In particular, the notion of state equivalence must be extended to accommodate threads, i.e. stacks of frames, rather than a single frame. The intuition is that two threads are indistinguishable if their frames are in one-to-one correspondence<sup>1</sup> until there is a call to a high method in both threads, or one of the threads enters in a high branch. Throughout this section, we assume that  $P$  is typable program whose type is  $S$ , i.e.  $S \vdash P$ .

We first define state equivalence. As the high threads cannot modify the low part of the heap, state equivalence only depends on low threads, and the heap.

**Definition 6 (State equivalence).** *Let  $\beta$  be a partial bijection on locations.*

- *Operand stack equivalence is defined inductively by the clauses:*

$$\frac{\text{highos}(os, st) \quad \text{highos}(os', st')}{os : s \overset{\text{os}}{\sim}_{\beta} os' : st'} \quad \frac{os : st \overset{\text{os}}{\sim}_{\beta} os' : st' \quad v \overset{v}{\sim}_{\beta, k} v'}{v :: os : k :: st \overset{\text{os}}{\sim}_{\beta} v' :: os' : k :: st'}$$

where  $\text{highos}(os, st)$  iff  $|os| = |st|$  and  $st(i) = H$  for all  $1 \leq i \leq |st|$ .

<sup>1</sup> This correspondence is expressed below by the auxiliary relation  $\overset{\ell}{\sim}$ .

– Frame equivalence is defined by the clause:

$$f \overset{s}{\sim}_{\beta} f' \text{ iff } f.\text{os} : S(f.\text{pc}) \overset{\text{os}}{\sim}_{\beta} f'.\text{os} : S(f'.\text{pc}) \text{ and} \\ f.\text{m} = f'.\text{m} \text{ and} \\ f.\text{env} \overset{\text{env}}{\sim}_{\beta, f.\text{m}} f'.\text{env}$$

where  $\rho \overset{\text{env}}{\sim}_{\beta, m} \rho'$  iff  $\rho(x) \overset{v}{\sim}_{\beta, k_v(x)} \rho'(x)$  and  $\text{mt}(m) = k_v \xrightarrow{k_h} k_r$ .

– Thread equivalence is defined inductively by the clauses:

$$\frac{}{\epsilon \overset{\ell}{\sim}_{\beta} \epsilon} \quad \frac{\text{high}(fs_0) \quad \text{high}(fs'_0) \quad fs \overset{\ell}{\sim}_{\beta} fs'}{fs_0 :: fs \overset{t}{\sim}_{\beta} fs'_0 :: fs'}$$

$$\frac{f \overset{s}{\sim}_{\beta} f' \quad \text{eff}(f.\text{m}) = L \quad fs \overset{\ell}{\sim}_{\beta} fs' \quad L(f) \wedge L(f') \Rightarrow f.\text{pc} = f'.\text{pc}}{f :: fs \overset{\ell}{\sim}_{\beta} f' :: fs'}$$

where  $\text{high}(fs)$  means that  $\text{eff}(fs[i].\text{m}) = H$  for all  $1 \leq i \leq |fs|$ .

– State equivalence is defined by the clause:

$$s \overset{c}{\sim}_{\beta} s' \text{ iff } s.\text{mem} \overset{h}{\sim}_{\beta} s'.\text{mem} \wedge \forall t \in s.\text{act} \cap \mathcal{T}_{L.s_t} \overset{t}{\sim}_{\beta} s'_t$$

The proof of the soundness theorem introduces three new execution relations, including a notion of visible step; the formal definitions are given in Figure 7. An intuition for these execution relations is:  $s, h \rightsquigarrow_{\text{vis}} s', h'$  iff neither  $s$  or  $s'$  have a critical thread, and  $s, h \rightsquigarrow s', h'$  or  $s, h \rightsquigarrow^* s', h'$  and all intermediate states have a critical thread. For the proofs, it is convenient to divide this relation into two disjoint ones, corresponding to the reasons under which a thread can be made critical. Specifically, we define  $s, h \rightsquigarrow_{\text{method}} s', h'$  iff  $s, h \rightsquigarrow^* s', h'$  and  $s, s'$  and all the intermediate states have a critical thread caused by a method call on a high object; and  $s, h \rightsquigarrow_{\text{branch}} s', h'$  iff  $s, h \rightsquigarrow^* s', h'$  and  $s, s'$  and all the intermediate states have a critical thread caused by a branching over a high value. The predicates  $C_m$  and  $C_b$  determine whether a thread is high because of a branch or a method call, and  $\text{NCT}$ ,  $\text{CT}_b$  and  $\text{CT}_m$  respectively indicate that the state has no critical thread; or a critical thread because of a branch; or a critical thread because of a method call.  $\text{VS}$  predicate is used to ensure threads with high identifiers do actually correspond to high threads.

Then, we prove unwinding lemmas for the newly introduced execution relations. The first two have the form of step-consistent lemmas, while the third has the form of a locally-respect lemma. For the sake of readability, we extend the next function to (critical) threads and (critical) states. For a critical thread  $t$ ,  $\text{next}(t)$  is defined as  $\text{next}(i)$  where  $i$  is the program point that caused the thread to become critical. Formally,

$$\text{next}(t) = \begin{cases} \text{next}(\text{hd}(t).\text{pc}) & \text{if } se(\text{hd}(t).\text{pc}) = H \wedge \text{eff}(\text{hd}(t).\text{m}) = L \wedge L(\text{tl}(t)) \\ \text{next}(\text{tl}(t)) & \text{otherwise} \end{cases}$$

Since the current thread (the one picked by the scheduler) can be deduced from the context, we further write  $\text{next}(s)$  for  $\text{next}(s_{ctid})$  where  $ctid$  is the current thread.

$$\begin{aligned}
 C_m(f) &\equiv \text{eff}(f.m) = H \wedge \text{highos}(f.pc, S(f.pc)) & C_m(t) &\equiv \exists i. C_m(t[i]) \wedge \forall j < i. L(t[j]) \\
 C_b(f) &\equiv se(f.pc) = H \wedge \text{eff}(f.m) = L & C_b(t) &\equiv \exists i. C_b(t[i]) \wedge \forall j < i. L(t[j]) \\
 CT_b(s) &\equiv s.\text{critT} = \{tid\} \wedge C_b(s_{tid}) \wedge VS(s) & VS(s) &\equiv \forall tid \in s.\text{act} \cap \mathcal{T}_H. C_m(s_{tid}) \\
 CT_m(s) &\equiv s.\text{critT} = \{tid\} \wedge C_m(s_{tid}) \wedge VS(s) & NCT(s) &\equiv s.\text{critT} = \emptyset \wedge VS(s)
 \end{aligned}$$

$$\begin{array}{c}
 \frac{NCT(s) \quad NCT(s') \quad s, h \rightsquigarrow s', h'}{s, h \rightsquigarrow_{\text{vis}} s', h'} \\
 \\
 \frac{CT_b(s) \quad CT_b(s') \quad s, h \rightsquigarrow s', h'}{s, h \rightsquigarrow_{\text{branch}} s', h'} \quad \frac{CT_m(s) \quad CT_m(s') \quad s, h \rightsquigarrow s', h'}{s, h \rightsquigarrow_{\text{method}} s', h'} \\
 \\
 \frac{s, h \rightsquigarrow_{\text{branch}} s', h' \quad s', h' \rightsquigarrow_{\text{branch}} s'', h''}{s, h \rightsquigarrow_{\text{branch}} s'', h''} \quad \frac{s, h \rightsquigarrow_{\text{method}} s', h' \quad s', h' \rightsquigarrow_{\text{method}} s'', h''}{s, h \rightsquigarrow_{\text{method}} s'', h''} \\
 \\
 \frac{s, h \rightsquigarrow s', h' \quad s', h' \rightsquigarrow_{\text{branch}} s'', h'' \quad s'', h'' \rightsquigarrow s''', h'''}{s, h \rightsquigarrow_{\text{vis}} s''', h'''} \\
 \\
 \frac{s, h \rightsquigarrow s', h' \quad s', h' \rightsquigarrow_{\text{method}} s'', h'' \quad s'', h'' \rightsquigarrow s''', h'''}{s, h \rightsquigarrow_{\text{vis}} s''', h'''}
 \end{array}$$

Fig. 7. Auxiliary execution relations

**Lemma 1.** *If  $s, h_s \rightsquigarrow_{\text{branch}} s', h_{s'}$  and  $s \stackrel{c}{\sim}_{\beta} t$  then we have  $s' \stackrel{c}{\sim}_{\beta} t$ ,  $h_s \stackrel{h}{\sim} h_{s'}$ , and  $\text{next}(s) = \text{next}(s')$ .*

**Lemma 2.** *If  $s, h_s \rightsquigarrow_{\text{method}} s', h_{s'}$  and  $s \stackrel{c}{\sim}_{\beta} t$  then we have  $s' \stackrel{c}{\sim}_{\beta} t$ , and  $h_s \stackrel{h}{\sim} h_{s'}$ .*

**Lemma 3.** *If  $s, h_s \rightsquigarrow_{\text{vis}} s', h_{s'}$  and  $t, h_t \rightsquigarrow_{\text{vis}} t', h_{t'}$  and  $s \stackrel{c}{\sim}_{\beta} t$  and  $h_s \stackrel{h}{\sim} h_t$  then we have  $s' \stackrel{c}{\sim}_{\beta'} t'$  and  $h_{s'} \stackrel{h}{\sim} h_{t'}$  for some  $\beta'$  such that  $\beta \subseteq \beta'$ .*

**Theorem 1 (Soundness).** *If  $P$  is typable with respect to a CDR structure  $\langle \text{region}, \text{jun} \rangle$ , then  $P$  is non-interfering.*

*Proof (Sketch of proof).* Suppose that  $\mu_1 \stackrel{\mu}{\sim}_{\beta} \mu_2$ ,  $P, \mu_1 \Downarrow \mu'_1$  and  $P, \mu_2 \Downarrow \mu'_2$ . Note that if  $P, \mu_1 \Downarrow \mu'_1$  then  $s_{\text{init}}(\mu_1), \epsilon (\rightsquigarrow_{\text{vis}})^* s_1, h_{s_1}$  with  $s_1.\text{lowT} = \emptyset$  and  $s_1.\text{mem} = \mu'_1$ . Analogously,  $s_{\text{init}}(\mu_2), \epsilon (\rightsquigarrow_{\text{vis}})^* s_2, h_{s_2}$  with  $s_2.\text{lowT} = \emptyset$  and  $s_2.\text{mem} = \mu'_2$ . Moreover, as the  $\rightsquigarrow_{\text{vis}}$  relation “steps over” the critical execution steps both executions take the same number of  $\rightsquigarrow_{\text{vis}}$  steps. By lemma 3, we conclude.

## 5 Type-Preserving Compilation

The purpose of this section is to define an information flow type system for a concurrent fragment of Java, and to show that compilation transforms typable programs into typable programs.

### 5.1 Source Language and Compiler

The expressions and commands of the source language are defined in Figure 8; the overall structure of a program remains identical. The compiler must produce, in addition to

$$\begin{aligned}
 e &::= n \mid x \mid \text{op } e \mid \text{new } C \\
 c &::= x = e \mid c ; c \mid \text{if } e \text{ then } c \text{ else } c \mid \text{while } e \text{ do } c \mid e.f = e \mid \\
 &\quad \text{return } e \mid e.\text{start}() \mid x = e.f \mid x = e.m(\bar{e}) \\
 p &::= [m(\bar{e})\{c ; \text{return } e\}]
 \end{aligned}$$

**Fig. 8.** Source language

the JVM program, a CDR structure. It is convenient for constructing the latter to consider a labeled version of the commands, and to define a mapping  $\mathcal{A}$  that automatically give the commands of the source program labels such that there are no two commands in the source that have the same label; the syntax of the labeled commands in Figure 9. The definition of  $\mathcal{A}$  can be found in the full version of the paper. The function  $\mathcal{A}$  takes an integer, the number from where it starts to label the subcommands, and a command to be labeled, and returns a labeled command. To label a program, we use  $\mathcal{A}$  in each method using as integer the length of the previous method compiled plus one and one for the first method to be compiled.

$$\begin{aligned}
 c &::= [x = e]^n \mid c ; c \mid [\text{if } e \text{ then } c \text{ else } c]^n \mid [\text{while } e \text{ do } c]^n \mid [e.f = e]^n \mid \\
 &\quad [\text{return } e]^n \mid [e.\text{start}()]^n \mid [x = e.f]^n \mid [x = e.m(\bar{e})]^n
 \end{aligned}$$

**Fig. 9.** Labeled syntax

The compiler functions  $\llbracket - \rrbracket$  transform expressions and commands of the source language as lists of JVM instructions; their definition is given in Figure 10. As in [3], we define for each branching command of the source language a *main instruction* in the target code. This instruction is a branching program point in the target program, and hence a program point for which we need to define the region and the junction point. Finally, programs are compiled by applying  $\llbracket - \rrbracket$  to each method body.

### 5.2 Source Type System

The type system for the source language is defined in Figure 11. As for the JVM, the type system assumes that the program comes equipped with a mapping of fields to security levels, and a security signature for each method. Following the standard approach for source type system, we first type the expressions and then the commands. A method is well-typed if its command is typed with a type that is compatible with the method signature. A program is well-typed if all its methods are well-typed.

$$\begin{aligned}
\llbracket x \rrbracket &= \text{load } x \\
\llbracket n \rrbracket &= \text{push } n \\
\llbracket e \text{ op } e' \rrbracket &= \llbracket e \rrbracket :: \llbracket e' \rrbracket :: \text{binop } op \\
\llbracket \text{new } C \rrbracket &= \text{new } C \\
\llbracket x = e \rrbracket &= \llbracket e \rrbracket :: \text{store } x \\
\llbracket c_1 ; c_2 \rrbracket &= \llbracket c_1 \rrbracket :: \llbracket c_2 \rrbracket \\
\llbracket e.f = e' \rrbracket &= \llbracket e' \rrbracket :: \llbracket e \rrbracket :: \text{putfield } f \\
\llbracket \text{return } e \rrbracket &= \llbracket e \rrbracket :: \text{return} \\
\llbracket \text{while } e \text{ do } c \rrbracket &= \text{goto } (pc + \#lc + 1) :: lc :: le :: \underline{\text{ifeq } (pc - \#lc - \#le)} \\
&\quad \text{where } le = \llbracket e \rrbracket, lc = \llbracket c \rrbracket \\
\llbracket \text{if } e \text{ then } c_1 \text{ else } c_2 \rrbracket &= le :: \underline{\text{ifeq } (pc + \#lc_2 + 2)} :: lc_2 :: \text{goto } (pc + \#lc_1 + 1) :: lc_1 \\
&\quad \text{where } le = \llbracket e \rrbracket, lc_1 = \llbracket c_1 \rrbracket, lc_2 = \llbracket c_2 \rrbracket \\
\llbracket e.\text{start}() \rrbracket &= \llbracket e \rrbracket :: \text{start} \\
\llbracket x = e.f \rrbracket &= \llbracket e \rrbracket :: \text{getfield } f :: \text{store } x \\
\llbracket x = e.m(\bar{e}) \rrbracket &= \llbracket \bar{e} \rrbracket :: \llbracket e \rrbracket :: \text{invokevirtual } m :: \text{store } x
\end{aligned}$$

Fig. 10. Compilation

### 5.3 Preservation Proof

As in [3], the preservation proof preserves in three steps: first, we define a function sregion for the source language, which is used later to obtain a CDR structure for the compiled program. Then, we define an intermediate type system for the source language, and prove its equivalence with the original type system; the intermediate type system is defined with help of the sregion function, and facilitates the third and final step, which is the proof of type-preserving compilation itself. We highlight some critical steps below.

We first define the sregion function, which is the equivalent of a CDR structure for the source language: for each branching command  $[c]^n$ , we let  $\text{sregion}(n)$  be the set of labels that appear in  $c$ , i.e. that label the subcommands of  $c$ .

**Proposition 2.** *For every program  $P$ ,  $\langle \text{tregion}, \text{tjun} \rangle$  is a CDR structure for the program  $\llbracket P \rrbracket$ , where for every branching instruction  $[c]^n$  in the source code,  $\text{tregion}(n)$  is defined as the set of instructions obtained by compiling commands  $[c']^{n'}$ , where  $n' \in \text{sregion}(n)$ , and  $\text{tjun}(n) = \max \{i : i \in \text{tregion}(n)\} + 1$ .*

The proof involves showing that every branching program point in the target language is the image of a branching program point in the source language, and derive the CDR properties from the definition of the compiler.

We now define a function  $ise$  that associates to each branching point in the compiled program the security level of the corresponding guard in the source language. Proposition 1 yields the existence of a function  $\text{next}$  that satisfies the **NeP** hypotheses w.r.t. the security environment built from  $ise$ .

The proof of type-preserving compilation between the source and target type systems is performed first on expressions, then on commands. Finally, we conclude that

$$\begin{array}{c}
\frac{k_v \vdash e : k \quad k \sqsubseteq k_v(x)}{k_v \xrightarrow{k_h} k_r \vdash x = e : k_v(x)} \quad \frac{k_v \vdash e : k \quad k_v \xrightarrow{k_h} k_r \vdash c : k}{k_v \xrightarrow{k_h} k_r \vdash \text{while } e \text{ do } c : k} \quad \frac{k_v \vdash e : k \quad k_v \xrightarrow{k_h} k_r \vdash c : k}{k_v \xrightarrow{k_h} k_r \vdash c' : k} \\
\frac{k_v \vdash e : k \quad k \sqsubseteq k_r \sqcap k_h}{k_v \xrightarrow{k_h} k_r \vdash \text{return } e : k} \quad \frac{k_v \vdash e : k \quad k \sqcup \text{ft}(f) \sqsubseteq k_v(x)}{k_v \xrightarrow{k_h} k_r \vdash x = e.f : k_v(x)} \quad \frac{k_v \vdash e : k \quad k_v \vdash e' : k' \quad k \sqcup k' \sqcup k_h \sqsubseteq \text{ft}(f)}{k_v \xrightarrow{k_h} k_r \vdash e.f = e' : \text{ft}(f)} \\
\frac{k_v \xrightarrow{k_h} k_r \vdash c : k \quad k' \sqsubseteq k}{k_v \xrightarrow{k_h} k_r \vdash c : k'} \quad \frac{k_v \vdash e : k \quad k \sqcup k_h \sqsubseteq \text{eff}(\text{run})}{k_v \xrightarrow{k_h} k_r \vdash e.\text{start}() : \text{eff}(\text{run})} \quad \frac{k_v \xrightarrow{k_h} k_r \vdash c : k \quad k_v \xrightarrow{k_h} k_r \vdash c' : k'}{k_v \xrightarrow{k_h} k_r \vdash c; c' : k \sqcap k'} \\
\frac{\forall i. k_v \vdash e[i] : k_v(\bar{e}[i]) \quad k \sqcup k_h \sqsubseteq k'_h \quad k_v \vdash e : k \quad \text{mt}(m) = k'_v \xrightarrow{k'_h} k'_r \quad k_r \sqsubseteq k'_v(x)}{k_v \xrightarrow{k_h} k_r \vdash x = e.m(\bar{e}) : k_v(x) \sqcap \text{eff}(m)}
\end{array}$$

Fig. 11. A high-level type system for the source language

a typable source program is transformed into a JVM program that is typable w.r.t. the same policy as the source program, the CDR structure generated by the compiler, and the security environment built from *ise*.

**Theorem 2.** *Suppose we have  $a \vdash SP$ , then  $\vdash \llbracket SP \rrbracket$ .*

## 6 A Toy Example

To explore and concretise some details of the type system, we introduce a toy example. Three classes are defined: A, B, and C, with fields defined  $\text{fieldsof}(A) = \{\text{varA}\}$ ,  $\text{fieldsof}(B) = \{\text{varB}\}$ ,  $\text{fieldsof}(C) = \{\text{varL}, \text{varH}\}$ , and  $\text{ft}$  defined by the rules  $\text{ft}(\text{varL}) = L$ , and  $\text{ft}(\text{varH}) = \text{ft}(\text{varA}) = \text{ft}(\text{varB}) = H$ . The methods are defined in Fig. 12.

From the type system,  $\text{o}$  is high as it is set inside a high-branch. Also, as  $\text{o}$  is high, from the start typing rule, we conclude that  $\text{eff}(\text{run})$  must be high as well. This is consistent with the typing rule for assignment, as the fields  $\text{varA}$  and  $\text{varB}$  are high too.

A small change would make the program become insecure and untypable. If we force one of  $\text{varA}$  or  $\text{varB}$  to be a low field, then the type system would reject the code as the value of  $\text{varH}$  could be determined by inspecting  $\text{varA}$  or  $\text{varB}$ .

From this example, we can learn that if a method implementing  $\text{run}$  is high, then all other instances are forced to be high too. While it seems that this could be fixed in part by considering whether  $\text{start}$  is launched from a high or low region/object, the problem

```

void A.run() {
  this.varA = 0;
}
void B.run() {
  this.varB = 1;
}
void C.run() {
  if (this.varH)
    o = new A();
  else
    o = new B();
  o.start();
  ...
}

```

Fig. 12. Code of example

is actually harder than this: we can arrive to a start instruction from different methods with different security levels.

## 7 Related Work

The prevalence of Java in mobile code applications makes it a natural target for information flow analysis, and there is a substantial body of literature on enforcement mechanisms for Java or JVM programs. In particular, Myers *et al* [9] have developed the Jif language, an information flow aware extension of (sequential) Java; Jif has been used for developing significant examples of secure applications, including the Civitas voting system. While there is no soundness proof for the whole Jif type system, Banerjee and Naumann [1] have developed a sound information flow type system for a significant fragment of Java. Barthe, Naumann and Rezk [3] extend the type system and soundness proof to a simplified form of exception. Our work adopts many notions and techniques from [1,9] for object-oriented programs, and [4,8] for bytecode languages. All these works are confined to sequential fragments of Java.

More recently, there have been many proposals for a dynamic approach to information flow, as developed by Le Guernic *et al* for a core imperative [7] and concurrent [6] languages. Nair *et al* [10] report on the implementation of the Trishul system, which dynamically tracks information flow in Java applications; however, no formal guarantee is established.

## 8 Conclusion

Developing sound and flexible information flow type systems for mobile code is an important goal for language-based security. This article leverages previous works on concurrent bytecode languages and on the sequential JVM to propose for the first time a sound information flow type system for a concurrent low-level object-oriented language. Our work, and in particular the notions of program equivalence we use, lay the foundations for designing flexible type systems for the full JVM. In the future, it would be interesting to extend the type system to exceptions (which are treated in [11]), declassification (in the style of [2]), and locks.

## References

1. Banerjee, A., Naumann, D.: Stack-based access control for secure information flow. *Journal of Functional Programming* 15, 131–177 (2005); Special Issue on Language-Based Security
2. Barthe, G., Cavadini, S., Rezk, T.: Tractable enforcement of declassification policies. In: *IEEE Computer Security Foundations Symposium* (June 2008)
3. Barthe, G., Naumann, D., Rezk, T.: Deriving an information flow checker and certifying compiler for Java. In: *Symposium on Security and Privacy*. IEEE Press (2006)
4. Barthe, G., Pichardie, D., Rezk, T.: A Certified Lightweight Non-interference Java Bytecode Verifier. In: De Nicola, R. (ed.) *ESOP 2007*. LNCS, vol. 4421, pp. 125–140. Springer, Heidelberg (2007)

5. Barthe, G., Rezk, T., Russo, A., Sabelfeld, A.: Security of multithreaded programs by compilation. *ACM Trans. Inf. Syst. Secur.* 13(3) (2010)
6. Le Guernic, G.: Automaton-based confidentiality monitoring of concurrent programs. In: *CSF*, pp. 218–232. IEEE Computer Society (2007)
7. Le Guernic, G., Banerjee, A., Jensen, T., Schmidt, D.A.: Automata-Based Confidentiality Monitoring. In: Okada, M., Satoh, I. (eds.) *ASIAN 2006*. LNCS, vol. 4435, pp. 75–89. Springer, Heidelberg (2007)
8. Kobayashi, N., Shirane, K.: Type-based information analysis for low-level languages. In: *Asian Programming Languages and Systems Symposium*, pp. 302–316 (2002)
9. Myers, A.C.: JFlow: Practical mostly-static information flow control. In: *Principles of Programming Languages*, pp. 228–241. ACM Press (1999), Ongoing development at <http://www.cs.cornell.edu/jif/>
10. Nair, S.K., Simpson, P.N.D., Crispo, B., Tanenbaum, A.S.: A virtual machine based information flow control system for policy enforcement. *Electr. Notes Theor. Comput. Sci.* 197(1), 3–16 (2008)
11. Rezk, T.: Verification of confidentiality policies for mobile code. PhD thesis, Université de Nice Sophia-Antipolis (2006)
12. Russo, A., Sabelfeld, A.: Securing interaction between threads and the scheduler. In: *Computer Security Foundations Workshop*, pp. 177–189 (2006)
13. Russo, A., Sabelfeld, A.: Security for Multithreaded Programs Under Cooperative Scheduling. In: Virbitskaite, I., Voronkov, A. (eds.) *PSI 2006*. LNCS, vol. 4378, pp. 474–480. Springer, Heidelberg (2007)
14. Smith, G., Volpano, D.: Secure Information Flow in a Multi-threaded Imperative Language. In: *Principles of Programming Languages*, pp. 355–364 (1998)