

Secure Information Flow and Pointer Confinement in a Java-like Language

Anindya Banerjee*
Computing and Information Sciences
Kansas State University
Manhattan KS 66506 USA
ab@cis.ksu.edu

David A. Naumann†
Computer Science
Stevens Institute of Technology
Hoboken NJ 07030 USA
naumann@cs.stevens-tech.edu

Abstract

We consider a sequential object-oriented language with pointers and mutable state, private fields and class-based visibility, dynamic binding and inheritance, recursive classes, casts and type tests, and recursive methods. Programs are annotated with security levels, constrained by security typing rules. A noninterference theorem shows how the rules ensure pointer confinement and secure information flow.

1. Introduction

There are many channels by which sensitive information can be leaked. This paper is concerned with information flows that arise in sequential object-oriented programs due to control flow, data flow, and dynamic memory allocation. Inspired by Denning’s work [10, 11], Volpano, Smith and Irvine devised an elegant, syntax-directed type system for annotating program variables, commands, and procedure parameters with security levels [34, 32]. Goguen and Meseguer [13] introduced noninterference, expressed in terms of suitable simulation relations, to formalize information flow policies. Volpano and Smith proved that their type system ensures noninterference [34, 32].

Subsequently, several researchers have given similar analyses for possibilistic and probabilistic noninterference for multi-threaded programs [28, 33, 25, 19, 27]. Barthe and Serpette prove noninterference for a purely functional instance-based object calculus [4]. For sequential programs, Abadi *et al.* [1], Sabelfeld and Sands [26] and Heintze and Riecke [15] consider higher order procedures. They also make explicit the connections between the relational formulation of noninterference and other dependency analyses such as slicing and binding time analysis [1, 4],

building noninterference properties into the semantics in the manner of Reynolds’ relationally parametric models [24]. However, it is difficult to extend such models in a tractable way to encompass language features such as recursive types and shared mutable objects which are extensively used in languages such as Java [2].

Our contribution is to deal with dynamic memory allocation and object-oriented constructs: we prove noninterference for a sequential object-oriented language with pointers and mutable state, private fields and class-based visibility, dynamic binding and inheritance, casts and type tests, and mutually recursive classes and methods. The security type system extends that of Volpano and Smith [32] to encompass data flow via mutable object fields and control flow in dynamically dispatched method calls.

Myers [20] gave a security typing system for a fragment of Java even richer than ours, but left open the problem of justifying the rules with a noninterference result. This is hardly surprising, as the rules are quite complicated. Some of the complications are inherent in the complexity of the language; others are introduced with the aim of accommodating dynamic access control and sophisticated security policies including declassification [12, 21, 20, 35].

In the present paper, we confine attention to the problem of proving noninterference for a “realistic” sequential language (not far from JavaCard [7]), using conventional annotations without declassification or dynamic access control. Our results are given in elementary terms. We eschew the elegant structures used in [15, 1, 26], but we can give detailed proofs in the space of a few pages. This may also help in extending our results to other language features.

We use the weak form of noninterference which does not consider termination to be observable. Strong noninterference is treated by Volpano and Smith [32] (and others), but for a sequential language this requires loop guards to have low security. We omit loops but include recursion which also admits nontermination. Thus to extend their ideas to our language we would require low security guards for conditionals that involve recursive calls, a complication

*Supported by NSF grants EIA-9806835 and CCR-0296182

†Supported by NSF grant INT-9813854

we choose to avoid here. The extension of strong noninterference to recursive procedures merits study in a simple setting before it is combined with the features of an object-oriented language.

Our work grew out of a study of data abstraction for Java [3]. We found that a straightforward compositional semantics is adequate even in the presence of recursive types and dynamically bound method calls (which are typically viewed as being akin to higher order procedures). The semantics is simple enough to extend easily to additional constructs, e.g., access control is included in [3].

Due to pointer aliasing, the language is not relationally parametric *per se*. But suitable confinement of pointers suffices to yield a strong representation-independence result for user-defined abstractions [3]. The term “confinement” appears to originate in the literature on operating system security, but its use is natural in object-oriented programming where pointer confinement has been proposed for encapsulation at the level of modules, classes, or instances [16, 8, 18, 5, 31]. For information flow, we impose a confinement condition on high-security pointers.

The following section is a detailed overview of the paper. Section 3 formalizes the language and its semantics. Section 4 gives the security typing system. Section 5 deals with confinement, which is then used in Section 6 to prove the noninterference theorem. Section 7 considers related work and prospects for further advance.

2. Overview

We consider a language that is quite complicated relative to those for which noninterference results have been proved before, but we use simple security annotations that generalize those of Volpano and Smith. We annotate local variables, fields, and method parameters using types (T, κ) where T is an ordinary program type and κ is one of the two security levels H and L . Generalization to a lattice of levels would complicate notations without adding illumination. We also annotate classes. The security level of a class is the security level of `this`, i.e., the target of a method call, and it is also used for confinement as described in the sequel.

In the following example there is a single field `f` of high security (level H), in a class named `C` of level L .

```
class C L extends D {
  (bool,H) f;

  (bool,H) m ((bool,L) x) L {
    if (x) this.f := not x;
    else  this.f := x;
    return (x == this.f); } }
```

Our typing system assigns to method `m` the type $x : (\text{bool}, L) \xrightarrow{L} (\text{bool}, H)$. This designates a method which takes a parameter of level L , has effect L on the heap, and returns a value of level H . The heap effect L is given in the method declaration, following the parameter list $((\text{bool}, L) x)$.

The type of the result expression $(x == \text{this.f})$ is (bool, H) because the type of `this.f` is (bool, H) . Although both `this` and `x` are L , field `f` is declared H .

In typing judgements, Volpano and Smith use a type H `cmd` for commands that assign only to H variables. Our typing judgements use a command type $(\text{com } \kappa_1, \kappa_2)$ for commands that assign only to variables of level at least κ_1 and to object fields of level at least κ_2 . That is, the heap effect designated by a command type is concerned with field levels. The same is true of the heap effect designated above the arrow in a method type like $x : (\text{bool}, L) \xrightarrow{L} (\text{bool}, H)$.

If we modify the example to declare a heap effect H for the method, resulting in a type $x : (\text{bool}, L) \xrightarrow{H} (\text{bool}, H)$, the class is still typable because the only effect is on a field of level H .

Consider the following variation.

```
class C H extends D {
  (bool,??) f;
  (bool,H) m ((bool,L) x) L {
    if (x) this.f := not x;
    else  this.f := x;
    return (x == this.f); } }
```

The type of `m` is again $x : (\text{bool}, L) \xrightarrow{L} (\text{bool}, H)$ but the class is H . What level, marked `??`, can be used for field `f`? The conditional statement needs to be typable in context $x : (\text{bool}, L), \text{this} : (C, H)$ because the label on a class designates the level of `this`. The conditional should be given type $(\text{com } ?, L)$, as the heap effect for the method is L . What about the effect, marked `?`, on the local environment (i.e., local variables and parameters)? Our rule for field update requires that the level of the field be at least the level of the assigned variable: as `this` is H , field `f` needs to be declared as H and then the conditional can be typed as $(\text{com } H, L)$.

Subclassing in Java is “invariant” in the sense that method signatures cannot be specialized in subclasses. Although other alternatives merit study, our typing system is the same: the declared security levels of a method cannot be changed in subclasses. However, we allow the subclass of an L class to be declared H . This has interesting consequences and it is one of the reasons why we need a semantic notion of pointer confinement, called *L*-confinement, as discussed later.

These examples are far from an exhaustive illustration of the interesting patterns that arise. Nor can a brief exposition

give a thorough justification for the rules. A few more examples appear in Section 4.1. Although our type system is quite general, it disallows some sensible programs such as those involving declassification. On the other hand, it admits some declarations that are sound but not very useful, such as variables of type (C, L) for H-class C .

In Section 3 we give the formal syntax and denotational semantics of the underlying language for which we consider security annotations. Java is sufficiently complicated that it is a challenge to formalize its syntax in a readable way. We adapt the formalization of a smaller fragment of Java, FJ [17], which mixes standard notations from type theory with Java-like syntax. We extend FJ by adding imperative features, and modify it by treating fields as private (class-visible) rather than public. The semantics is defined in terms of an ordinary typing system; it does not depend on security annotations.

In Section 4 we give the typing system for security annotations. In practice, one would want to specify security policy by labelling certain inputs and outputs, leaving the rest to automated inference, but that is beyond the scope of the paper, as is label polymorphism. For expository purposes, it is convenient to use separate subsumption and subtyping rules that express, for example, that an L expression can always be used in a context where an H expression is allowed. But for proofs it is more convenient to use syntax-directed rules that incorporate subsumption. For lack of space we give only the syntax-directed rules for our system, at the cost that the rules have rather many constraints on security levels. Readers unfamiliar with this style of specifying a flow analysis are encouraged to read Volpano and Smith’s clear and succinct presentation [32].

Our main result, in Section 6, is a noninterference theorem: For a program with annotations satisfying the rules, fields, variables, and method results with label L do not depend on those with label H. As in work cited in the introduction, noninterference is formalized in terms of an equivalence relation on states. Here a state consists of an environment η and a heap h . For two states (η, h) and (η', h') to be equivalent, written $(\eta, h) \sim (\eta', h')$, means that they agree on L-variables and on L-fields of heap objects. The noninterference theorem says that if a typable program is executed in the two related states (η, h) and (η', h') , the outcomes are also related. What this means is that changes to high security inputs cannot be observed at low outputs.

The main noninterference result pertains to the methods of a complete program consisting of a collection of classes that can involve mutually recursive fields and methods. The semantics of such a program is given as a fixpoint and the main result is proved by induction. This proof depends on a main lemma saying that commands are noninterfering—*safe*, for short—under the assumption that the methods they call are safe. Safety depends on confinement.

Section 5 is concerned with confinement. In [32], noninterference is proved on the basis of two lemmas called simple security and confinement. Simple security corresponds roughly to our L-confinement and safety Lemmas 5.1 and 6.1 for expressions. These results say that an L-expression cannot distinguish between related states. Confinement in [32] corresponds to what we call H-confinement, which pertains to H-commands: such commands do not assign to L-variables or fields. Noninterference for conditionals depends on H-confinement: A conditional with H-guard can behave differently in equivalent states, so the statement parts must be restricted to be H-commands.

To deal with the heap, we use a property called L-confinement. The language includes subclassing and subsumption, and we allow an L-class to have an H-subclass. As a result, simple typing does not prevent certain H to L flows. Moreover, an L-object can be aliased by both an L- and an H-variable. Hence, we must show that a typable L-expression never denotes an H-pointer and that commands preserve the invariant that L-variables and L-fields contain no H-pointers.

Besides controlling direct information flows, L-confinement is also needed to treat dynamic memory allocation and the equivalence \sim for heaps. An environment maps variable identifiers to values, whereas a heap maps locations to values. One need only define $\eta \sim \eta'$ for η and η' with the same domain. In a simple imperative language, an H-command can distinguish between \sim -related states but the distinction appears in terms of the states of variables (and termination). In a language with heap allocation, there is also the possibility of differing allocations. Thus it is unreasonable for $h \sim h'$ to require identical domains. But allocation must depend on the domain of the heap, to ensure freshness. If $h \sim h'$ allows different domains then an L-command could potentially distinguish by observing the behavior of the allocator.

In a language with pointer arithmetic, comparison using $<$ is a real problem. In Java, pointers are abstract: they can be compared only by $=$. Although two related heaps could result in different allocations, one can argue that no information is leaked to L-commands because they cannot make useful distinctions between pointer values. One could formalize this idea by requiring that, for the L part of related heaps, the domains need not be equal but rather there should be a bijection between them so that corresponding objects are related. From related states, allocation of a new L-object would add a pair to the bijection. Although this approach appears viable, we have avoided it for two reasons. First, the added complication would pervade all definitions and results. We prefer to follow the lead of Volpano and Smith in using simple standard semantic notions to the extent possible. Second, although it is easy to see how to maintain the bijection in a small step semantics, it is not as simple in a

compositional semantics. Some other possibilities for relating heaps are in [29, 30], though we have not pursued these ideas.

Our approach depends on the allocator satisfying a mild parametricity condition which is also needed for the abstraction theorem of [3]. The condition says that the choice of a fresh location for an object of class C depends only on currently-allocated C objects. Capability-based systems provide similar abstractions. The benefit here is that the definition of $h \sim h'$ can simply require that h and h' have the same domain for L-objects. We have not imposed a condition that the heaps are garbage-free. Garbage in the initial state has no influence on the final state. One might have thought that we would need to garbage-collect in order for the final states to be related, but this is not necessary.

3. Language: syntax and semantics

This section presents the language without security annotations; it is this language for which the semantics is defined. The language is the same as the core language in [3], from which we have also borrowed expository material. To make the paper more self-contained, we give complete typing rules and semantic definitions. For further explanation and justification of the definitions, the reader is advised to consult [3].

The grammar is based on given sets of class names (with typical element C), field names (f), method names (m), and variable/parameter names x (including *this*). Barred identifiers like \bar{T} indicate finite lists, e.g., $\bar{T} \bar{f}$ stands for a list \bar{f} of field names with corresponding types \bar{T} .

$$\begin{aligned}
T & ::= \text{bool} \mid \text{unit} \mid C \\
CL & ::= \text{class } C \text{ extends } C \{ \bar{T} \bar{f}; \bar{M} \} \\
M & ::= T m(\bar{T} \bar{x}) \{S; \text{return } e\} \\
S & ::= x := e \mid x.f := e \mid x := \text{new } C() \mid e.m(\bar{e}) \mid \\
& \quad \text{if } e S \text{ else } S \mid \text{var } T x := e \text{ in } S \mid S; S \\
e & ::= x \mid e.f \mid e.m(\bar{e}) \mid \\
& \quad e == e \mid (C) e \mid \text{null} \mid e \text{ instanceof } C
\end{aligned}$$

Without formalizing it precisely, we assume there is a class `Object` with no fields or methods which can be used as a superclass. Additional base types, such as integers, can be treated in the same way as `bool` and `unit`.

Table 1 gives the typing rules. A typing environment Γ is a finite function from variable names to types. A judgement of the form $\Gamma; C \vdash e : T$ says that e has type T in the context of a method of class C , with parameters and local variables declared by Γ . A judgement $\Gamma; C \vdash S : \text{com}$ says that S is a command in the same context. To simplify the typing rules and semantic definitions, we assume that variable and parameter names are not re-used.

A complete program is given as a *class table* CT that associates each declared class name with its declaration. The typing rules make use of auxiliary notions that are defined in terms of CT , so the typing relation \vdash depends on CT but this is elided in the notation. Because typing of each class is done in the context of the full table, methods can be mutually recursive, and so can field types.

Methods and classes are considered public. The rules for field access and update enforce visibility: fields of a class are accessible only to methods of that class, as in Java's `private` visibility.

Subsumption is built in to the rules using the subtyping relation \leq on T specified as follows. For base types, `bool` \leq `bool` and `unit` \leq `unit`. For classes C and D , we have $C \leq D$ iff either $C = D$ or the class declaration for C is `class C extends B { ... }` for some $B \leq D$.

To define some auxiliary notations, let

$$CT(C) = \text{class } C \text{ extends } D \{ \bar{T}_1 \bar{f}; \bar{M} \}$$

and let M be in the list \bar{M} of method declarations, with $M = T m(\bar{T}_2 \bar{x}) \{S; \text{return } e\}$. We record the typing information by defining $mtype(m, C) = (\bar{x} : \bar{T}_2) \rightarrow T$. For the declared fields, we define $dfields C = \bar{T}_1 \bar{f}$ and $type(\bar{f}, C) = \bar{T}_1$. To include inherited fields, we define $fields C = dfields C \cup fields D$, and assume \bar{f} is disjoint from the names in $fields D$. The built-in class `Object` has no methods or fields. Note that $mtype(m, C)$ is defined only if m is declared or inherited in C .

A class table is well formed if each of its method declarations is well formed according to the following rule.

$$\frac{
\begin{aligned}
& (\bar{x} : \bar{T}, \text{this} : C); C \vdash S : \text{com} \\
& (\bar{x} : \bar{T}, \text{this} : C); C \vdash e : U \quad U \leq T \\
& mtype(m, D) \text{ is undefined or equals } (\bar{x} : \bar{T}) \rightarrow T
\end{aligned}
}{
C \text{ extends } D \vdash T m(\bar{T} \bar{x}) \{S; \text{return } e\}
}$$

Turning to semantics, the state of a method in execution is comprised of a *heap* h , which is a finite partial function from locations to object states, and an *environment* η , which assigns locations and primitive values to local variables and parameters. Every environment of interest includes the distinguished variable *this* which points to the target object. A command denotes a function from initial state to either a final state or the error value \perp .

For locations, we assume that a countable set Loc is given, along with a distinguished entity `nil` not in Loc . We treat object states as mappings from field names to values. To track the object's class we assume given a function $loctype : Loc \rightarrow ClassNames$ such that for each C there are infinitely many locations ℓ with $loctype \ell = C$. We write $locs C$ for $\{\ell \mid loctype \ell = C\}$. The assumption about $loctype$ ensures an adequate supply of fresh locations, given that the domain of any heap is finite.

$\frac{\Gamma; C \vdash e_1 : T \quad \Gamma; C \vdash e_2 : T}{\Gamma; C \vdash e_1 == e_2 : \text{bool}}$	$\frac{Tf \in \text{dfields } C \quad \Gamma; C \vdash e : C}{\Gamma; C \vdash e.f : T}$	$\frac{\text{mtype}(m, D) = (\bar{x} : \bar{T}) \rightarrow T \quad \Gamma; C \vdash e : D \quad \Gamma; C \vdash \bar{e} : \bar{U} \quad \bar{U} \leq \bar{T}}{\Gamma; C \vdash e.m(\bar{e}) : T}$
$\frac{\Gamma; C \vdash e : D \quad B \leq D}{\Gamma; C \vdash (B) e : B}$	$\frac{\Gamma; C \vdash e : D \quad B \leq D}{\Gamma; C \vdash e \text{ instanceof } B : \text{bool}}$	$\Gamma; C \vdash \text{unit} : \text{unit}$
$\frac{x \neq \text{this} \quad \Gamma; C \vdash e : T \quad T \leq \Gamma x}{\Gamma; C \vdash x := e : \text{com}}$	$\frac{\Gamma x = C \quad Tf \in \text{dfields } C \quad \Gamma; C \vdash e : U \quad U \leq T}{\Gamma; C \vdash x.f := e : \text{com}}$	$\frac{x \neq \text{this} \quad B \leq \Gamma x}{\Gamma; C \vdash x := \text{new } B() : \text{com}}$
$\frac{\text{mtype}(m, D) = (\bar{x} : \bar{T}) \rightarrow T \quad \Gamma; C \vdash e : D \quad \Gamma; C \vdash \bar{e} : \bar{U} \quad \bar{U} \leq \bar{T}}{\Gamma; C \vdash e.m(\bar{e}) : \text{com}}$	$\frac{\Gamma; C \vdash e : \text{bool} \quad \Gamma; C \vdash S_1 : \text{com} \quad \Gamma; C \vdash S_2 : \text{com}}{\Gamma; C \vdash \text{if } e S_1 \text{ else } S_2 : \text{com}}$	
$\frac{\Gamma; C \vdash S_1 : \text{com} \quad \Gamma; C \vdash S_2 : \text{com}}{\Gamma; C \vdash S_1; S_2 : \text{com}}$		$\frac{\Gamma; C \vdash e : U \quad (\Gamma, x : T); C \vdash S : \text{com} \quad U \leq T}{\Gamma; C \vdash \text{var } T x := e \text{ in } S : \text{com}}$

Table 1. Typing rules for expressions and commands.

$\llbracket \text{bool} \rrbracket = \{\text{true}, \text{false}\}$	$\llbracket \text{unit} \rrbracket = \{\bullet\}$	$\llbracket C \rrbracket = \{\text{nil}\} \cup \{\ell \mid \ell \in \text{Loc} \wedge \text{loctype } \ell \leq C\}$
$\eta \in \llbracket \Gamma \rrbracket$	\Leftrightarrow	$\text{dom } \eta = \text{dom } \Gamma \wedge \forall x \in \text{dom } \eta. \eta x \in \llbracket \Gamma x \rrbracket$
$s \in \llbracket C \text{ state} \rrbracket$	\Leftrightarrow	$\text{dom } s = \text{fields } C \wedge \forall f \in \text{fields } C. sf \in \llbracket \text{type}(f, C) \rrbracket$
$h \in \llbracket \text{Heap} \rrbracket$	\Leftrightarrow	$\text{dom } h \subseteq_{\text{fin}} \text{Loc} \wedge \forall \ell \in \text{dom } h. h\ell \in \llbracket (\text{loctype } \ell) \text{ state} \rrbracket$
$\llbracket C, (\bar{x} : \bar{T}) \rightarrow T \rrbracket$	$=$	$\llbracket \bar{x} : \bar{T}, \text{this} : C \rrbracket \rightarrow \llbracket \text{Heap} \rrbracket \rightarrow (\llbracket T \rrbracket \times \llbracket \text{Heap} \rrbracket)_{\perp}$
$\llbracket MEnv \rrbracket$	\subseteq	$(C : \text{ClassNames}) \rightarrow (m : \text{MethodNames}) \rightarrow \llbracket C, \text{mtype}(m, C) \rrbracket$

Table 2. Semantic domains.

Methods are associated with classes, in a *method environment*, rather than with instances. For this reason the semantic domains, given in Table 2, are rather simple. There are no recursive domain equations to be solved. In addition to domains like $\llbracket T \rrbracket$ and $\llbracket \Gamma \rrbracket$ that correspond directly to syntactic notations, we use the following domains: $\llbracket \text{Heap} \rrbracket$ is the set of heaps, $\llbracket C \text{ state} \rrbracket$ is the set of states of objects of class C , $\llbracket MEnv \rrbracket$ is the set of method environments (we write \rightarrow for finite partial functions), and $\llbracket (C, (\bar{x} : \bar{T}) \rightarrow T) \rrbracket$ is the set of meanings for methods of class C with result T and parameters $\bar{x} : \bar{T}$.

The sets $\llbracket \text{Heap} \rrbracket$, $\llbracket \text{bool} \rrbracket$, $\llbracket C \rrbracket$, and $\llbracket C \text{ state} \rrbracket$ are ordered by equality. We write \rightarrow for continuous function space, ordered pointwise, and X_{\perp} for domain X with added bottom element \perp . Each set $\llbracket (C, (\bar{x} : \bar{T}) \rightarrow T) \rrbracket$ has a least

element (the constantly- \perp function) and least upper bounds of ascending chains, and this suffices for the fixpoint semantics. Without giving a precise formalization, we require that $\llbracket MEnv \rrbracket$ contain exactly the partial functions μ such that $\mu C m$ is defined for all classes C declared in CT and methods m declared or inherited in C .

The semantics is defined for an arbitrary allocator, but the noninterference theorem depends on parametricity.

Definition 1 (Allocator, parametric)

An *allocator* is a location-valued function fresh such that $\text{loctype}(\text{fresh}(C, h)) = C$ and $\text{fresh}(C, h) \notin \text{dom } h$, for all C, h . An allocator is *parametric* if $\text{dom } h_1 \cap \text{locs } C = \text{dom } h_2 \cap \text{locs } C$ implies $\text{fresh}(C, h_1) = \text{fresh}(C, h_2)$. \square

For example, if $\text{Loc} = \mathbb{N}$ the function $\text{fresh}(C, h) =$

$\min\{\ell \mid \text{loctype } \ell = C \wedge \ell \notin \text{dom } h\}$ is parametric.

It is straightforward to show that, as in Java, no program constructs create dangling pointers, but it is slightly simpler to formulate the definitions to allow dangling pointers. Like cast failures, dereferences of dangling pointers and nil are considered an error. We identify all errors, and divergence, with the improper value \perp .

The semantics is defined by induction on typing judgements, and for all typings for e and S we have

$$\begin{aligned} \llbracket \Gamma; C \vdash e : T \rrbracket &\in \llbracket MEnv \rrbracket \rightarrow \llbracket \Gamma \rrbracket \rightarrow \llbracket Heap \rrbracket \rightarrow \llbracket T \rrbracket_{\perp} \\ \llbracket \Gamma; C \vdash S : \text{com} \rrbracket &\in \\ &\llbracket MEnv \rrbracket \rightarrow \llbracket \Gamma \rrbracket \rightarrow \llbracket Heap \rrbracket \rightarrow (\llbracket \Gamma \rrbracket \times \llbracket Heap \rrbracket)_{\perp} \end{aligned}$$

To streamline the treatment of \perp in the semantic definitions, we use a metalanguage construct, **let** $d = E_1$ **in** E_2 , with the following meaning: If the value of E_1 is \perp then that is the value of the entire let expression; otherwise, its value is the value of E_2 with d bound to the value of E_1 . This construct is only exploited in the semantic definitions; later, in definitions of properties, we handle \perp explicitly.

Function update is written, e.g., $[\eta \mid x \mapsto d]$. In the semantics of local variables, we write \downarrow for domain restriction: if x is in the domain of function η then $\eta \downarrow x$ is the function like η but without x in its domain.

Table 3 gives the semantics of expressions and Table 4 gives the semantics of commands. The definitions are straightforward renderings of the operational semantics [2]. For example, the value of $e.f$ in state (η, h) is \perp if the value of e is \perp or is not in $\text{dom } h$; otherwise, the value of e is some location $\ell \in \text{dom } h$, so the object state $h\ell$ is a finite map with $f \in \text{dom } (h\ell)$ and the value of $e.f$ is $h\ell f$. Field update $x.f := e$ in state (η, h) does not change the environment; the new heap $[h \mid \ell \mapsto [h\ell \mid f \mapsto d]]$ updates h by replacing $h\ell$ with the object state $[h\ell \mid f \mapsto d]$ obtained by updating field f to have the value d of e .

For method call as an expression, $e.m(\bar{e})$, the value is \perp unless the value of e is some $\ell \in \text{dom } h$. In that case, let d be the method meaning given by μ for method m at the dynamic type (*loctype* ℓ) of e . The result of the call is obtained by applying d to the initial heap h and to the environment $[\bar{x} \mapsto \bar{d}, \text{this} \mapsto \ell]$ where \bar{d} is the list of values of arguments \bar{e} . The result, if not \perp , is a pair (d_0, h_0) ; for method call as expression, the value of $e.m(\bar{e})$ is d_0 . Note that h_0 is discarded; for expository simplicity we do not model side effects of expressions (see Section 7 for a discussion). For method call as command, d_0 is discarded and the new state is η, h_0 , as the call has no effect on the environment η of the caller.

The semantics of a class table is the method environment, $\hat{\mu}$, given as the least upper bound of the ascending

$$\begin{aligned} \llbracket \Gamma; C \vdash x : T \rrbracket_{\mu\eta h} &= \eta x \\ \llbracket \Gamma; C \vdash \text{null} : B \rrbracket_{\mu\eta h} &= \text{nil} \\ \llbracket \Gamma; C \vdash \text{unit} : \text{unit} \rrbracket_{\mu\eta h} &= \bullet \\ \llbracket \Gamma; C \vdash e_1 == e_2 : \text{bool} \rrbracket_{\mu\eta h} \\ &= \text{let } d_1 = \llbracket \Gamma; C \vdash e_1 : T \rrbracket_{\mu\eta h} \text{ in} \\ &\quad \text{let } d_2 = \llbracket \Gamma; C \vdash e_2 : T \rrbracket_{\mu\eta h} \text{ in } (d_1 = d_2) \\ \llbracket \Gamma; C \vdash e.f : T \rrbracket_{\mu\eta h} \\ &= \text{let } \ell = \llbracket \Gamma; C \vdash e : C \rrbracket_{\mu\eta h} \text{ in} \\ &\quad \text{if } \ell \notin \text{dom } h \text{ then } \perp \text{ else } h\ell f \\ \llbracket \Gamma; C \vdash e.m(\bar{e}) : T \rrbracket_{\mu\eta h} \\ &= \text{let } \ell = \llbracket \Gamma; C \vdash e : D \rrbracket_{\mu\eta h} \text{ in} \\ &\quad \text{if } \ell \notin \text{dom } h \text{ then } \perp \text{ else} \\ &\quad \text{let } (\bar{x} : \bar{T}) \rightarrow T = \text{mtype}(m, D) \text{ in} \\ &\quad \text{let } d = \mu(\text{loctype } \ell)m \text{ in} \\ &\quad \text{let } \bar{d} = \llbracket \Gamma; C \vdash \bar{e} : \bar{U} \rrbracket_{\mu\eta h} \text{ in} \\ &\quad \text{let } (d_0, h_0) = d[\bar{x} \mapsto \bar{d}, \text{this} \mapsto \ell]h \text{ in } d_0 \\ \llbracket \Gamma; C \vdash (B) e : B \rrbracket_{\mu\eta h} \\ &= \text{let } \ell = \llbracket \Gamma; C \vdash e : D \rrbracket_{\mu\eta h} \text{ in} \\ &\quad \text{if } \ell \in \text{dom } h \wedge \text{loctype } \ell \leq B \text{ then } \ell \text{ else } \perp \\ \llbracket \Gamma; C \vdash e \text{ instanceof } B : \text{bool} \rrbracket_{\mu\eta h} \\ &= \text{let } \ell = \llbracket \Gamma; C \vdash e : D \rrbracket_{\mu\eta h} \text{ in} \\ &\quad \ell \in \text{dom } h \wedge \text{loctype } \ell \leq B \end{aligned}$$

Table 3. Semantics of expressions.

chain $\mu \in \mathbb{N} \rightarrow \llbracket MEnv \rrbracket$ defined as follows.

$$\begin{aligned} \mu_0 C m &= \lambda\eta. \lambda h. \perp \\ \mu_{j+1} C m &= \llbracket M \rrbracket_{\mu_j} \quad \text{if } m \text{ is declared as } M \text{ in } C \\ \mu_{j+1} C m &= \mu_{j+1} B m \quad \text{if } m \text{ is inherited from } B \text{ in } C \\ &\llbracket M \rrbracket_{\mu\eta h} \\ &= \text{let } (\eta_0, h_0) = \llbracket (\bar{x} : \bar{T}, \text{this} : C); C \vdash S : \text{com} \rrbracket_{\mu\eta h} \text{ in} \\ &\quad \text{let } d = \llbracket (\bar{x} : \bar{T}, \text{this} : C); C \vdash e : T \rrbracket_{\mu\eta_0 h_0} \text{ in } (d, h_0) \\ &\quad \text{where in class } C \text{ we have } M = T m(\bar{T} \bar{x})\{S; \text{return } e\}. \end{aligned}$$

4. Security typing

In this section we annotate the syntax of Section 3 with security labels. Where a type T could occur, i.e., in declarations of fields, parameters, and local variables, we use pairs (T, κ) where κ is a security level. Such a pair, written τ , is called a *security type*. The security levels are L and H , ordered $L \leq H$. We write $\kappa_1 \sqcup \kappa_2$ to denote least upper

$$\begin{aligned}
& \llbracket \Gamma; C \vdash x := e : \text{com} \rrbracket \mu \eta h \\
& = \text{let } d = \llbracket \Gamma; C \vdash e : T \rrbracket \mu \eta h \text{ in } ([\eta \mid x \mapsto d], h) \\
& \llbracket \Gamma; C \vdash x.f := e : \text{com} \rrbracket \mu \eta h \\
& = \text{let } \ell = \eta x \text{ in if } \ell \notin \text{dom } h \text{ then } \perp \text{ else} \\
& \quad \text{let } d = \llbracket \Gamma; C \vdash e : U \rrbracket \mu \eta h \text{ in } (\eta, [h \mid \ell \mapsto [h\ell \mid f \mapsto d]]) \\
& \llbracket \Gamma; C \vdash x := \text{new } B() : \text{com} \rrbracket \mu \eta h \\
& = \text{let } \ell = \text{fresh}(B, h) \text{ in} \\
& \quad ([\eta \mid x \mapsto \ell], [h \mid \ell \mapsto [\text{fields } B \mapsto \text{defaults}]]) \\
& \llbracket \Gamma; C \vdash e.m(\bar{e}) : \text{com} \rrbracket \mu \eta h \\
& = \text{let } \ell = \llbracket \Gamma; C \vdash e : D \rrbracket \mu \eta h \text{ in} \\
& \quad \text{if } \ell \notin \text{dom } h \text{ then } \perp \text{ else} \\
& \quad \text{let } (\bar{x} : \bar{T}) \rightarrow T = \text{mtype}(m, D) \text{ in} \\
& \quad \text{let } d = \mu(\text{loctype } \ell)m \text{ in} \\
& \quad \text{let } \bar{d} = \llbracket \Gamma; C \vdash \bar{e} : \bar{U} \rrbracket \mu \eta h \text{ in} \\
& \quad \text{let } (d_0, h_0) = d[\bar{x} \mapsto \bar{d}, \text{this} \mapsto \ell]h \text{ in } (\eta, h_0) \\
& \llbracket \Gamma; C \vdash S_1; S_2 : \text{com} \rrbracket \mu \eta h \\
& = \text{let } (\eta_0, h_0) = \llbracket \Gamma; C \vdash S_1 : \text{com} \rrbracket \mu \eta h \text{ in} \\
& \quad \llbracket \Gamma; C \vdash S_2 : \text{com} \rrbracket \mu \eta_0 h_0 \\
& \llbracket \Gamma; C \vdash \text{if } e S_1 \text{ else } S_2 : \text{com} \rrbracket \mu \eta h \\
& = \text{let } b = \llbracket \Gamma; C \vdash e : \text{bool} \rrbracket \mu \eta h \text{ in} \\
& \quad \text{if } b \text{ then } \llbracket \Gamma; C \vdash S_1 : \text{com} \rrbracket \mu \eta h \\
& \quad \text{else } \llbracket \Gamma; C \vdash S_2 : \text{com} \rrbracket \mu \eta h \\
& \llbracket \Gamma; C \vdash \text{var } T x := e \text{ in } S : \text{com} \rrbracket \mu \eta h \\
& = \text{let } d = \llbracket \Gamma; C \vdash e : U \rrbracket \mu \eta h \text{ in} \\
& \quad \text{let } (\eta_0, h_0) = \llbracket (\Gamma, x : T); C \vdash S \rrbracket \mu [\eta \mid x \mapsto d]h \text{ in} \\
& \quad (\eta_0 \setminus x, h_0)
\end{aligned}$$

Table 4. Semantics of commands.

bound. The grammar is revised as follows.

$$\begin{aligned}
\kappa & ::= L \mid H \\
\tau & ::= (T, \kappa) \\
CL & ::= \text{class } C \kappa \text{ extends } C \{ \bar{\tau} \bar{f}; \bar{M} \} \\
M & ::= \tau m(\bar{\tau} \bar{x}) \kappa \{ S; \text{return } e \} \\
S & ::= \dots \mid \text{var } \tau x := e \text{ in } S \mid \dots
\end{aligned}$$

Tables 5 and 6 give typing rules for annotated programs. We write Δ for typing environments that assign security types. A judgement $\Delta; C \vdash S : (\text{com } \kappa_1, \kappa_2)$ says that S assigns only to variables (locals and parameters) of level $\geq \kappa_1$ and to object fields of level $\geq \kappa_2$ (see Lemma 5.4).

In Section 2 we noted that security typing can be presented more perspicuously using separate subtyping and

subsumption rules. For example, our rule for $x := e$ in Table 6 could give the command the type $(\text{com } \kappa_1, H)$, rather than including the unconstrained κ_4 and the constraint $\kappa_3 \leq \kappa_1$, as there is no heap effect and the most precise environment information is κ_1 . Even for a syntax-directed system, it may be possible to avoid subsumptions for primitive commands, given those built into rules for command constructs like sequencing. But care needs to be taken because levels in a method declaration cannot be changed in overriding declarations, and primitive commands can occur as method bodies. We err on the side of generality.

The rules use versions of the auxiliary functions *mtype* etc. that take security levels into account. Let

$$CT(C) = \text{class } C \kappa_1 \text{ extends } D \{ \bar{\tau}_1 \bar{f}; \bar{M} \}$$

and let M be in the list \bar{M} of method declarations, with

$$M = \tau_1 m(\bar{\tau}_2 \bar{x}) \kappa_2 \{ S; \text{return } e \}$$

The security version of *mtype* is defined by $\text{smtype}(m, C) = (\bar{x} : \bar{\tau}_2) \xrightarrow{\kappa_2} \tau_1$. Corresponding to *dfields*, *fields* and *type*, we define *sdfields*, *sfields* and *stype* which differ only in that they give security types, e.g., $\text{sdfields } C = \bar{\tau}_1 \bar{f}$.

We also need a function *level* that gives the level associated with the class itself: for the declaration above, $\text{level } C = \kappa_1$. Define $\text{level Object} = L$. For locations, define $\text{level } \ell = \text{level}(\text{loctype } \ell)$.

The rule for method declaration imposes the condition that an overriding definition cannot change the parameter or return types, nor the heap effect. The rule for class declaration restricts inheritance of methods. This is discussed in Section 4.1 below.

We use the symbol \dagger to erase annotations: $(T, \kappa)^\dagger = T$, and this extends to erasure for typing environments, commands, and method declarations in the obvious way. For example, if M is the method declaration

$$(\text{bool}, L) m((\text{bool}, L) x) L \{ x := x; \text{return } x \}$$

then M^\dagger is $\text{bool } m(\text{bool } x) \{ x := x; \text{return } x \}$.

For any judgement $\Delta; C \vdash S : (\text{com } \kappa_1, \kappa_2)$ derivable using the rules in Tables 5 and 6, the erased judgement $\Delta^\dagger; C \vdash S^\dagger : \text{com}$ is derivable using the rules of Table 1. Conversely, any program typable using the rules of Table 1 can be annotated everywhere by L and typed by the rules in Tables 5 and 6.

4.1. Examples

We consider a program in which the HIV status of a medical patient is to be kept confidential. A basic patient record looks as follows. We assume that there is an L-class *String*.

$\Delta; C \vdash x : \Delta x$	$\Delta; C \vdash \text{null} : (D, \kappa)$	$\Delta; C \vdash \text{unit} : (\text{unit}, \kappa)$
$\frac{\Delta; C \vdash e_1 : (T, \kappa_1) \quad \Delta; C \vdash e_2 : (T, \kappa_2)}{\Delta; C \vdash e_1 == e_2 : (\text{bool}, \kappa_1 \sqcup \kappa_2)}$	$\frac{(T, \kappa_1) f \in \text{sfields } C \quad \Delta; C \vdash e : (C, \kappa_2)}{\Delta; C \vdash e.f : (T, \kappa_1 \sqcup \kappa_2)}$	$\frac{\Delta; C \vdash e : (D, \kappa) \quad B \leq D}{\Delta; C \vdash (B) e : (B, \kappa)}$
$\frac{\text{smtpe}(m, D) = \bar{x} : (\bar{T}, \bar{\kappa}) \xrightarrow{\kappa_3} (T, \kappa_2) \quad \Delta; C \vdash e : (D, \kappa_4) \quad \Delta; C \vdash \bar{e} : (\bar{U}, \bar{\kappa}_5) \quad \bar{U} \leq \bar{T} \quad \bar{\kappa}_5 \leq \bar{\kappa}}{\Delta; C \vdash e.m(\bar{e}) : (T, \kappa_2 \sqcup \kappa_4)}$		$\frac{\Delta; C \vdash e : (D, \kappa) \quad B \leq D}{\Delta; C \vdash e \text{ instanceof } B : (\text{bool}, \kappa)}$

Table 5. Security typing rules for expressions.

```
class LPatient L extends Object {
  (String,L) name;
  (String,L) getName() H {
    return this.name
  }
  (unit,L) setName((String,L) n) L {
    this.name := n
  } }
```

Note that `setName` has L-effect, as it assigns to an L-field. Method `getName` can be declared as having L or H effect; we choose the latter. A confidential field is added for HIV status in the following class.

```
class XPatient L extends LPatient {
  (String,H) hiv;
  (unit,L) setHIV((String,H) s) H {
    this.hiv := s
  }
  (String,H) getHIV() H {
    return this.hiv
  } }
```

Here is an application using such objects, written using slight abuses of the official syntax.

```
class Main L extends Object {
  (unit,L) main() L {
    var (String,L) Lbuf := null;
    var (String,H) Hbuf := null;
    var (LPatient,L) lp:=readFile(...);
    var (XPatient,L) xp:=new XPatient();
    Lbuf := lp.getName();
    Hbuf := lp.getName();
    (1) xp.setName(Lbuf);
    Hbuf := readFromTrustedChan(...);
    (2) xp.setHIV(Hbuf); } }
```

After execution of line (1), `Lbuf`, `Hbuf`, and `xp.name` all reference the same object. After execution of line (2), `Lbuf` aliases `xp.name`, and `Hbuf` aliases `xp.hiv`. All the code above is allowed by the security typing rules, assuming appropriate types for the read methods. The following are not allowed in `main`.

```
(3) Lbuf := Hbuf;
(4) lp.setName(xp.getHiv());
```

In both (3) and (4), there is a direct flow from H to L. Let us consider an H-subclass of `XPatient`.

```
class HPatient H extends XPatient {
  (String,H) f;
  (unit,L) setf((String,H) v) H {
    this.f := v
  }
  (String,H) getf() H {
    return this.f
  } }
```

By contrast with `setName`, the effect of `setf` can be declared H as it assigns only to an H-field.

There is, however, a problem with inherited methods. Suppose `LPatient` declared a method

```
(String,L) passSelf() H {
  ...o.m(this)... }
```

where `o.m` is a call on some object `o` of a method `m` with L parameter. This is typable in `LPatient`, although it would not be typable in `HPatient` where the level of `this` is H. For a variable `hp` of class `HPatient`, an invocation `hp.passSelf()` of the inherited method would result in a flow from H to L in the call `o.m(this)`.

We solve the problem in a simple way: The rule for classes in Table 6 requires that an H-subclass of an L-class overrides all methods of its superclass. This is unnecessarily, and unacceptably, restrictive. It requires, for example, that `getName` be overridden although it poses no risk of bad flow; and an overriding declaration does not have access to the private field `name`. The problem can be solved in a less restrictive way using the notion of “anonymous method” of Vitek and Bokowski [31]: such methods do not leak the receiver *this*. The static analyses described in [31, 14] restrict inheritance only for methods that leak *this*. We expect that their constraints can be adapted easily to our setting.

$\frac{x \neq \text{this} \quad T_2 \leq T_1 \quad \kappa_2 \leq \kappa_1 \quad \kappa_3 \leq \kappa_1}{\Delta, x : (T_1, \kappa_1); C \vdash e : (T_2, \kappa_2)}$	$\frac{(T, \kappa_2) f \in \text{sdfields } C \quad \Delta, x : (C, \kappa_1); C \vdash e : (U, \kappa_3) \quad U \leq T \quad \kappa_1 \sqcup \kappa_3 \sqcup \kappa_5 \leq \kappa_2}{\Delta, x : (C, \kappa_1); C \vdash x.f := e : (\text{com } \kappa_4, \kappa_5)}$
$\frac{x \neq \text{this} \quad B \leq D \quad \text{level } B \sqcup \kappa_2 \leq \kappa_1 \quad \kappa_3 \leq \text{level } B}{\Delta, x : (D, \kappa_1); C \vdash x := \text{new } B() : (\text{com } \kappa_2, \kappa_3)}$	
$\frac{\text{smtype}(m, D) = \bar{x} : (\bar{T}, \bar{\kappa}) \xrightarrow{\kappa_3} (T, \kappa_2) \quad \Delta; C \vdash e : (D, \kappa_4) \quad \Delta; C \vdash \bar{e} : (\bar{U}, \bar{\kappa}_5) \quad \bar{U} \leq \bar{T} \quad \bar{\kappa}_5 \leq \bar{\kappa} \quad \kappa_4 \sqcup \kappa_7 \leq \kappa_3}{\Delta; C \vdash e.m(\bar{e}) : (\text{com } \kappa_6, \kappa_7)}$	$\frac{\Delta; C \vdash S_1 : (\text{com } \kappa_1, \kappa_2) \quad \Delta; C \vdash S_2 : (\text{com } \kappa_3, \kappa_4) \quad \kappa_5 \leq \kappa_1 \sqcap \kappa_3 \quad \kappa_6 \leq \kappa_2 \sqcap \kappa_4}{\Delta; C \vdash S_1; S_2 : (\text{com } \kappa_5, \kappa_6)}$
$\frac{\Delta; C \vdash e : (\text{bool}, \kappa_5) \quad \Delta; C \vdash S_1 : (\text{com } \kappa_1, \kappa_3) \quad \Delta; C \vdash S_2 : (\text{com } \kappa_2, \kappa_4) \quad \kappa_5 \leq \kappa_6 \sqcap \kappa_7 \quad \kappa_6 \leq \kappa_1 \sqcap \kappa_2 \quad \kappa_7 \leq \kappa_3 \sqcap \kappa_4}{\Delta; C \vdash \text{if } e \text{ } S_1 \text{ else } S_2 : (\text{com } \kappa_6, \kappa_7)}$	$\frac{\Delta; C \vdash e : (U, \kappa_4) \quad \Delta, x : (T, \kappa_1); C \vdash S : (\text{com } \kappa_5, \kappa_6) \quad U \leq T \quad \kappa_4 \leq \kappa_1 \quad \kappa_2 \leq \kappa_5 \quad \kappa_3 \leq \kappa_6}{\Delta; C \vdash \text{var } (T, \kappa_1) \text{ } x := e \text{ in } S : (\text{com } \kappa_2, \kappa_3)}$
$\frac{\bar{x} : (\bar{T}, \bar{\kappa}), \text{this} : (C, \kappa_1); C \vdash S : (\text{com } \kappa_2, \kappa_3) \quad \bar{x} : (\bar{T}, \bar{\kappa}), \text{this} : (C, \kappa_1); C \vdash e : (T, \kappa_4) \quad \text{smtype}(m, D) \text{ is undefined or equals } \bar{x} : (\bar{T}, \bar{\kappa}) \xrightarrow{\kappa_3} (T, \kappa_4)}{C \kappa_1 \text{ extends } D \vdash (T, \kappa_4) \text{ } m((\bar{T}, \bar{\kappa}) \bar{x}) \kappa_3 \{S; \text{return } e\}}$	$\frac{\text{level } D \leq \kappa \quad C \kappa \text{ extends } D \vdash M \text{ for each } M \in \bar{M} \quad \text{If level } D \neq \kappa \text{ then every } m \text{ with } \text{smtype}(m, D) \text{ defined is overridden in } C \text{ by some } M \in \bar{M}.}{\vdash \text{class } C \kappa \text{ extends } D \{ \bar{\tau} \bar{f}; \bar{M} \}}$

Table 6. Security typing rules for commands, method declarations, and classes.

Our last examples involve information leaks via the conditional control flow implicit in dynamically bound method calls. Consider these three classes.

```
class YN L extends Object {
  (bool,L) val(){ return true; } }
class Y L extends YN {
  (bool,L) val(){ return true; } }
class N L extends YN {
  (bool,L) val(){ return false; } }
```

We add the following method to XPatient. The typing rules force the result to have level H; otherwise there would be a bad data flow.

```
(YN,H) leak() H {
  var (YN,H) o;
  if (this.hiv) o := new Y();
  else          o := new N();
  return o; }
```

In main, the expression `xp.leak().val()` has level H due to the return type of `leak`.

Here is a similar example but using the heap and method call commands rather than expressions.

```
class YNh L extends Object {
  (bool,H) v;
  (bool,H) val(){ return v } }
  (unit,L) setv((bool,H) w) H {
    this.v:=w }
  (unit,L) set() H {this.setv(true)} }
class Yh L extends YNh {
  (unit,L) set() H {this.setv(true)} }
class Nh L extends YNh {
  (unit,L) set() H {this.setv(false)} }
```

The leak method is the same as before but using YNh and its subclasses. Consider `x` of type `(YNh, H)` in

```
x:=xp.leak(); x.set(); ...x.val()...
```

The declared level of `x` must be H because it is assigned from `leak`. The method call rule then requires for `x.set` that the heap effect of `set` be H, which in turn forces the level of field `v` to be H. Indeed, if its level was L then the call `x.set()` would violate noninterference.

4.2. Remarks about proofs

Proofs in the sequel involve detailed analysis of the semantics and the security typing rules. For each specific case, the semantic definition may involve several values (e.g., the value of e is needed in the semantics of $x := e$), and the rule may involve several types and security labels. In writing a given proof case, we found it convenient to write down both the rule and the semantics for reference. It is impractical to include such redundancy in the paper, however. Instead, when it comes to proving something about a particular construct we make free use of identifiers in the typing rule (in Table 5 or 6), for types and labels, and identifiers in the semantic definition for semantic values (in Table 3 or 4). We explicitly introduce identifiers for types or values only when necessity or perspicuity demands it.

Note that the semantic definition may use different identifiers for types, as the semantics is based on the typing rules in Table 1 rather than the security rules in Tables 5 and 6.

We streamline the proofs by ignoring \perp outcomes in many cases. Most of the results only pertain to non- \perp outcomes, and the constructs are mostly strict in \perp . Without comment we assume various intermediate values are non- \perp unless confusion could result.

5. Confinement

This section shows that typable programs maintain the invariant that L fields and variables never hold H locations. The formalization uses the indistinguishability relation \sim also used in the main results of Section 6.

In formalizing the absence of L -variables that refer to H -objects, we take advantage of the fact that $\text{nil} \notin \text{Loc}$ and $\perp \notin \text{Loc}$. We use the short name “ ok ” for L -confinement.

Definition 2 (L -confinement (ok))

- Define $LLoc = \{\ell \in \text{Loc} \mid \text{level } \ell = L\}$.
- For heaps, define $ok\ h$ iff for all $\ell \in \text{dom } h$ and every $f \in \text{fields}(\text{loctype } \ell)$, if $\text{styp}(f, \text{loctype } \ell) = (T, L)$ for some T and $h\ell f \in \text{Loc}$ then $h\ell f \in LLoc$.
- For environments, define $ok\ \Delta\ \eta$ iff for every x with $\Delta x = (T, L)$ for some T , if $\eta x \in \text{Loc}$ then $\eta x \in LLoc$.
- For method environments, define $ok\ \mu$ iff the following holds: for every m, C, η, h , if $ok\ h$, $ok\ \Delta\ \eta$, and $\mu C m \eta h \neq \perp$ then $ok\ h_0$ and $\kappa_3 = L \wedge d \in \text{Loc} \Rightarrow d \in LLoc$,
where $\text{smtyp}(m, C) = \bar{x} : (\bar{T}, \bar{\kappa}) \xrightarrow{\kappa_2} (T, \kappa_3)$
 $\Delta = \bar{x} : (\bar{T}, \bar{\kappa}), \text{this} : (C, \text{level } C)$
 $(d, h_0) = \mu C m \eta h$

Lemma 5.1 (L -confinement of expressions)

Let $\Delta; C \vdash e : (T, L)$ and let $d = \llbracket \Delta^\dagger; C \vdash e : T \rrbracket_{\mu\eta h}$. If $ok\ \mu$, $ok\ \Delta\ \eta$, and $ok\ h$ then $d \in \text{Loc} \Rightarrow d \in LLoc$.

Proof: By induction on the derivation of $\Delta; C \vdash e : (T, L)$ (for brevity: “induction on e ”). Recall from Section 4.2 that throughout the proofs we ignore the \perp cases.

- Case of x : Then $d = \llbracket \Delta^\dagger; C \vdash x : T \rrbracket_{\mu\eta h} = \eta x$. The result follows directly from assumption $ok\ \Delta\ \eta$.
- $e.f$: By typing, $\kappa_1 = L = \kappa_2$. Because $\kappa_2 = L$ we can use induction on e ; this, together with the assumption that the semantics is non- \perp , yields that there is $\ell = \llbracket \Delta^\dagger; C \vdash e : C \rrbracket_{\mu\eta h}$ and $\ell \in LLoc$ and $\ell \in \text{dom } h$. Now the result follows using $\kappa_1 = L$ and assumption $ok\ h$.
- $e.m(\bar{e})$: Let $\Delta_0 = [\bar{x} : (\bar{T}, \bar{\kappa}), \text{this} : (\text{loctype } \ell, \text{level } \ell)]$, and $\eta_0 = [\bar{x} \mapsto \bar{d}, \text{this} \mapsto \ell]$. We claim that $ok\ \Delta_0\ \eta_0$; then we get the result by $ok\ \mu$ because by typing $\kappa_2 = \kappa_4 = L$. It remains to prove the claim.

If $\bar{\kappa} = L$, then by the typing rule, $\bar{\kappa}_5 = L$. So by induction on \bar{e} , and since the semantics is non- \perp , we get $\bar{d} \in \text{Loc} \Rightarrow \bar{d} \in LLoc$. Hence $ok\ (\bar{x} : (\bar{T}, \bar{\kappa})) [\bar{x} \mapsto \bar{d}]$. We get $ok\ (\text{this} : (\text{loctype } \ell, \text{level } \ell)) [\text{this} \mapsto \ell]$ directly from the definitions of ok and $LLoc$, as $\text{level } \ell = L$ iff $\ell \in LLoc$.

- $(B)\ e$: By $\kappa = L$, and because we are considering the case where $\llbracket (B)\ e \rrbracket_{\mu\eta h} \neq \perp$, we can use induction on e to obtain $\ell \in \text{Loc} \Rightarrow \ell \in LLoc$. Moreover, ℓ must be in $\text{dom } h$ otherwise the semantics is \perp . Now the result follows directly.
- e instance of $B, e == e', \text{null}, \text{unit}$: The result returned in each case is not in Loc . This falsifies the antecedent in the lemma. \square

Lemma 5.2 (L -confinement of commands)

Let $\Delta; C \vdash S : (\text{com } \kappa_1, \kappa_2)$. If $ok\ \mu$, $ok\ h$, $ok\ \Delta\ \eta$, and $\llbracket \Delta^\dagger; C \vdash S^\dagger : \text{com} \rrbracket_{\mu\eta h} \neq \perp$ then $ok\ \Delta\ \eta_0$ and $ok\ h_0$, where $(\eta_0, h_0) = \llbracket \Delta^\dagger; C \vdash S^\dagger : \text{com} \rrbracket_{\mu\eta h}$.

Proof: By induction on the derivation of $\Delta; C \vdash S : (\text{com } \kappa_1, \kappa_2)$, using the assumptions of the Lemma.

- $x := e$: This has no heap effect: the result heap h_0 is h and $ok\ h$ holds by assumption. We only need to show $ok\ (\Delta, x : (T_1, \kappa_1)) [\eta \mid x \mapsto d]$, where $d = \llbracket \Delta^\dagger; C \vdash e : T_2 \rrbracket_{\mu\eta h}$. Accordingly, assume that $\kappa_1 = L$. Then by typing, $\kappa_2 = L$ so Lemma 5.1 for e yields $d \in \text{Loc} \Rightarrow d \in LLoc$.
- $x.f := e$: The result environment η_0 is just η , and $ok\ \Delta\ \eta$ holds by assumption. We only need to show $ok\ h_0$,

where $h_0 = [h \mid \ell \mapsto [h\ell \mid f \mapsto d]]$, $\ell = \eta x$ and $d = \llbracket \Delta^\dagger; C \vdash e : U \rrbracket \mu \eta h$. Since the semantics is non- \perp , we have $\ell \in \text{dom } h$. Now assume that the level κ_2 of field f is L . Then by the typing rule we have $\kappa_1 = \kappa_3 = L$. Now applying Lemma 5.1 on e , we obtain, $d \in \text{Loc} \Rightarrow d \in \text{LLoc}$. This proves $ok h_0$.

- $x := \text{new } B(\)$: For the environment, we must show $ok(\Delta, x : (D_1, \kappa_1))[\eta \mid x \mapsto \ell]$, where $\ell = \text{fresh}(B, h)$. Since by assumption $ok \Delta \eta$, it is enough to deal with x ; that is, if $\kappa_1 = L$ we must show $\ell \in \text{Loc} \Rightarrow \ell \in \text{LLoc}$. Indeed, the allocator fresh yields $\ell \in \text{Loc}$. But by the typing rule, $\kappa_1 = L$ implies $\text{level } B = L$, hence $\ell \in \text{LLoc}$ as required. Finally, we get $ok h_0$, where $h_0 = [h \mid \ell \mapsto [\text{fields } B \mapsto \text{defaults}]]$, because $ok h$ by hypothesis and defaults contains no locations (the defaults are false and nil).
- $e.m(\bar{e})$: We have $\eta_0 = \eta$, and $ok \Delta \eta$ by assumption, so it suffices to show $ok h_0$. Let $\Delta_0 = \bar{x} : (\bar{T}, \bar{\kappa}), \text{this} : (\text{loctype } \ell, \text{level } \ell)$, and $\eta_0 = [\bar{x} \mapsto \bar{d}, \text{this} \mapsto \ell]$. We claim $ok \Delta_0 \eta_0$; then we get the result by $ok \mu$. The claim is proved by the same argument as for method calls in the proof of Lemma 5.1 (literally the same argument, owing to the fact that the relevant identifiers are the same in the typing rules for method call as expression and as command).
- $\text{if } e S_1 \text{ else } S_2$: Let $b = \llbracket \Delta^\dagger; C \vdash e : \text{bool} \rrbracket \mu \eta h$. Then if $b = \text{true}$, the result follows by induction on S_1 and if $b = \text{false}$, the result follows by induction on S_2 .
- $\text{var } (T, \kappa_1) x := e \text{ in } S$: First, we have $ok(\Delta, x : (T, \kappa_1))[\eta \mid x \mapsto d]$ where $d = \llbracket \Delta^\dagger; C \vdash e : U \rrbracket \mu \eta h$. This is because if $\kappa_1 = L$, then by typing $\kappa_4 = L$, so by Lemma 5.1 for e , $d \in \text{Loc} \Rightarrow d \in \text{LLoc}$. Induction on S yields $ok(\Delta, x : (T, \kappa_1)) \eta_1$ and $ok h_0$, where $(\eta_1, h_0) = \llbracket (\Gamma, x : T); C \vdash S \rrbracket \mu[\eta \mid x \mapsto d]h$. Hence $ok \Delta(\eta_1 \mid x)$.
- $S_1; S_2$: Use induction on S_1 , then on S_2 . \square

Lemma 5.3 (L -confinement of method environments)

For each i we have $ok \mu_i$, and $ok \hat{\mu}$. \square

The proof is by induction on i , using Lemmas 5.1 and 5.2, and then fixpoint induction for $\hat{\mu}$. It follows the pattern of the proof of Theorem 6.3, and is given in the full paper.

Object states are indistinguishable by L if their L -fields are equal, and environments are indistinguishable if their L -variables are equal. In the case of heaps and object states, the relevant levels are determined by the field declarations in the class table. By contrast, the levels for environments

are determined by parameter and local variable declarations, hence the dependence is explicit in the notation \sim_Δ . It is straightforward to show that each of these is an equivalence relation.

Definition 3 (Indistinguishable by L)

- For $s, s' \in \llbracket C \text{ state} \rrbracket$, define $s \sim s'$ iff $\forall f \in \text{fields } C . \text{let } (T, \kappa) = \text{stype}(f, C) \text{ in } (\kappa = L \Rightarrow sf = s'f)$.
- For $h, h' \in \llbracket \text{Heap} \rrbracket$, define $h \sim h'$ iff $\text{dom } h \cap \text{LLoc} = \text{dom } h' \cap \text{LLoc}$ and $\forall \ell \in \text{dom } h \cap \text{LLoc} . h\ell \sim h'\ell$.
- For $\eta, \eta' \in \llbracket \Delta^\dagger \rrbracket$, define $\eta \sim_\Delta \eta'$ iff $\forall x \in \text{dom } \Delta . \text{let } (T, \kappa) = \Delta x \text{ in } (\kappa = L \Rightarrow \eta x = \eta' x)$. \square

If a command is typable as $(\text{com } H, \kappa)$ it does not assign to L -variables, and if it is typable as $(\text{com } \kappa_2, H)$ it does not assign to L -fields of objects.

Definition 4 (H -confined method environment)

Method environment μ is H -confined, written $H\text{conf } \mu$, if $\mu C m \eta h \neq \perp \Rightarrow h_0 \sim h$, where $(d, h_0) = \mu C m \eta h$, for all C, m with $\text{smype}(m, C) = \bar{x} : (\bar{T}, \bar{\kappa}) \xrightarrow{H} (T, \kappa)$. \square

Lemma 5.4 (H -confinement of commands)

Let $\Delta; C \vdash S : (\text{com } \kappa_1, \kappa_2)$. Then for all μ, η, h such that $H\text{conf } \mu$ and $\llbracket \Delta^\dagger; C \vdash S^\dagger : \text{com} \rrbracket \mu \eta h \neq \perp$ we have

- if $\kappa_1 = H$ and $(\eta_0, h_0) = \llbracket \Delta^\dagger; C \vdash S^\dagger : \text{com} \rrbracket \mu \eta h$ then $\eta \sim_\Delta \eta_0$.
- if $\kappa_2 = H$ and $(\eta_0, h_0) = \llbracket \Delta^\dagger; C \vdash S^\dagger : \text{com} \rrbracket \mu \eta h$ then $h \sim h_0$.

Proof: By induction on the derivation of $\Delta; C \vdash S : (\text{com } \kappa_1, \kappa_2)$. As usual, we follow the conventions described in Section 4.2; in particular, level identifiers in the proof are those in the relevant rules, *not* κ_1, κ_2 as used in the statement of the Lemma.

- $x := e$: This has no effect on the heap. We need to show that $\kappa_3 = H \Rightarrow \eta \sim_{(\Delta, x : (T_1, \kappa_1))} \eta_0$, where $\eta_0 = [\eta \mid x \mapsto d]$ and $d = \llbracket \Delta^\dagger; C \vdash e : T \rrbracket \mu \eta h$. Assuming $\kappa_3 = H$, we obtain $\kappa_1 = H$ by the typing rule. Now the result follows using definition $\sim_{(\Delta, x : (T_1, H))}$.
- $x.f := e$: This has no effect on the environment, so it suffices to show that $\kappa_5 = H \Rightarrow h \sim h_0$, where $h_0 = [h \mid \ell \mapsto [h\ell \mid f \mapsto d]]$, $\ell = \eta x$ and $d = \llbracket \Delta^\dagger; C \vdash e : U \rrbracket \mu \eta h$. If $\ell \notin \text{LLoc}$ then $h \sim h_0$ because the two heaps are identical on LLoc . If $\ell \in \text{LLoc}$ then we must consider the updated field $h\ell f$. Assuming $\kappa_5 = H$, the typing rule forces the level κ_2 of field f to be H ; nothing else is updated, so $h \sim h_0$ holds by definition of \sim .

- $x := \text{new } B(\)$: First, assume $\kappa_2 = H$. Then by typing, $\kappa_1 = H$. We must show $\eta \sim_{(\Delta, x: (D, \kappa_1))} \eta_0$ where $\eta_0 = [\eta \mid x \mapsto \ell]$ and $\ell = \text{fresh}(B, h)$. But this follows by definition of $\sim_{(\Delta, x: (D, \kappa_1))}$ since $\kappa_1 = H$.

Next assume $\kappa_3 = H$. Then by typing, $\text{level } B = H$. Hence $\ell \notin \text{LLoc}$. We must show $h \sim h_0$, where $h_0 = [h \mid \ell \mapsto [\text{fields } B \mapsto \text{defaults}]]$. But this follows by definition of \sim because h and h_0 are identical on LLoc .

- $e.m(\bar{e})$: This has no effect on the environment. For the heap, suppose $\kappa_7 = H$. Then we must show $h \sim h_0$, where $(d_0, h_0) = \mu(\text{loctype } \ell)m[\bar{x} \mapsto \bar{d}, \text{this} \mapsto \ell]h$, and $\ell = \llbracket \Delta; C \vdash e : D \rrbracket \mu\eta h$ and $\bar{d} = \llbracket \Delta; C \vdash \bar{e} : \bar{U} \rrbracket \mu\eta h$. Because $\kappa_7 = H$, we have by the typing rule, $\kappa_3 = H$. Moreover, $\text{smtype}(m, (\text{loctype } \ell)) = \bar{x} : (\bar{T}, \bar{\kappa}) \xrightarrow{H} (T, \kappa_2)$. Hence, by assumption $H\text{conf } \mu$, we get $h \sim h_0$.
- $S_1; S_2$: First assume $\kappa_5 = H$. Then by the typing rule, $\kappa_1 = H = \kappa_3$. By induction on S_1 we get $\eta \sim_{\Delta} \eta_1$ where we write (η_1, h_1) for the intermediate state. Then by induction on S_2 , we get $\eta_1 \sim_{\Delta} \eta_0$ where $(\eta_0, h_0) = \llbracket \Delta^\dagger; C \vdash S_2 : \text{com} \rrbracket \mu\eta_1 h_1$. Hence $\eta \sim_{\Delta} \eta_0$ by transitivity.

Next, assume $\kappa_6 = H$. Then $\kappa_2 = H = \kappa_4$. And, $h \sim h_0$ follows by induction on S_1 and S_2 and transitivity.

- $\text{if } e \ S_1 \ \text{else } S_2$: First, assume $\kappa_6 = H$. Then by the typing rule, $\kappa_1 = H$ and $\kappa_2 = H$. Let $b = \llbracket \Delta^\dagger; C \vdash e : \text{bool} \rrbracket \mu\eta h$. Then if $b = \text{true}$, the result follows by induction on S_1 and if $b = \text{false}$, the result follows by induction on S_2 .

Next, assume $\kappa_7 = H$. Then by the typing rule, $\kappa_3 = H$ and $\kappa_4 = H$. Again, the result follows by induction on S_1 if $b = \text{true}$ and by induction on S_2 if $b = \text{false}$.

- $\text{var } (T, \kappa_1) \ x := e \ \text{in } S$: First, assume $\kappa_2 = H$. Then by the typing rule, $\kappa_5 = H$. Hence by induction on S , $[\eta \mid x \mapsto d] \sim_{(\Delta, x: (T, \kappa_1))} \eta_0$ where $(\eta_0, h_0) = \llbracket (\Delta^\dagger, x : T); C \vdash S : \text{com} \rrbracket \mu[\eta \mid x \mapsto d]h$ and $d = \llbracket \Delta^\dagger; C \vdash e : U \rrbracket \mu\eta h$. Hence $\eta \sim_{\Delta} (\eta_0 \upharpoonright x)$.

Next, assume $\kappa_3 = H$. Then by the typing rule, $\kappa_6 = H$. And, by induction on S we get $h \sim h_0$. \square

Note that a command $\text{var } (T, L) \ x := e \ \text{in } S$ can be typed as $\Delta; C \vdash S : (\text{com } H, \kappa_2)$ so the Lemma applies to such commands. But the typing rule ensures that no L -variable is assigned in S . Moreover, if $\kappa_2 = H$ then no L -fields are assigned. So x is of limited use.

Lemma 5.5 (H -confinement of method environments)

For each i we have $H\text{conf } \mu_i$, and $H\text{conf } \hat{\mu}$. \square

The proof is by induction on i , using Lemma 5.4, and then fixpoint induction. It follows the pattern of the proof of Theorem 6.3, and is given in the full paper.

6. Noninterference

A method meaning is safe, i.e., noninterfering, provided that, for terminating computations, L -indistinguishable initial heaps and environments lead to L -indistinguishable results.

Definition 5 (Safe method environment)

We define *safe* μ iff for all C, m and all h, h', η, η' the following holds: If $\text{ok } h, \text{ok } h', \text{ok } \Delta \eta$, and $\text{ok } \Delta \eta'$ then

$$h \sim h' \wedge \eta \sim_{\Delta} \eta' \wedge \mu C m \eta h \neq \perp \neq \mu C m \eta' h' \\ \Rightarrow h_0 \sim h'_0 \wedge (\kappa_3 = L \Rightarrow d = d')$$

$$\text{where } \text{smtype}(m, C) = \bar{x} : (\bar{T}, \bar{\kappa}) \xrightarrow{\kappa_2} (T, \kappa_3) \\ \Delta = \bar{x} : (\bar{T}, \bar{\kappa}), \text{this} : (C, \text{level } C) \\ (d, h_0) = \mu C m \eta h \\ (d', h'_0) = \mu C m \eta' h'$$

\square

Our main result is that the method environment denoted by a secure class table is safe. The proof uses the following two results which express noninterference for the expression and command constructs, respectively.

If an expression can be typed $\Delta; C \vdash e : (T, L)$ then its meaning is the same in two L -indistinguishable states, provided that it diverges in neither state.

Lemma 6.1 (Safe expressions)

Suppose $\Delta; C \vdash e : (T, L)$. Suppose that $h \sim h', \eta \sim_{\Delta} \eta', \text{safe } \mu, \text{ok } \mu, \text{ok } h, \text{ok } h', \text{ok } \Delta \eta, \text{ok } \Delta \eta'$, and $\llbracket \Delta^\dagger; C \vdash e : T \rrbracket \mu\eta h \neq \perp \neq \llbracket \Delta^\dagger; C \vdash e : T \rrbracket \mu\eta' h'$. Then $\llbracket \Delta^\dagger; C \vdash e : T \rrbracket \mu\eta h = \llbracket \Delta^\dagger; C \vdash e : T \rrbracket \mu\eta' h'$.

In this proof and subsequent ones, we extend the convention described in Section 4.2. When comparing semantics for a pair of states (η, h) and (η', h') , we use corresponding primes on identifiers in the semantic definitions. For example, the semantic definition of $\llbracket \Delta^\dagger; C \vdash x := e : \text{com} \rrbracket \mu\eta h$ involves value d denoted by e in state η, h , so we write d' for the corresponding value for $\llbracket \Delta^\dagger; C \vdash x := e : \text{com} \rrbracket \mu\eta' h'$.

Proof: By induction on the derivation of $\Delta; C \vdash e : (T, L)$. Using the assumptions of the Lemma, we show $\llbracket \Delta^\dagger; C \vdash e : T \rrbracket \mu\eta h = \llbracket \Delta^\dagger; C \vdash e : T \rrbracket \mu\eta' h'$ by cases on e .

- x : Then $\llbracket \Delta^\dagger; C \vdash x : T \rrbracket \mu\eta h = \eta x$ and $\llbracket \Delta^\dagger; C \vdash x : T \rrbracket \mu\eta' h' = \eta' x$. By assumption $\eta \sim_{\Delta} \eta'$ and $\Delta x = (T, L)$ we have $\eta x = \eta' x$.
- $e_1 == e_2$: By typing, $\kappa_1 = L = \kappa_2$. Thus we can use induction on e_1 to obtain $\llbracket \Delta^\dagger; C \vdash e_1 : T \rrbracket \mu\eta h = \llbracket \Delta^\dagger; C \vdash e_1 : T \rrbracket \mu\eta' h'$ (we are considering the case where $\llbracket e_1 == e_2 \rrbracket \mu\eta h \neq \perp$ and $\llbracket e_1 == e_2 \rrbracket \mu\eta' h' \neq \perp$, so by semantics the value of e_1 is also non- \perp). Similarly, induction on e_2 yields $\llbracket \Delta^\dagger; C \vdash e_2 : T \rrbracket \mu\eta h = \llbracket \Delta^\dagger; C \vdash e_2 : T \rrbracket \mu\eta' h'$. The result follows directly.

- $e.f$: By typing, $\kappa_1 = L = \kappa_2$. Because $\kappa_2 = L$ we can use induction on e ; this yields that there is ℓ with $\llbracket \Delta^\dagger; C \vdash e : C \rrbracket \mu \eta h = \ell = \llbracket \Delta^\dagger; C \vdash e : C \rrbracket \mu \eta' h'$, as we only consider the case that both semantics are non- \perp . For the same reason, ℓ is in the domain of both h and h' . By $\kappa_2 = L$ and Lemma 5.1 we have $\ell \in LLoc$ so, by assumption $h \sim h'$, we get $h\ell \sim h'\ell$; this implies $h\ell f = h'\ell f$ because field f has label $\kappa_1 = L$.
- $e.m(\bar{e})$: By the security typing rule we have $e : (D, L)$, so by induction $\ell = \ell'$. Let $\Delta_0 = \bar{x} : (\bar{T}, \bar{\kappa}), this : (loctype \ell, level \ell)$, $\eta_0 = [\bar{x} \mapsto \bar{d}, this \mapsto \ell]$, and $\eta'_0 = [\bar{x} \mapsto \bar{d}', this \mapsto \ell']$. We claim that $ok \Delta_0 \eta$ and $ok \Delta_0 \eta'$ and $\eta_0 \sim_{\Delta_0} \eta'_0$. Then we get the result $d = d'$ by safe μ . It remains to prove the claims. We give the argument for the case that \bar{x} is a single identifier, as the generalization is obvious but awkward to put into words.
For $\eta_0 \sim_{\Delta_0} \eta'_0$, note that since $\ell = \ell'$ it suffices to deal with \bar{d}, \bar{d}' regardless of whether $loctype \ell = L$. If $\bar{\kappa} = L$ then we need $\bar{d} = \bar{d}'$. Now $\bar{\kappa} = L$ implies $\overline{\bar{\kappa}} = L$ by the security typing rule, and then we get $\bar{d} = \bar{d}'$ by induction on \bar{e} ; moreover Lemma 5.1 yields $ok(\bar{x} : (\bar{T}, \bar{\kappa}))[\bar{x} \mapsto \bar{d}]$. Thus $ok \Delta_0 \eta$ because $ok(this : (loctype \ell, level \ell))[this \mapsto \ell]$ holds for any ℓ . We have $ok \Delta_0 \eta'$ *mutatis mutandis*.
- $(B) e$: By $\kappa = L$, we can use induction on e to obtain $\ell = \ell'$. Moreover, as we are considering the case where $\llbracket (B) e \rrbracket \mu \eta h \neq \perp \neq \llbracket (B) e \rrbracket \mu \eta' h'$, we have that ℓ is in both $dom h$ and $dom h'$. The result follows directly.
- e instance of $B, null, unit$: All are easy. \square

Lemma 6.2 (Safe commands)

Suppose $\Delta; C \vdash S : (com \kappa_1, \kappa_2)$. Suppose also $ok \mu, ok h, ok h', ok \Delta \eta, ok \Delta \eta', safe \mu, Hconf \mu, \eta \sim_{\Delta} \eta', h \sim h'$, and $\llbracket \Delta^\dagger; C \vdash S^\dagger : com \rrbracket \mu \eta h \neq \perp \neq \llbracket \Delta^\dagger; C \vdash S^\dagger : com \rrbracket \mu \eta' h'$. Then $\eta_0 \sim_{\Delta} \eta'_0$ and $h_0 \sim h'_0$, where $(\eta_0, h_0) = \llbracket \Delta^\dagger; C \vdash S^\dagger : com \rrbracket \mu \eta h$ and $(\eta'_0, h'_0) = \llbracket \Delta^\dagger; C \vdash S^\dagger : com \rrbracket \mu \eta' h'$.

Proof: By induction on the derivation of $\Delta; C \vdash S : (com \kappa_1, \kappa_2)$. Under the assumptions of the Lemma, we show $\llbracket \Delta^\dagger; C \vdash S : com \rrbracket \mu \eta h = \llbracket \Delta^\dagger; C \vdash S : com \rrbracket \mu \eta' h'$ by cases on S .

- $x := e$: This has no effect on the heap; we only need to show $[\eta \mid x \mapsto d] \sim_{(\Delta, x: (T_1, \kappa_1))} [\eta' \mid x \mapsto d']$, where $d = \llbracket \Delta^\dagger; C \vdash e : T_2 \rrbracket \mu \eta h$ and $d' = \llbracket \Delta^\dagger; C \vdash e : T_2 \rrbracket \mu \eta' h'$. Under the assumption $\eta \sim_{(\Delta, x: (T_1, \kappa_1))} \eta'$, it remains to show that $\kappa_1 = L$ implies $d = d'$. If $\kappa_1 = L$ then $\kappa_2 = L$, by typing, and then Lemma 6.1 yields $d = d'$. (Use of Lemma 6.1 depends on the assumptions $ok h$ etc.)

- $x.f := e$: This has no effect on the environment, so we only need to show that the result heaps are related. Let $h_0 = [h \mid \ell \mapsto [h\ell \mid f \mapsto d]]$ and let $h'_0 = [h' \mid \ell' \mapsto [h'\ell' \mid f \mapsto d']]$. We must show $h_0 \sim h'_0$. If the level κ_2 of f is H , then $h_0 \sim h$ (and $h'_0 \sim h'$) because h_0 and h are identical except for the H field f . So the result follows by transitivity of \sim . For the other case, $\kappa_2 = L$, we have $\kappa_1 = \kappa_3 = L$ by the typing rule. Since $\eta \sim_{(\Delta, x: (C_1, \kappa_1))} \eta'$, we obtain $\ell = \eta x = \eta' x = \ell'$. So it remains to show that $h_0 \ell f \sim h'_0 \ell' f$, i.e., $d = d'$. And this holds by Lemma 6.1 for e , using that $\kappa_3 = L$.
- $x := new B()$: For the environment, we must show $[\eta \mid x \mapsto \ell] \sim_{\Delta, x: (D_1, \kappa_1)} [\eta' \mid x \mapsto \ell']$. By assumption $\eta \sim_{\Delta, x: (D_1, \kappa_1)} \eta'$ it is enough to deal with x ; that is, if $\kappa_1 = L$ we need $\ell = \ell'$. By the typing rule, $\kappa_1 = L$ implies $level B = L$. Thus, by $h \sim h'$, we have $dom h \cap locs B = dom h' \cap locs B$; then $\ell = \ell'$ by parametricity of the allocator (Definition 1).

Finally, we get $h_0 \sim h'_0$ as follows. If either ℓ or ℓ' is in $LLoc$ then $level B = L$ so by parametricity of the allocator we get $\ell = \ell'$, satisfying the domain condition for $h_0 \sim h'_0$. For the range, i.e., $h\ell \sim h'\ell$, the result holds because the new object states are identical.

- $e.m(\bar{e})$: By semantics, the command has no effect on the environment, so it suffices to show $h_0 \sim h'_0$. (This is fortunate, because the statement of the Lemma uses identifiers η_0, η'_0 that are used differently in the semantics.)

We show $h_0 \sim h'_0$ by cases on κ_4 . If $\kappa_4 = H$ then it is possible that $\ell \neq \ell'$ and thus the two calls can have different behavior. But by the typing constraint $\kappa_4 \leq \kappa_3$ we have $\kappa_3 = H$ and thus $Hconf \mu$ yields $h_0 \sim h \sim h' \sim h'_0$. It remains to consider the case $\kappa_4 = L$. In this case, we have $\ell = \ell'$ by Lemma 6.1. Now let $\Delta_0 = \bar{x} : (\bar{T}, \bar{\kappa}), this : (loctype \ell, level \ell)$, $\eta_0 = [\bar{x} \mapsto \bar{d}, this \mapsto \ell]$, and $\eta'_0 = [\bar{x} \mapsto \bar{d}', this \mapsto \ell']$. We claim that $ok \Delta_0 \eta$ and $ok \Delta_0 \eta'$ and $\eta_0 \sim_{\Delta_0} \eta'_0$. Then we get the result $h_0 = h'_0$ by safe μ . Owing to our choice of identifiers in the typing rules, the proof of the claim is just the same as in the case for method call as expression (see the proof of Lemma 6.1).

- $S_1; S_2$: Use induction and L -confinement (Lemma 5.2) for S_1 ; then induction on S_2 .
- $if e S_1 else S_2$: We proceed by cases on level κ_5 of the guard e . Suppose $\kappa_5 = L$. Then by Lemma 6.1 for $e, b = b'$. If $b = true$, the result follows by induction on S_1 and if $b = false$, the result follows by induction on S_2 . Consider the other case, $\kappa_5 = H$. By typing, $\kappa_6 = H = \kappa_7$ and $\kappa_1 = \kappa_2 = \kappa_3 = \kappa_4 = H$. Let $(\eta_0, h_0) = \llbracket \Delta^\dagger; C \vdash if e S_1 else S_2 : com \rrbracket \mu \eta h$ and $(\eta'_0, h'_0) = \llbracket \Delta^\dagger; C \vdash if e S_1 else S_2 : com \rrbracket \mu \eta' h'$.

By H -confinement Lemma 5.4 we have $\eta \sim_{\Delta} \eta_0$, $\eta' \sim_{\Delta} \eta'_0$, $h \sim h_0$, and $h' \sim h'_0$. Using assumptions $\eta \sim_{\Delta} \eta'$ and $h \sim h'$ we get $\eta_0 \sim_{\Delta} \eta'_0$ and $h_0 \sim h'_0$ by transitivity.

- $\text{var } (T, \kappa_1) \ x := e \text{ in } S$: First, we have $[\eta \mid x \mapsto d] \sim_{(\Delta, x: (T, \kappa_1))} [\eta' \mid x \mapsto d']$ because if $\kappa_1 = L$ then by typing $\kappa_4 = L$, so, by Lemma 6.1, $d = d'$. Second, we have $ok(\Delta, x : (T, \kappa_1)) [\eta \mid x \mapsto d]$ by Lemma 5.1. So we can use induction on S to get $\eta_0 \sim_{(\Delta, x: (T, \kappa_1))} \eta'_0$ and $h_0 \sim h'_0$, and hence $(\eta_0 \mid x) \sim_{\Delta} (\eta'_0 \mid x)$. \square

Theorem 6.3 (Noninterfering programs) The meaning $\hat{\mu}$ of a well-formed class table is safe: *safe* $\hat{\mu}$.

Proof: Because *safe* $\hat{\mu}$ is defined as a fixpoint, we first show that *safe* μ_i for all i , by induction on i . Then the result follows by fixpoint induction.

We have *safe* μ_0 because $\mu_0 C m$ is constantly \perp .

Suppose *safe* μ_i , to show *safe* μ_{i+1} . By definition, we must show *safe* $\mu_{i+1} C m$ for each C, m . There are two cases, depending on whether m is declared or inherited.

Suppose m has declaration

$M = \tau_1 \ \kappa_1 \ m(\bar{\tau} \ \bar{x}) \ \kappa_2 \ \{S; \text{return } e\}$ in C and let $\Delta = \bar{x} : \bar{\tau}, \text{this} : (C, \text{level } C)$. By Lemmas 5.3 and 5.5 we have $ok \ \mu_i$ and $Hconf \ \mu_i$. Suppose $ok \ \Delta \ \eta$ and $ok \ h$, and let $(\eta_0, h_0) = \llbracket (\bar{x} : \bar{\tau}, \text{this} : C); C \vdash S : \text{com} \rrbracket_{\mu_i} \eta h$ (if the outcome is \perp there is nothing more to prove). By Lemma 5.2, L -confinement of commands, we have $ok \ \Delta \ \eta_0$, $ok \ \Delta \ \eta'_0$, $ok \ h$, and $ok \ h'$. By Lemma 6.2, safety for commands, we have $h_0 \sim h'_0$ and $\eta_0 \sim_{\Delta} \eta'_0$. It remains to show that if the result level κ_4 for m is L we have $d = d'$. But by security typing for method declarations, if $\kappa_4 = L$ then the return expression e is typed L . Using $ok \ \mu_i$, $\eta_0 \sim_{\Delta} \eta'_0$, etc., Lemma 6.1, safety for expressions, yields $d = d'$. This concludes the proof of *safe* $\mu_{i+1} C m$.

Suppose m is inherited in C from superclass D . Let $\Delta_C = \bar{x} : \bar{\tau}, \text{this} : (C, \text{level } C)$ and $\Delta_D = \bar{x} : \bar{\tau}, \text{this} : (D, \text{level } D)$. We claim that, for any η, η' , $ok \ \Delta_C \ \eta \Rightarrow ok \ \Delta_D \ \eta$ and $\eta \sim_{\Delta_C} \eta' \Rightarrow \eta \sim_{\Delta_D} \eta'$. Then *safe* $\mu_{i+1} C m$ follows from the claim and *safe* $\mu_{i+1} D m$ which was already proved. (Strictly speaking we are using secondary induction on inheritance chains.)

For the claim, we only need to consider *this*, as otherwise Δ_C and Δ_D are the same. For *this*, $ok \ \Delta_D \ \eta$ requires $\text{level } D = L \Rightarrow \eta \text{this} \in LLoc$. From $C \leq D$ we get $\text{level } D \leq \text{level } C$ by the typing rule for classes. Moreover, since m is inherited from C the rule requires $\text{level } D = \text{level } C$ so we are done. \square

7. Discussion

Beyond the progress reported here, much remains to be done. Non-interference is an attractive property because it can be easily formalized and can provide a precise description of end-to-end security in a system. By itself, however, noninterference as an information flow policy can be rather restrictive. As has been shown by several researchers, a controlled amount of declassification or downgrading of sensitive information is needed in realistic systems for them to be useful. As a future extension of this work, we expect to formalize noninterference in the presence of declassification following the work of Zdancewic and Myers [35].

One direction of work that we have already pursued is adding Java's access control mechanism to the core language [3]. We plan to add information flow annotations to this language. Then, allowing declassification may lead to leakage of information, but the access control mechanism can possibly be used to obtain a noninterference result.

Java has quite a few features beyond the language treated here. To extend our language to the remaining features of JavaCard [7], the semantics can be extended using standard techniques. To treat expressions with side effects, both the environment and the heap would be threaded through the semantics of expressions. We have avoided this in the current paper because it is unilluminating. Exceptional control flow would add further semantic complications of a similar kind. The other missing features have to do with scope and visibility: protected fields, private and protected classes, interfaces, and packages. These features can be treated in the typing rules, similarly to our treatment of private fields, and the semantic consequences could perhaps be exploited to reduce the need for security annotations.

Features of Java beyond those of JavaCard pose a bigger challenge: threads, class loading [9], reflection, and serialization. Specifying noninterference for such constructs would probably go hand-in-hand with specification of pointer confinement and data abstraction properties.

As a step towards more general pointer confinement and abstraction, we are already studying polymorphic classes as in GJ [6]. Label polymorphism is also desirable [21, 20], e.g., for library classes. Label polymorphism might lessen the practical need for H-subclasses of L-classes, which in turn would allow simplification of the security typing rules.

An important implementation issue is which security annotations can be left implicit, to be inferred by a type reconstruction algorithm. We have not addressed type reconstruction in the current work, but expect that techniques from Pottier *et al.* can be adapted [22, 23].

Acknowledgement: To Geoffrey Smith for discussions that helped clarify noninterference and to the anonymous referees for their comments.

References

- [1] M. Abadi, A. Banerjee, N. Heintze, and J. G. Riecke. A core calculus of dependency. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 147–160, 1999.
- [2] K. Arnold and J. Gosling. *The Java Programming Language, second edition*. Addison-Wesley, 1998.
- [3] A. Banerjee and D. A. Naumann. Representation independence, confinement and access control. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 166–177, 2002.
- [4] G. Barthe and B. Serpette. Partial evaluation and non-interference for object calculi. In A. Middeldorp and T. Sato, editors, *Proceedings of FLOPS'99*, volume 1722 of *LNCS*, pages 53–67. Springer-Verlag, 1999.
- [5] J. Boyland, J. Noble, and W. Retert. Capabilities for sharing: A generalisation of uniqueness and read-only. In *ECOOP 2001 - Object-Oriented Programming*, pages 2–27, 2001.
- [6] G. Bracha, M. Odersky, D. Stoutamire, and P. Wadler. Making the future safe for the past: Adding genericity to the Java programming language. In *ACM Symposium on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, pages 183–200, Oct. 1998.
- [7] Z. Chen. *Java Card Technology for Smart Cards*. Addison-Wesley, 2000.
- [8] D. G. Clarke, J. Noble, and J. M. Potter. Simple ownership types for object containment. In J. L. Knudsen, editor, *ECOOP 2001 - Object Oriented Programming*, pages 53–76, 2001.
- [9] R. D. Dean. *Formal Aspects of Mobile Code Security*. PhD thesis, Princeton University, 1999.
- [10] D. Denning. A lattice model of secure information flow. *Communications of the ACM*, 19(5):236–242, 1976.
- [11] D. Denning and P. Denning. Certification of programs for secure information flow. *Communications of the ACM*, 20(7):504–513, 1977.
- [12] E. Ferrari, P. Samarati, E. Bertino, and S. Jajodia. Providing flexibility in information flow control for object-oriented systems. In *Proceedings of IEEE Symposium on Security and Privacy*, pages 130–140, 1997.
- [13] J. Goguen and J. Meseguer. Security policies and security models. In *Proceedings of the 1982 IEEE Symposium on Security and Privacy*, pages 11–20, 1982.
- [14] C. Grothoff, J. Palsberg, and J. Vitek. Encapsulating objects with confined types. In *ACM Symposium on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, pages 241–253, 2001.
- [15] N. Heintze and J. G. Riecke. The SLam calculus: programming with secrecy and integrity. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 365–377, 1998.
- [16] J. Hogg. Islands: Aliasing protection in object-oriented languages. In *ACM Symposium on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, pages 271–285, 1991.
- [17] A. Igarashi, B. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. *ACM Trans. Prog. Lang. Syst.*, 23(3):396–459, May 2001.
- [18] K. R. M. Leino and G. Nelson. Data abstraction and information hiding. Technical Report 160, COMPAQ Systems Research Center, Nov. 2000. To appear in TOPLAS.
- [19] H. Mantel and A. Sabelfeld. A generic approach to the security of multi-threaded programs. In *Proceedings of 14th IEEE Computer Security Foundations Workshop*, pages 126–142, 2001.
- [20] A. C. Myers. JFlow: Practical mostly-static information flow control. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 228–241, 1999.
- [21] A. C. Myers and B. Liskov. Complete, safe information flow with decentralized labels. In *Proceedings, IEEE Symposium on Security and Privacy*, pages 186–197, 1998.
- [22] F. Pottier and S. Conchon. Information flow inference for free. In *Proceedings of the fifth ACM International Conference on Functional Programming*, pages 46–57, 2000.
- [23] F. Pottier and V. Simonet. Information flow inference for ML. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 319–330, 2002.
- [24] J. C. Reynolds. Types, abstraction, and parametric polymorphism. In R. Mason, editor, *Information Processing '83*, pages 513–523. North-Holland, 1984.
- [25] A. Sabelfeld and D. Sands. Probabilistic noninterference for multi-threaded programs. In *Proceedings of 13th IEEE Computer Security Foundations Workshop*, pages 200–215, 2000.
- [26] A. Sabelfeld and D. Sands. A Per model of secure information flow in sequential programs. *Higher-order and Symbolic Computation*, 14(1):59–91, 2001.
- [27] G. Smith. A new type system for secure information flow. In *Proceedings of 14th IEEE Computer Security Foundations Workshop*, pages 115–125, 2001.
- [28] G. Smith and D. Volpano. Secure information flow in a multi-threaded imperative language. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 355–364, 1998.
- [29] I. Stark. *Names and Higher-Order Functions*. PhD thesis, University of Cambridge, 1994. <http://www.dcs.ed.ac.uk/home/stark/publications/thesis.html>.
- [30] E. Sumii and B. Pierce. Logical relations for encryption. In *14th IEEE Computer Security Foundations Workshop*, pages 256–269, 2001.
- [31] J. Vitek and B. Bokowski. Confined types in java. *Software Practice and Experience*, 31(6):507–532, 2001.
- [32] D. Volpano and G. Smith. A type-based approach to program security. In *Proceedings of TAPSOFT'97*, number 1214 in *LNCS*, pages 607–621. Springer-Verlag, 1997.
- [33] D. Volpano and G. Smith. Confinement properties for multi-threaded programs. *Electronic Notes in Theoretical Computer Science*, 20, 1999.
- [34] D. Volpano, G. Smith, and C. Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3):167–187, 1996.
- [35] S. Zdancewic and A. Myers. Robust declassification. In *Proceedings of 14th IEEE Computer Security Foundations Workshop*, pages 15–23, 2001.