Succour to the Confused Deputy

Types for Capabilities

Radha Jagadeesan, Corin Pitcher, and James Riely

DePaul University

Abstract. The possession of secrets is a recurrent theme in security literature and practice. We present a refinement type system, based on indexed intuitonist S4 necessity, for an object calculus with explicit locations (corresponding to principals) to control the principals that may possess a secret. Type safety ensures that if the execution of a well-typed program leads to a configuration with an object p located at principal a, then a possesses the capability to p. We illustrate the type system with simple examples drawn from web applications, including an illustration of how Cross-Site Request Forgery (CSRF) vulnerabilities may manifest themselves as absurd refinements on object declarations during type checking.

A fuller version of the paper is available at fpl.cs.depaul.edu/jriely/papers/ 2012-aplas.pdf.

1   Introduction

Many systems depend upon a prescribed usage of secrets to enforce policies that incorporate secrecy, integrity, authentication, authorization, and auditing concerns. Nevertheless, it may be computationally expensive, or impossible in some adversarial models, to control the *use* of secrets directly. For this reason, it is common to control the *possession* of secrets instead of their use. However, invariants about the possession of secrets can fail due to inadequately-specified interfaces or a lack of agreement between software components. We illustrate some of the issues with two examples.

Object References The Java security manager permits access control checks based upon permissions assigned to code [19]. This allows control over systems composed of code from different sources. The `java.io.FileOutputStream` system class utilizes access control checks in the following manner:

- The `FileOutputStream` constructor checks for the relevant file write permission.
- For performance, `FileOutputStream` methods do not have access control checks.

The lack of access control checks after construction means that references to instances of `FileOutputStream` can be used to write a file by untrusted code, *if* the reference is made available to the untrusted code. For this reason, sensitive object references must be confined to trusted code.

Cross-Site Request Forgery and the Confused Deputy Cross-Site Request Forgery (CSRF) attacks [13] are acknowledged as an instance of the Confused Deputy problem [22]. A principal is a Confused Deputy if it uses its authority to mistakenly act

on behalf of an initiating principal. The capability-based solution [22] to the Confused Deputy problem requires the initiating principal to provide a capability to the Deputy, which the Deputy requires to complete its actions. Since capabilities authorize access to a resource without further checks, their possession must be constrained programmatically. For example, one might ask:

- If the initiating principal hands a capability to the Deputy, to whom can the Deputy pass the capability?
- Is the Deputy permitted to add its own capabilities to any request from the initiating principal?

In the case of a web browser, acting as a Deputy for both user and JavaScript behavior on web pages, it is permitted to add cookies to outgoing HTTP requests based on the URLs determined from the web pages. Several browser extensions provide a more restrictive policy on forwarding cookies for cross-site requests to prevent misunderstandings in web applications vulnerable to CSRF attacks.

In this paper, we address control over the possession of secrets via the use of logical specifications embedded in the types of a distributed programming language. Static analysis is used to verify that programs comply with possession policies, yielding an upper bound on the principals that may possess a secret.

Approach. The main contribution of this paper is the application of a *refinement type system for a distributed object-oriented language*. The type system controls possession of object references, representing secrets, via specifications in a principal-indexed variant of intuitionist S4.

We specify possession of secrets using intuitionist S4 logic [6,28]. Instead of a single modality, we consider modalities indexed by principals. An indexed modality of the form $\Box_a \Phi$ represents a predicate Φ that is permissible for principal a [17,10]. We use a "may possess" predicate $\mathtt{mp}(\mathtt{s})$ representing possession of a secret $\mathtt{s}$. Thus $\Box_a \mathtt{mp}(\mathtt{s})$ means that principal a is permitted to possess secret $\mathtt{s}$.

It is key to our approach that indexed modalities allow different principals to have different possession policies. In particular, $\Box_a \mathtt{mp}(\mathtt{s})$ and $\Box_b \mathtt{mp}(\mathtt{s})$ are independent statements about whether the different principals a and b may possess $\mathtt{s}$. The underlying logic then provides relationships between uses of modalities:

- Indexed modalities commute, i.e., $(\Box_a \Box_b \Phi) \Rightarrow (\Box_b \Box_a \Phi)$. This permits a sequence of indexed modalities to be treated as a multiset.
- The counit $(\Box_a \Phi) \Rightarrow \Phi$ allows a modality to be eliminated. The converse does not hold in general. Consequently, a right for principal a can be forgotten during logical deduction, but cannot be manufactured.
- From $\Box_a(\Phi_1 \Rightarrow \Phi_2)$ and $\Box_a \Phi_1$, we can deduce $\Box_a \Phi_2$; thus, the possessions of a principal are closed under deduction. From comultiplication, $(\Box_a \Phi) \Rightarrow (\Box_a \Box_a \Phi)$, we deduce that the deductions in the scope of a principal a includes the knowledge of a's posessions.
- If b is less secure than a and $\Box_b \Phi$ then we can deduce $\Box_a \Phi$; so, by this principle of Principal naturality, more secure principals have access to more secrets.

These relationships yield an indexed intuitionist S4 necessity modality, representing layers of permission for principals, over the underlying logic. This distinction in the

logic between permissions and who has those permissions, represented by principal-indexed modalities, greatly reduces the need to quantify over principals during reasoning. For example, if a policy states that $s2$ may be possessed if $s1$ may be possessed, written $mp(s1) \Rightarrow mp(s2)$, then the indexed intuitionist S4 necessity modality structure allows this implication to be lifted to any principal a as $(\Box_a mp(s1)) \Rightarrow (\Box_a mp(s2))$.

Noninterference theorems [25] justify the use of indexed intuitionist S4 necessity modalities in this modeling. In that paper, we show that noninterference captures the idea that there is no information flow between differently indexed modalities. Let α be a modality free formula. The intuitive idea behind non interference is that if $\Box_a \alpha$ is derivable from some deductively closed set of hypothesis, then it is derivable from a subset of those hypothesis that are in the scope of the modality indexed by a, i.e. the formulas of the form $\Box_a \cdot$. In particular, noninterference implies the unprovability of the following formulas:

- $\Box_a mp(s1) \Rightarrow \Box_b mp(s1)$
- $(\Box_a(mp(s1) \Rightarrow mp(s2)) \wedge \Box_b mp(s1)) \Rightarrow \Box_a mp(s2)$

The unprovability of $\Box_a mp(s1) \Rightarrow \Box_b mp(s1)$ shows that the logical reasoning does not transfer capabilities unrestrictedly between principals. The unprovability of the second formula $(\Box_a(mp(s1) \Rightarrow mp(s2)) \wedge \Box_b mp(s1)) \Rightarrow \Box_a mp(s2)$ ensures that the acquisition of a new capability ($s1$) by another principal (b) does not create new capabilities for principal a by purely logical reasoning. Thus, non-interference facilitates distribution and decentralized enforcement of policies in the following sense. The reference monitor at a location uses logical reasoning to deduce whether a principal has sufficient capabilities to access the resource available at the location. Noninterference ensures that this reasoning is not dependent on other principals; so, the reference monitor at a location can function without knowledge of the principals at other locations.

We present three analyses to establish the utility of our approach:

- Sealed objects (Section 2) that demonstrate modeling of symmetric cryptography [2].
- An object encoding (Section 5) of Hardy's Confused Deputy [22].
- A web browser and server model to explore browser security policies and Cross-Site Request Forgery prevention solutions. For space reasons, this example is in the full version of this paper at fpl.cs.depaul.edu/jriely/papers/2012-aplas.pdf.

Related Work

Capability-based systems. Capabilities have been used to realize security policies in a variety of systems, e.g., [23,29,5] to name but a few. Distributed object languages such as E [12] illustrate the "capabilities-as-object references" paradigm where both subjects and resources are represented uniformly as objects, and classical object-oriented mechanisms are used to structure the exchange and invocation of capabilities. This viewpoint underlies Caja, a safe subset of Javascript. Caja eschews direct references to DOM objects, instead providing references to wrappers that restrict the capabilities provided on DOM objects. [26] formalize a notion of capability-safety, show that the subset Cajita satisfies this property and derive that Cajita programs have inter-component isolation.

Type systems for secrecy, confinement, and access control. In object-oriented languages, ownership and confinement types (see [11] for a survey of ownership type system) aim to delimit the portions of the object reference graph that can have references

to the objects under consideration. In this paper, we generalize from confinement types to multi-party secrecy types using refinement types built on intuitionist S4 to express dependencies.

Abadi [4] describes a type system for controlling secret keys in the spi calculus, using a binary division of code as either fully trusted or untrusted. This paper explores an idea stated there: "distinguish various principals within the system, and enable us to discuss which of these principals have a given piece of data".

Language-based approaches to access control have long been studied in the setting of process calculi, though these approaches are not based explicitly in logic; two early references are [27,24]. In [15], Fournet, Gordon and Maffeis validate authorization policies statically using a specification language with "expect" assertions in a Datalog-style language. The says family of principal-indexed modalities is used in logics for reasoning about authorization statements made by different principals [18,1].The says modality has a monadic structure, as exemplified by the unit law ($\Phi \Rightarrow a$ says Φ). In our prior work [9] we develop a type system based on authorization logic to capture provenance in a distributed object calculus. [14] explores the impact of compromised principals on authorization policies in a distributed setting.

In this paper, we carry out a similar program, albeit in the logical setting of intutionist S4, by reusing the infrastructure of *refinement types* [16] developed in the literature: policies (and therefore types) may quantify over object references of a given class. Object references (and variables) appear in logical formulae in equality predicates and in the "may-possess" predicate mp(.) described previously. Our semantics and notion of safety are from [9] and derive, ultimately, from [20] and [21].

2   Sealed Objects

In this section, we introduce the computational model and logic, by way of an example. The details of the logic can be found in a companion paper [25]; here we summarize the properties of the logic required to understand the example.

Computation is based on threads that communicate via a shared heap. Threads are "located" at the principal for which the thread is running; similarly objects are "located" at the principal that created the object. We use the terms "principal" and "location" interchangeably. For an object p, the location is available to the programmer via the pseudo-field p.loc.

Neither threads nor objects can change location; however, object references can be communicated between threads using shared objects. A method invocation on an object leads to code execution at the location of the callee object. Thus, when the caller and callee objects are located at different locations, method invocation leads to a change of location context.

We conflate opponents, representing them all via $\bot$. Threads acting on behalf of opponents can only instantiate classes with trivial invariants, discussed below. Threads acting on behalf of non-opponents must obey a global policy. All threads must be well typed according to typical object-oriented programming rules, e.g., as in Java. Additionally, our type system controls communication of object references by non-opponent threads.

Principals are ordered by a partial order with least principal being the Opponent $\bot$. Principal naturality allows that whenever $\Box_\bot \Phi$ is deducible, then so is $\Box_a \Phi$, for any a. In particular, this means that any of the Opponent capabilities are available to all principals. Thus, our type system does not impose any restrictions more than those of usual object oriented programming on Opponent programs.

A program is *safe* if every object reference that is available to a principal at runtime is permitted by the global specification of permitted capabilities. Our type system ensures that safe well-typed programs remain safe under evaluation in the face of arbitrary opponent processes.

Consider `javax.crypto.SealedObject`. It permits a serializable object to be encrypted with a secret key and a symmetric-key cipher. The constructor is responsible for serialization and encryption. The resulting `SealedObject` contains only ciphertext. The original object can be recovered by passing the same secret key to `getObject`. We model `SealedObject` as:

```
class SealedObject {
  private final SecretKey key;
  private final Object contents;
  public SealedObject (SecretKey key, Object contents) {
    this.key = key; this.contents = contents;
  }
  public Object getObject (Section key) {
    if (key == this.key) return this.contents;
    else return null;
  }
}[ □⊥(mp(this.key) ⇒ mp(this.contents)) ]
```

By controlling possession of the `key`, one controls access to the `contents`. This code uses private fields guarded by object equality rather than encryption. This is sufficient since the type system enforces that the caller of `getObject` must possess `key`.

Specifications in our system are divided between a *global policy* and a set of *class invariants*. Intuitively, the combinatation of these policies indicates upper bounds on the capabilities that can be possessed by a principal. Our safety theorem shows that at any stage in the evolution of a system, even in the presence of opponents, any principal only possesses references that are provided for in the policy.

The global policy describes the distribution of initial secrets, and also any potential relationships between classes. It is informative to consider the following extremal global policies. Suppose that all class invariants are trivial (i.e., `tt`).

– The extremely permissive global policy $\forall \eta.\ \Box_\bot \mathtt{mp}(\eta)$ does not forbid any transmission of objects. Thus, typing under this global policy is essentially the same as standard object-oriented typing.
– The extremely restrictive global policy `tt` in the case where there are only two principals — Opponent ($\bot$) and Secret ($\top$) — forbids all transmission of objects from $\top$ to $\bot$. Thus, typing under this global policy is essentially the same as standard information flow.

The class invariant is intended to describe the private internals of a single class. The mutable state in our objects is only in the form of private instance variables. The class

invariant is written at the end of each class, in square brackets. Because we are in a concurrent setting, we make the simplifying assumption that only final fields may be mentioned in the class invariant and that constructors may do nothing but assign fields — we also disallow reassignment of method parameters and local variables. Thus, the class invariant holds for every object at the point its constructor terminates.

References to `SealedObjects` can be safely sent anywhere because they do not leak their contents arbitrarily. The fact that they are allowed anywhere is exemplified by the global policy $(\forall o\texttt{:SealedObject}.\ \Box_\bot \texttt{mp}(\texttt{o}))$ — type-sorted quantification is shorthand for quantification using a "type" predicate on objects. This policy allows `SealedObjects` to be given to opponents; however, they can only retrieve the contents if they have the matching key. More restrictive policies are also possible.

The class invariant of `SealedObject` indicates that any principal that may possess `this.key` may also possess `this.contents`. Opponents cannot create secrets, and therefore are restricted to creating instances of "global" classes with invariants (i.e. "true") that are trivially satisfied. The invariant of `SealedObject` is nontrivial, and therefore opponents may not create instances of the class.

The invariant must be statically justifiable by any code that creates an instance of the class. For example, consider the code new SealedObject (key, acct), where `key` is an instance of `SecretKey` and `acct` is an instance of a `BankAccount` class. We must establish that every principal that may possess `key` may also possess `acct`, written $\Box_\bot(\texttt{mp}(\texttt{key}) \Rightarrow \texttt{mp}(\texttt{acct}))$. This might be accomplished using a global policy that allows `acct` to be possessed anywhere, written $\Box_\bot(\texttt{mp}(\texttt{acct}))$. Stricter policies could be specified pairwise, including $\Box_\bot(\texttt{mp}(\texttt{key}) \Rightarrow \texttt{mp}(\texttt{acct}))$ as a fact. More flexible arrangements are also possible, for example, using the invariant of the factory class that creates new keys. In any case, the implication must be deduced from the available policy in order to instantiate `SealedObject`. In all non-trivial cases, the initial ability to create `SealedObjects` is specified as part of the global policy; indeed, the non-interference theorems ensure that there is no possible creation of `SealedObjects` otherwise.

In order to justify safety of the `getObject` method, we first observe that the caller to `getObject` must possess the key, i.e., $\Box_{\text{caller}}(\texttt{mp}(\texttt{key}))$. From the `SealedObject` class invariant, we know that:

$$\Box_\bot(\texttt{mp}(\texttt{this.key}) \Rightarrow \texttt{mp}(\texttt{this.contents}))$$

From which we can deduce that (note the principal on $\Box$):

$$\Box_{\text{caller}}(\texttt{mp}(\texttt{this.key}) \Rightarrow \texttt{mp}(\texttt{this.contents}))$$

After the reference equality test (key==this.key), the callee knows key = this.key. Moreover, equality can be lifted to comodalities, and we have $\Box_{\text{caller}}(\texttt{key} = \texttt{this.key})$. From $\Box_{\text{caller}}(\texttt{mp}(\texttt{key}))$ and $\Box_{\text{caller}}(\texttt{key} = \texttt{this.key})$, we deduce $\Box_{\text{caller}}(\texttt{mp}(\texttt{this.key}))$. In conjunction with the implication above, we find that $\Box_{\text{caller}}(\texttt{mp}(\texttt{this.contents}))$. This justifies return of `this.contents` to the caller. In the case where the equality test fails, we use the property that `null` may be possessed anywhere, written $\Box_\bot(\texttt{mp}(\texttt{null}))$.

If the `SealedObject` class had public fields, then a higher threshold must be met to instantiate the class. In this case, one would also need to establish that any principal that may possess the object may also possess the values placed into the public fields.

It is worth noting that other symmetric cryptography schemes can be encoded as simple variants of `SealedObject`. For example, using nested conditionals, one can encode an object requiring n keys to encrypt and $k \leq n$ keys to decrypt.

3 Language

To formalize the preceding discussion, we first describe a distributed class-based language with mutable objects [9]. The operational semantics borrows heavily from [20], adding distribution [7,8] and classes. We consider typing in Section 4.

Syntax Names for classes (c, d), methods (ℓ), fields (f, g), variables (x, y, z), objects (p, q) and principals (a, b) are drawn from separate namespaces, as usual. Predicate variables (α, β) and predicate constructors (γ) occur in static annotations used during type-checking.

The reserved words of the language include: the variable name "this"; the principal "caller"; the class name Object; the predicate constructors "tt", "ff", "⇒", "∧", "∨", "¬" and "□". We write binary constructors infix.

The language is explicitly typed. Object types ($c{<}\vec{\phi}{>}$) include the actual predicate parameters $\vec{\phi}$, which we treat formally as *extended values*. Value types include objects (C), principals (Prin) and Unit. Extended value types include predicate types (P), which are resolved during typechecking. The process type (Proc) has no values.

One may write classes and methods that are generic in the predicate variables, achieving ML-style polymorphism with respect to effects. Class declarations thus include the formal predicate parameters $\vec{\alpha}$, which may occur in the effect Φ (see next table) associated with instances of the class. In addition to effects, class declarations include field and method declarations, but omit implicit constructor declarations. Fields include mutability annotations that are used in the statics. The syntax is as follows[1].

Types, Annotations, Class and Method Declarations

C, D ::= $c{<}\vec{\phi}{>}$		Object Types
T, S ::= C | Prin | Unit		Value Types
P, Q ::= Pred($\vec{\mathscr{T}}$)		Predicate Types
$\mathscr{T}, \mathscr{S}$ ::= T | P | Proc		Types
μ ::= private final | private mutable | public final		
$\mathscr{D}$ ::= class $c{<}\vec{\alpha} : \vec{P}{>} \triangleleft D\{\vec{\mu}\ \vec{T}\ \vec{f};\ \vec{\mathscr{M}}\}[\Phi]$		
$\mathscr{M}$ ::= ${<}\vec{\beta} : \vec{Q}{>}S\,\ell(\vec{T}\ \vec{x})\{M\}$		

Values, Terms, Evaluation Contexts

$V, W, U, A, B, \phi, \psi$ ::= x | p | a | unit | α | γ | $\phi(\vec{V})$	

---

[1] When writing definitions using classes and methods, we sometimes omit irrelevant bits of syntax, e.g., we leave out the parameters to classes when empty, such as writing Object rather than Object<∅>. We identify syntax up to renaming of bound names, and write $M\{\!|^V\!/\!x|\!\}$ for substitution of V for x in M (and similarly for other categories). We often omit type information. We use standard syntactic sugar in place of explicit sequencing. For example, we may write "$y.f.g$" to abbreviate "let $x = y.f$; $x.g$".

$$M,N,L,\Phi,\Psi ::= V \mid V.f \mid V.\text{loc}$$
$$\quad \text{if } V=W \text{ then } M \text{ else } N \mid \text{let } x=N; M \mid N \parallel M$$
$$\quad V.f:=W \mid \text{let } x=\text{new } c\langle\vec{\phi}\rangle(\vec{V}); M \mid \text{let } x=V.\ell\langle\vec{\phi}\rangle(\vec{W}); M$$
$$\quad p:C\{\vec{f}=\vec{V}\} \mid (\nu p:C) M \mid a[M]_c^b$$
$$\mathbb{E} ::= [-] \mid a[\mathbb{E}]_c^b \mid \text{let } x=\mathbb{E}; M \mid \mathbb{E} \parallel N \mid M \parallel \mathbb{E} \mid (\nu p) \mathbb{E}$$

We use the metavariables ϕ, ψ, Φ and Ψ to represent values and terms of predicate type, and the other metavariables to represent runtime values and terms, with A and B reserved for values of principal type. Predicates are static annotations used in type-checking; they play no role in the dynamics. An *expectation* "expect Φ" as in [15] can be coded as "new Proof<Φ>()", where class Proof is defined "class Proof<α : Pred>{}[α]".

The last three constructs in the definition of terms — $p:C\{\vec{f}=\vec{V}\}$, $(\nu p:C) M$, and $a[M]_c^b$ — are *dynamic constructs*. These constructs are not allowed in method declarations or initial code.

With the exception of $V.\text{loc}$, $N \parallel M$, and the terms on the last line of the definition, the constructs of the language are standard for class-based languages with generics.

The special "field" loc returns the location of an object. Concurrent composition ($\parallel$) is asymmetric. In $N \parallel M$, the returned value comes from M; the term N is available only for side effects. The terms on the last line are are not allowed to appear in declarations, as they model the runtime heap and call stack. These include heap elements $p:C\{\cdots\}$ (indicating that p is located at a with actual class C and fields $\vec{f}=\vec{V}$), name restriction (νp) (indicating that p is a fresh name) and frames $a[M]_c^b$ (indicating that M is running under authority of principle a and class c, with result available to b). We write irreducible frames simply as $a[M]$.

Evaluation is defined using a structural congruence on terms. Let $\equiv$ be the least congruence on terms that satisfies the following axioms. The rules for concurrent composition are from [20]. They capture properties of concurrent composition, including semi-associativity and the interaction with let. The rules for distribution are inspired by [8]. The interpretation of a value is independent of the location at which it occurs and the computation of a frame does not depend upon the location from which the frame was invoked (eg. $a[b[M]_d^{b'}]_c^{a'} \equiv b[M]_d^{b'}$) and axiomatize the interaction of let with distribution (eg. $a[\text{let } x=N; M]_c^{a'} \equiv \text{let } x=a[N]_c^{a'}; a[M]_c^{a'}$).

Structural Congruence $(M \equiv M')$ (where $p \notin \mathit{fn}(M)$)

$(M \parallel N) \parallel L \equiv M \parallel (N \parallel L)$	
$(M \parallel N) \parallel L \equiv (N \parallel M) \parallel L$	$a[\text{let } x=b[V]_d^a; M]_c^{a'} \equiv a[\text{let } x=V; M]_c^{a'}$
$((\nu p) N) \parallel M \equiv (\nu p)(N \parallel M)$	$a[b[M]_d^{b'}]_c^{a'} \equiv b[M]_d^{b'}$
$M \parallel ((\nu p) N) \equiv (\nu p)(M \parallel N)$	$a[N \parallel M]_c^{a'} \equiv a[N]_c^{a'} \parallel a[M]_c^{a'}$
$\text{let } x=(L \parallel N); M \equiv L \parallel (\text{let } x=N; M)$	$a[(\nu p) N]_c^{a'} \equiv (\nu p) a[N]_c^{a'}$
$\text{let } x=((\nu p) N); M \equiv (\nu p)(\text{let } x=N; M)$	$a[\text{let } x=N; M]_c^{a'} \equiv \text{let } x=a[N]_c^{a'}; a[M]_c^{a'}$

The evaluation relation is defined with respect to an arbitrary fixed class table. The class table is referenced indirectly in the semantics through the lookup functions *fields* and *body*. We refer the reader to the full paper for the routine definitions of these functions.

Term Evaluation $(M \to M')$

$$\text{let } y = \text{new } C(\vec{V}); L \to (\nu p{:}C)(p{:}C\{\vec{f}=\vec{V}\} \mathbin{\Vert} L\{\!|^p\!/_y|\!\})$$
$$\quad \text{if } \mathit{fields}(C) = \vec{f} \;\; \text{and} \;\; |\vec{f}| = |\vec{V}|$$
$$b[p{:}C\{\cdots\}] \mathbin{\Vert} a[\text{let } y = p.\ell(\vec{W}); L]^{a'}_d \to b[p{:}C\{\cdots\}] \mathbin{\Vert} a[\text{let } y = b[M']^a_c; L']^{a'}_d$$
$$\quad \text{if } \mathit{body}(C.\ell) = (\vec{x})\{M\} \;\; \text{and} \;\; |\vec{x}| = |\vec{W}| \;\; \text{and} \;\; M' = M\{\!|^a\!/_{\mathsf{caller}}|\!\}\{\!|^p\!/_{\mathsf{this}}|\!\}\{\!|^{\vec{W}}\!/_{\vec{x}}|\!\} \;\; \text{and} \;\; C = c{<}\cdots{>}$$
$$b[p{:}C\{\cdots\}] \qquad \mathbin{\Vert} p.\mathsf{loc} \;\to b[p{:}C\{\cdots\}] \qquad \mathbin{\Vert} b$$
$$b[p{:}C\{f=V\cdots\}] \mathbin{\Vert} p.f := W \to b[p{:}C\{f=W\cdots\}] \mathbin{\Vert} \mathsf{unit}$$
$$b[p{:}C\{f=V\cdots\}] \mathbin{\Vert} p.f \qquad \to b[p{:}C\{f=V\cdots\}] \mathbin{\Vert} V$$
$$\text{if } V = V \text{ then } M \text{ else } N \to M$$
$$\text{if } V = W \text{ then } M \text{ else } N \to N \;\; \text{if } V \neq W \qquad\qquad \frac{M \equiv N \to N' \equiv M'}{M \to M'} \qquad \frac{M \to M'}{\mathbb{E}[M] \to \mathbb{E}[M']}$$
$$\text{let } x = V; M \to M\{\!|^V\!/_x|\!\}$$

The new construct creates an object and returns a reference to it. The result is a concurrent composition: the new object appears on the left, the return value on the right. Method invocation happens at the callee site, and thus a new frame is introduced in the consequent $b[M']^a_c$; the result of the method call will be made available to a. In M', the distinguished variables caller and this are bound to the calling principal and the object upon which the method is invoked respectively.

4   Types

The type system controls the distribution of object references via logical policies. We follow [15], as adapted to distributed OO languages with localities in [9].

By allowing predicates to include open values, we can reason about terms that include variables, such as x; however, we cannot reason about $x.f$. Thus we extend the type system to include equations between terms and values. Allowing any term is unsound, however, since our language includes mutability. Thus we identify a subset of *pure* terms which do not include mutable features. In addition, we require that evaluation of pure terms must terminate, and therefore we disallow method calls in pure terms. To shorten some definitions, we define a category of *identifiers*, η, which include bound names and principals.

$$\eta ::= x \mid p \mid a \mid \alpha$$

Environments have two types of data: type bindings for names (as usual) and logical phrases, including equalities and predicates. Define $\mathit{dom}(\Delta) = \{\eta \mid \eta : \mathscr{T} \in \Delta\}$.

$$\Delta ::= \emptyset \mid \Delta, \eta : \mathscr{T} \mid \Delta, \Phi \mid \Delta, V = M$$

Predicate lookup $(\mathit{effect}(C) = \Phi)$ is similar to method lookup. Here "$\Phi_D \wedge \Phi\{\!|^{\vec{\phi}}\!/_{\vec{\alpha}}|\!\}$" is sugar for "let $x = \Phi_D$; let $y = \Phi\{\!|^{\vec{\phi}}\!/_{\vec{\alpha}}|\!\}$; $x \wedge y$".

$$\frac{}{\mathit{effect}(\mathsf{Object}) = \mathsf{true}} \qquad \frac{\vec{\mathscr{D}} \ni \mathsf{class}\ c{<}\vec{\alpha}:\vec{P}{>} \triangleleft D\{\cdots\}[\Phi] \quad \mathit{effect}(D\{\!|^{\vec{\phi}}\!/_{\vec{\alpha}}|\!\}) = \Phi_D}{\mathit{effect}(c{<}\vec{\phi}{>}) = \Phi_D \wedge \Phi\{\!|^{\vec{\phi}}\!/_{\vec{\alpha}}|\!\}}$$

We also define a function $(\mathit{env}_a(M) = \Delta)$ to create an environment from a term.

$$\mathit{env}_a(\eta{:}C\{\vec{f}=\vec{V}\}) = \square_a\mathtt{mp}(\eta), a = \eta.\mathsf{loc}, V_1 = \eta.f_1, \ldots, V_n = \eta.f_n$$

$$env_a(\text{let } x = N; M) = env_a(N) \qquad env_a(N \,\Vert\, M) = env_a(N), env_a(M)$$
$$env_a(b\,[M]_c^{a'}) = env_b(M) \quad env_a(\,(\nu p : C)\, M) = p : C, env_a(M) \qquad env_a(M) = \emptyset,\ otherwise$$

The type system is parameterized with respect to a semantic entailment relation $(\Delta \vDash \Psi)$. In addition to the rules arising from indexed intuitionist necessity modalities, we expect the relation to support domain specific axioms and satisfy the following properties. Let σ stand for substitutions of pure terms M for x.

1. If $\Delta \vDash \Psi$ then $\Delta\sigma \vDash \Psi\sigma$, for any substitution σ from variables to values, or from principals to principals.
2. If $\Delta, V = V, \Delta' \vDash \Psi$ then $\Delta, \Delta' \vDash \Psi$.
3. If $\Delta, x : T, x = M, \square_a\text{mp}(x), \Delta' \vDash \Psi$ and $\Delta, \Delta' \vDash \square_a\text{mp}(M)$ then $\Delta, \Delta' \vDash \Psi\{\!\!\{^M/_x\}\!\!\}$.

In examples, we assume that whenever $\square_a\text{mp}(\eta)$ and $\eta : C$ are deducible, then so is $\square_a\text{mp}(\eta.f)$ for every public field of C.

The standard judgements required for the type system are relegated to the full paper, including subtyping $(\vdash \mathscr{T}' <: \mathscr{T})$, well-formed overriding $(\vdash <\vec{\beta} : \vec{Q}>S(\vec{T})$ overrides $D.\ell)$, well-formed types $(\Delta \vdash \mathscr{T})$, and well-formed environments $(\Delta \vdash \diamond)$. The only noteworthy aspect of these definitions is that the implication of the effects for the same base class also yields subtyping:

$$\frac{\vec{\mathscr{D}} \ni \text{class } c{<}\vec{\alpha}{>} \quad \vec{\phi} \vDash \vec{\psi} \quad |\vec{\alpha}| = |\vec{\phi}| = |\vec{\psi}|}{\vdash c{<}\vec{\phi}{>} <: c{<}\vec{\psi}{>}}$$

The judgments for declarations have the standard format. The judgment for values include a script a, indicating that the value is well typed at a specific location. The judgment for terms carries additional structure. In $\Delta \vdash_a^{a'} M : \mathscr{T}\ \rho\ d$, a should be read as the location of the term, a' as the location of the caller, $\mathscr{T}$ as the type of the resulting value, d as the class from which the code is derived, and $\rho \in \{\text{Pure}, \text{Impure}\}$ as a *purity annotation*.

The effect on a class must be a pure term of type Pred. The rule for typing methods uses a standard well-formed overriding definition. The typing of the method body occurs in the context of an abstract principal a that is constrained to coincide with the location of the ambient object. Similarly, the abstract principal caller is constrained to coincide with the annotation on the typing of the body of the method. In typing the method body, one can use the logical variables of the class, the method declaration and assume that the caller was permitted to possess the arguments.

Well-Formed Declarations $\quad (\Delta \vdash \mathscr{D}) \quad (\Delta \vdash \mathscr{M}\ in\ c{<}\vec{\alpha} : \vec{P}{>}{\triangleleft}D)$

$$\frac{\Delta, \vec{\alpha} : \vec{P} \vdash D, \vec{T} \quad \Delta, \vec{\alpha} : \vec{P}, a : \text{Prin}, \text{this} : c{<}\vec{\alpha}{>}, a = \text{this.loc}, \square_a\text{mp}(\text{this}) \vdash_a^{|a} \Phi : \text{Pred Pure } c}{\Delta \vdash \mathscr{M}\ in\ c{<}\vec{\alpha} : \vec{P}{>}{\triangleleft}D \quad \textit{fields}(D) = \vec{\mu}_D\ \vec{T}_D \vec{f}_D \quad \vec{f}_D \cap \vec{f} = \emptyset \qquad\qquad a \notin \textit{fn}(M)}{\Delta \vdash \text{class } c{<}\vec{\alpha} : \vec{P}{>}{\triangleleft}D\{\vec{\mu}\ \vec{T}\ \vec{f};\ \mathscr{M}\}[\Phi]}$$

$$\frac{\begin{array}{l}\Delta, \vec{\alpha} : \vec{P} \vdash D, \vec{T} \quad \Delta, \vec{\alpha} : \vec{P}, a : \text{Prin}, \text{this} : c{<}\vec{\alpha}{>}, a = \text{this.loc}, \square_a\text{mp}(\text{this}) \vdash_a^{|a} \Phi : \text{Pred Pure } c \\ \Delta \vdash \mathscr{M}\ in\ c{<}\vec{\alpha} : \vec{P}{>}{\triangleleft}D \quad \textit{fields}(D) = \vec{\mu}_D\ \vec{T}_D \vec{f}_D \quad \vec{f}_D \cap \vec{f} = \emptyset \qquad\qquad a \notin \textit{fn}(M)\end{array}}{\Delta \vdash \text{class } c{<}\vec{\alpha} : \vec{P}{>}{\triangleleft}D\{\vec{\mu}\ \vec{T}\ \vec{f};\ \mathscr{M}\}[\Phi]}$$

$$\frac{\begin{array}{l}\Delta, \vec{\alpha} : \vec{P}, \vec{\beta} : \vec{Q} \vdash S, \vec{T} \quad \vdash S' <: S \quad \vdash <\vec{\beta}>S(\vec{T})\ \textit{overrides } D.\ell \\ \Delta, \vec{\alpha} : \vec{P}, \vec{\beta} : \vec{Q}, \vec{x} : \vec{T}, a : \text{Prin}, \text{this} : c{<}\vec{\alpha}{>}, a = \text{this.loc}, \square_a(\text{mp}(\text{this}) \wedge \text{mp}(\vec{x})), \\ \qquad\qquad\qquad\qquad\qquad\qquad \text{caller} : \text{Prin}, \square_{\text{caller}}\text{mp}(\vec{x}) \vdash_a^{\text{caller}} M : S'\ \rho\ c \quad a \notin \textit{fn}(M)\end{array}}{\Delta \vdash <\vec{\beta} : \vec{Q}>S\ \ell\,(\vec{T}\ \vec{x})\,\{M\}\ in\ c{<}\vec{\alpha} : \vec{P}{>}{\triangleleft}D}$$

The judgment for values requires that well-formed objects satisfy their class invariants. In addition, the object value, as well as the objects held in its public fields must be permitted at the given location.

Well-Formed Values and Terms $(\Delta \vDash_a V : \mathcal{T})$ $(\Delta \vDash_a^{a'} M : \mathcal{T} \, \rho \, d)$ $(\rho ::= \mathsf{Pure} \mid \mathsf{Impure})$

$$\frac{\Delta \ni b : \mathsf{Prin}}{\Delta \vDash_a b : \mathsf{Prin}} \quad \frac{\Delta \ni x : T \quad \Delta \vDash \Box_a \mathrm{mp}(x)}{\Delta \vDash_a x : T} \quad \frac{\Delta \ni p : C \quad \Delta \vDash \Box_a \mathrm{mp}(p)}{\Delta \vDash_a p : C} \qquad \frac{}{\Delta \vDash_a \mathsf{unit} : \mathsf{Unit}}$$

$$\frac{\Delta \ni \alpha : \mathsf{Pred}(\vec{\mathscr{T}})}{\Delta \vDash_a \alpha : \mathsf{Pred}(\vec{\mathscr{T}})} \quad \frac{arity(\gamma) = \vec{\mathscr{T}}}{\Delta \vDash_a \gamma : \mathsf{Pred}(\vec{\mathscr{T}})} \quad \frac{\Delta \vDash_a \phi : \mathsf{Pred}(\vec{\mathscr{T}}) \ \ \Delta \vDash_a \vec{V} : \vec{\mathscr{T}}}{\Delta \vDash_a \phi(\vec{V}) : \mathsf{Pred}}$$

$$\frac{\Delta \vdash \diamond \quad \Delta \vDash_a p : C \quad fields(C) = \vec{\mu}\,\vec{T}\,\vec{f} \quad \Delta \vDash_a \vec{V} : \vec{T'} \quad \vdash \vec{T'} <: \vec{T}}{\Delta, env_a(p : C\{\vec{f} = \vec{V}\}) \vDash effect(C)\{^p/_{\mathsf{this}}\}}{\Delta \vDash_a^{a'} p : C\{\vec{f} = \vec{V}\} : \mathsf{Proc}\, \rho\, d}$$

$$\frac{\Delta \vdash \diamond \quad \Delta \vdash C \quad fields(C) = \vec{\mu}\,\vec{T}\,\vec{f} \quad \Delta \vDash_a \vec{V} : \vec{T'} \quad \vdash \vec{T'} <: \vec{T}}{\Delta, env_a(x : C\{\vec{f} = \vec{V}\}) \vDash effect(C)\{^x/_{\mathsf{this}}\} \quad \Delta, x : C, env_a(x : C\{\vec{f} = \vec{V}\}) \vDash_a^{a'} M : \mathscr{T}\,\rho\,d}{\Delta \vDash_a^{a'} \mathsf{let}\, x = \mathsf{new}\, C(\vec{V}); M : \mathscr{T}\ \mathsf{Impure}\, d}$$

$$\frac{\Delta \vdash \diamond \quad \Delta \vDash_a V : C \quad body(C.\ell) = \langle \vec{\beta} : \vec{Q} \rangle S(\vec{T}) \quad \Delta \vDash_a \vec{\phi} : \vec{Q} \quad \Delta \vDash_a \vec{W} : \vec{T'} \quad \vdash \vec{T'} <: \vec{T}\{\vec{\phi}/\vec{\beta}\}}{\Delta, b : \mathsf{Prin}, b = V.\mathsf{loc} \vDash \Box_b \mathrm{mp}(\vec{W}) \quad b \notin dom(\Delta) \quad \Delta, x : S\{\vec{\phi}/\vec{\beta}\}, \Box_a \mathrm{mp}(x) \vDash_a^{a'} M : \mathscr{T}\,\rho\,d}{\Delta \vDash_a^{a'} \mathsf{let}\, x = V.\ell \langle \vec{\phi} \rangle(\vec{W}); M : \mathscr{T}\ \mathsf{Impure}\, d}$$

$$\frac{\Delta \vdash \diamond \quad \Delta \vDash_a V : d \langle \vec{\phi} \rangle \quad fields(d \langle \vec{\phi} \rangle) = \vec{\mu}\,\vec{T}\,\vec{f} \quad \mu_i = \mathsf{private\ mutable} \quad \Delta \vDash_a W : T' \quad \vdash T' <: T_i}{\Delta \vDash_a^{a'} V.f_i := W : \mathsf{Unit}\ \mathsf{Impure}\, d}$$

$$\frac{\Delta \vdash \diamond \quad \Delta \vDash_a V : c \langle \vec{\phi} \rangle \quad fields(c \langle \vec{\phi} \rangle) = \vec{\mu}\,\vec{T}\,\vec{f} \quad \text{If } \mu_i \ni \mathsf{private\ then}\, c = d}{\text{If } \mu_i \ni \mathsf{mutable\ then}\, \rho = \mathsf{Impure}}{\Delta \vDash_a^{a'} V.f_i : T_i\, \rho\, d}$$

$$\frac{\Delta \vdash \diamond \quad \Delta \vDash_a V : T \quad \Delta \vDash_a W : S \quad \Delta, V = W \vDash_a^{a'} M : \mathscr{T}\,\rho\,d \quad \Delta \vDash_a^{a'} N : \mathscr{T'}\,\rho\,d}{\text{Either } \vdash \mathscr{T'} <: \mathscr{T} = \mathscr{T''} \text{ or } \vdash \mathscr{T} <: \mathscr{T'} = \mathscr{T''}}{\Delta \vDash_a^{a'} \mathsf{if}\, V = W \mathsf{\ then\ } M \mathsf{\ else\ } N : \mathscr{T''}\,\rho\,d}$$

$$\frac{\Delta \vDash_a^a N : T\ \mathsf{Impure}\, d}{\Delta, env_a(N), x : T, \Box_a \mathrm{mp}(x) \vDash_a^{a'} M : \mathscr{T}\,\rho\,d}{\Delta \vDash_a^{a'} \mathsf{let}\, x = N; M : \mathscr{T}\ \mathsf{Impure}\, d}$$

$right(N) = N'$

$$\frac{\Delta \vDash_a^a N : T\,\rho\,d \quad \Delta, env_a(N) \vDash_a^a N' : T\ \mathsf{Pure}\, d}{\Delta, env_a(N), x : T, x = N', \Box_a \mathrm{mp}(x) \vDash_a^{a'} M : \mathscr{T}\,\rho\,d}{\Delta \vDash_a^{a'} \mathsf{let}\, x = N; M : \mathscr{T}\,\rho\,d}$$

$$\frac{\Delta, env_a(M) \vDash_a^{a''} N : \mathscr{T'}\,\rho\,d \quad \Delta, env_a(N) \vDash_a^{a'} M : \mathscr{T}\,\rho\,d}{\Delta \vDash_a^{a'} N \parallel M : \mathscr{T}\,\rho\,d} \qquad \frac{\Delta, p : C \vDash_a^{a'} M : \mathscr{T}\,\rho\,d}{\Delta \vDash_a^{a'} (\nu p : C)\, M : \mathscr{T}\,\rho\,d}$$

$$\frac{\Delta \vdash \diamond \quad \Delta \vDash_{a'} V : \mathscr{T}}{\Delta \vDash_a^{a'} V : \mathscr{T}\,\rho\,d} \quad \frac{\Delta \vdash \diamond \quad \Delta \vDash_a V : C}{\Delta \vDash_a^{a'} V.\mathsf{loc} : \mathsf{Prin}\,\rho\,d} \quad \frac{\Delta \vDash_a b : \mathsf{Prin} \quad \Delta \vDash_b^{b'} M : \mathscr{T}\,\rho\,c}{\Delta \vDash_a^{a'} b[M]_c^{b'} : \mathscr{T}\,\rho\,d}$$

The typing rules for terms are designed to establish several invariants, which we now discuss.

Well located. The rules for terms use the value judgment to ensure that value occurrences are available at a given location. The rule for located terms switches principals as expected.

Purity annotations. Field updates and constructor calls are impure because they mutate the heap. Field accesss to mutable fields are impure because they rely on mutable state. Method invocations are impure because they might not terminate. In all other cases,

the purity annotation is constructed inductively. For example, a let is pure only if both terms involved are pure. Similarly for concurrent composition and conditionals.

Equations. The rule for pure let terms uses the function *right*. Intuitively, for any term N, $right(N)$ returns the rightmost subterm of N after it has been rewritten to a normal form. Routine details are omitted. Conditionals and let expression on pure terms introduce equations to the environment. Equations are also generated by the rules for heap objects ($\eta : C\{\cdots\}$) and new.

Caller annotations. The caller annotation is carried inductively through all rules but two. In the rule for the concurrent composition, only the right term is constrained; the value of the left term is ignored. The purpose of the caller annotation is revealed by the rule for values which appear as terms — these are the return values. The rule ensures that the caller principal is permitted to have a reference to the value.

Checking effects and the $\mathrm{mp}(\cdot)$ *predicate.* The rule for new illustrates the methodology. (The rule for heap objects enforces similar proof obligations.) In this rule, the hypothesis for typing fields is standard. The lookup of the effect obligation via *effect*(C) yields a conjunction of the effects for this class and all its superclasses. The proof obligations ensure that the created object conforms to the class predicate, and that the reference and its public fields are permitted to be at the principal at which the object is located. The facts used to discharge this proof obligation are derived from the environment via Δ which accumulates the benefits derivable from the objects declared in the environment and the equations accumulated in the environment via lets and conditionals. The parameters to the constructor have to be available at the current location a.

In field update and lookup, the class annotation on the typing judgment is ensured to be the class of the object if the field is private.

In the rule for "generic" methods, we substitute concrete formulas for the logical variables being carried in the method definition. Since methods are executed at the location of the callee, we check to ensure that the location of the callee object possesses the right to hold references to the objects being passed in as actual parameters.

Conjoining specifications The rule for concurrent composition reflects the ideas from conjoining specifications of concurrent systems [3] — each component can assume the information exposed by the other component.

Results. An *initial program* is one that contains no dynamic constructs.

An *opponent class* is one whose effect is trivial, i.e., $\mathtt{tt}$. An *opponent program* is one that can be typed only allowing the constructor rule for opponent classes. In typing opponents, we allow the assumption $\forall \eta.\ \Box_\perp \mathrm{mp}(\eta)$. Thus opponents are typed using a restricted class table, but under a permissive policy. This permissive policy is essentially the same as standard object-oriented typing.

An opponent can instantiate opponent classes. By Principal Naturality, the opponent can unconstrainedly pass arguments or return results in method invocations. Thus, the opponent typability requirement in the following safety result means only that the opponent program is typable in the sense of classic object-oriented programming.

Recall that a frame is a term of the form $a\,[M]_c^b$.

Definition 1. A term M is *safe for* Δ if whenever $M \to^* N$, $N \equiv \mathbb{E}[a\,[N']_c^b]$, N' contains no frames and $p \in fn(N)$ then $\Delta, env_\perp(N) \vDash \Box_a \mathrm{mp}(p)$. $\hfill\square$

```
class User {
  private final Compiler compiler;
  private FileOutputStream fDebug;
  void action () {
    ...
    this.compiler.compile (this.fDebug, ...); // Invoke with current fDebug.
  }
  void setDebug (FileOutputStream fDebug) {
    this.fDebug = fDebug;
  }
}[ ∀o : FileOutputStream.□_this.loc mp(o) ⇒ □_this.compiler.loc mp(o) ]

class Compiler {
  private final FileOutputStream fStats;
  public void compile (FileOutputStream fDebug, String source) {
    ...
    this.fStats.write (...); // Write statistics to fStats.
    fDebug.write (...); // Write debugging output to fDebug.
  }
}
```

Fig. 1. User and Compiler Code

Proposition 2. Suppose that $\Delta \vdash_\top M$ and $\Delta, env_\perp(M), (\forall \eta . \, \square_\perp \mathtt{mp}(\eta)) \vdash_\top N$ for an initial opponent program N. Then $N \mathbin{\|} M$ is safe for Δ. $\qquad\square$

The safety result ensures that well-typed trustworthy programs are safe when combined with arbitrary (typed but untrustworthy) opponents.

5 The Confused Deputy

In this section, we examine how to typecheck code that addresses the Confused Deputy problem using object references as capabilities.

Hardy [22] discusses a system with a compiler invoked by a user. The compiler writes two files, in addition to any generated code. The first is a statistics file. The name of the statistics file is hardcoded into the compiler, and the compiler is explicitly granted permission to write to that file. The second is a debugging file, chosen by the user. In order to write to the user's choice of debugging file, the compiler must be granted a broad permission. Hardy describes an occasion when a user selected a sensitive file— subsequently overwritten by the compiler—and dubs the compiler a Confused Deputy.

Hardy's solution requires the user to obtain a capability to write to the debugging file, and to send that capability to the compiler. The compiler can use the user's capability to write to the debugging file.

Modeling. We model Hardy's solution using the code in Figure 1. Following the object references as capabilities paradigm, the capabilities to write to files are represented by a FileOutputStream class (as in Java).

The `User` class invokes a compiler, passing the `FileOutputStream` contents of the `fDebug` field. The `User` class allows its `fDebug` field to be updated via a method `setDebug`—we examine the typing consequences below.

The `Compiler` class must be initialized with a final `FileOutputStream` field `fStats` at construction. When it compiles, it uses its own `fStats` field and the `fDebug` method parameter supplied by the caller to write to the statistics and debugging files.

Controlling capabilities. The capability solution improves upon the Confused Deputy situation, with respect to the principle of least privilege, because the compiler lacks the broad permission to write to many files in the object/capability solution. Hardy observes that achieving a comparable system with a traditional access control policy for the compiler is challenging, e.g., because the compiler may be invoked by different users with access to different files.

However, the capability solution is not entirely satisfactory. As discussed in the introduction, an untrustworthy compiler might forward capabilities that it receives to objects at different locations (principals).

Type assignment. We now consider how to typecheck the code of Figure 1 in a way that allows the `Compiler` to receive the `fDebug` object reference but not forward it to another location. We omit discussion of the source code given to the compiler, and any executable output, for reasons of space.

The typing of the compiler's use of `FileOutputStream` references is straightforward. The compiler receives permission to possess the field `fStats` implicitly. More generally, our type system implicitly allows every object to access its own fields. On the other hand, code that constructs a `Compiler` instance is responsible for ensuring that the chosen location of the compiler is able to possess `fStats`. That is, if a newly created `Compiler` instance is referenced via `c`, then the proof obligation is $\Box_{\texttt{c.loc}}\texttt{mp}(\texttt{c.fStats})$. Similarly, our type system implicitly grants the compiler permission to possess the method parameter `fDebug`, and the obligation lies with the caller to ensure that the location of the callee may possess `fDebug`.

Typechecking the body of the `compile` method does not introduce proof obligations involving `fStats` or `fDebug`, because those values are passed as `this`, and the type system automatically validates $(\forall \texttt{o}.\Box_{\texttt{o.loc}}\texttt{mp}(\texttt{o}))$. That is, a location may possess a reference to any object stored at that location.

The user has a more interesting policy, because it has to permit forwarding of `fDebug` to the compiler. The form of the policy hinges upon the mutability of the `fDebug` field. For example, if `fDebug` was a final field, it could be referred to in the class invariant, e.g., with form:

$$\Box_{\texttt{this.compiler.loc}}\Box_{\texttt{this.loc}}\texttt{mp}(\texttt{this.fDebug})$$

which entails:

$$(\Box_{\texttt{this.loc}}\texttt{mp}(\texttt{this.fDebug})) \wedge (\Box_{\texttt{this.compiler.loc}}\texttt{mp}(\texttt{this.fDebug}))$$

To demonstrate a more flexible alternative, we chose to make `fDebug` non-final (mutable) in the code of Figure 1. Since we can no longer refer to the field in a class invariant, we instead state that all `FileOutputStream` references that may be possessed at the

location of `User` may also be possessed at the location of the corresponding compiler. This class invariant is written:

$$\forall o : \texttt{FileOutputStream}.\Box_{\texttt{this.loc}}\texttt{mp}(o) \Rightarrow \Box_{\texttt{this.compiler.loc}}\texttt{mp}(o)$$

With this class invariant, typechecking justifies forwarding of `this.fDebug` to `this.compiler` using implication together with the facts that: (1) `this.loc` may possess `this.fDebug` (the location of an object implicitly possesses its fields); and (2) `this.fDebug` is declared to be an instance of `FileOutputStream`.

Finally, code that constructs an instance of `User` has an obligation to show that the associated `Compiler` instance is usable with any `FileOutputStream` object reference that the `User` receives.

6   Conclusion

The control of the possession and transmission of secrets is a recurrent theme in security literature and practice. The policies on possession in this paper describe an upper bound on the principals who can possess a secret. We describe a static analysis to ensure programs in a distributed object-oriented language comply with such policies. Our static analysis takes the form of a refinement type system, based on indexed necessity modalities from intuitionist S4, for an object calculus with locations. The safety result ensures that in the configurations that arise from the execution of well-typed programs, objects are only accessible to principals who are permitted to do so by the system policy, even in the presence of attackers who try to subvert the policies by inserting malicious objects and code into the system. Our results suggest that type systems are a practical tool to debug secrecy errors in the design of user-defined APIs in distributed systems.

References

1. Abadi, M.: Access control in a core calculus of dependency. ENTCS. 172, 5–31 (2007)
2. Abadi, M., Gordon, A.D.: A calculus for cryptographic protocols: The spi calculus. Information and Computation 148, 36–47 (1999)
3. Abadi, M., Lamport, L.: Conjoining specifications. ACM Trans. Program. Lang. Syst. 17(3), 507–535 (1995)
4. Abadi, M.: Secrecy by typing in security protocols. Journal of the ACM 46, 611–638 (1998)
5. Anderson, M., Pose, R.D., Wallace, C.S.: A password-capability system. Comput. J. 29(1), 1–8 (1986)
6. Bierman, G.M., de Paiva, V.C.V.: On an intuitionistic modal logic. Studia Logica 65 (2001)
7. Cardelli, L.: A language with distributed scope. POPL. pp. 286–297 (1995)
8. Castellani, I.: Process algebras with localities. Handbook of Process Algebra, chap. 15, pp. 945–1045 (2001)
9. Cirillo, A., Jagadeesan, R., Pitcher, C., Riely, J.: TAPIDO: Trust and authorization via provenance and integrity in distributed objects. ESOP. pp. 208–223 (2008)

10. DeYoung, H., , Pfenning, F.: Reasoning about the consequences of authorization policies in a linear epistemic logic. Tech. Rep. 1213, CMU (2009)
11. Drossopoulou, S.: Ten years of ownership types or the benefits of putting objects into boxes (2008), invited talk at BCS. Talk available at http://www.doc.ic.ac.uk/~scd/BCS.pdf
12. E: Open source distributed capabilities, http://www.erights.org
13. Feil, R., Nyffenegger, L.: Evolution of cross site request forgery attacks. Journal in Computer Virology 4(1), 61–71 (Nov 2007)
14. Fournet, C., Gordon, A.D., Maffeis, S.: A type discipline for authorization in distributed systems. CSF (2007)
15. Fournet, C., Gordon, A.D., Maffeis, S.: A type discipline for authorization policies. ACM Trans. Program. Lang. Syst. 29(5) (2007)
16. Freeman, T., Pfenning, F.: Refinement types for ML. pp. 268–277. PLDI '91, ACM, New York, NY, USA (1991)
17. Garg, D., Bauer, L., Bowers, K.D., Pfenning, F., Reiter, M.K.: A linear logic of authorization and knowledge. ESORICS. LNCS, vol. 4189, pp. 297–312 (2006)
18. Garg, D., Pfenning, F.: Non-interference in constructive authorization logic. CSFW. pp. 283–296 (2006)
19. Gong, L., Mueller, M., Prafullch, H.: Going beyond the sandbox: An overview of the new security architecture in the Java Development Kit 1.2. USENIX Symposium on Internet Technologies and Systems. pp. 103–112 (1997)
20. Gordon, A.D., Hankin, P.D.: A concurrent object calculus: Reduction and typing. Proceedings HLCL'98 (1998)
21. Gordon, A.D., Jeffrey, A.: Authenticity by typing for security protocols. Journal of Computer Security 11(4), 451–520 (2003)
22. Hardy, N.: The confused deputy: (or why capabilities might have been invented). SIGOPS Oper. Syst. Rev. 22, 36–38 (October 1988)
23. Hardy, N.: KeyKOS architecture. SIGOPS Oper. Syst. Rev. 19, 8–25 (October 1985)
24. Hennessy, M., Riely, J.: Resource access control in systems of mobile agents. Information and Computation 173, 2002 (1998)
25. Jagadeesan, R., Pitcher, C., Riely, J.: Non interference for intuitionist necessity. Tech. Rep. 12-003, School of Computing, DePaul University (2012)
26. Maffeis, S., Mitchell, J.C., Taly, A.: Object capabilities and isolation of untrusted web applications. IEEE Symposium on Security and Privacy. pp. 125–140 (2010)
27. Nicola, R.D., Ferrari, G., Pugliese, R.: Klaim: a kernel language for agents interaction and mobility. IEEE Transactions on Software Engineering 24, 315–330 (1997)
28. Pfenning, F., Wong, H.C.: On a modal λ-calculus for S4. Proceedings of MFOS. New Orleans, Louisiana (Mar 1995), *ENTCS*, Volume 1, Elsevier
29. Shapiro, J.S., Smith, J.M., Farber, D.J.: EROS: a fast capability system. SIGOPS Oper. Syst. Rev. 33, 170–185 (Dec 1999)