

Towards Capability Policy Specification and Verification

Abstract. The object-capability model is a de-facto industry standard widely adopted for the implementation of security policies for web-based software. Unfortunately, code written using capabilities tends to concentrate on the low-level *mechanism* rather than the high-level *policy*, and the parts implementing the policy are tangled with parts implementing the functionality.

In this paper we argue that the policies intended by programs using object capabilities should be made explicit and written separately from the code implementing them. We also argue that the specification of such capability policies requires concepts that go beyond the features of current specification languages. Moreover, we argue that we need methodologies with which to prove that programs adhere to their capability policies as specified.

We sketch such a capability policy specification language, and propose capability triples, which are a generalisation of Hoare logic triples. Capability triples consist of conditions, code and conclusions, but interestingly they allow the conditions as well as the conclusion to talk about the state before and after execution of the code, they allow the code to be existentially or universally quantified, and their interpretation quantifies over all modules extending the current module.

Based on the famous mint example [22], we outline its capability policy specifications, and demonstrate how we can reason that it satisfies the capability policies. Interestingly, the reasoning makes use of restrictions imposed by the type system, such as `final` and `private`.

1 Introduction

Capabilities — unforgeable authentication tokens — have been used to provide security and task separation on multi-user machines since the 60s [7], *e.g.* PDP-1, operating systems *e.g.* CAL-TSS [16], and the CAP computer and operating system [47]. In capability-based security resources can only be accessed via capabilities: possessing a capability gives the right to access the resource represented by that capability.

Object capabilities [24] apply the concept of capability to object-oriented programming languages. In an object capability system, an object is a capability for the services the object provides: any part of a program that has a reference to an object can always use all the services of that object. To restrict authority over an object, one can create, instead of “naked” reference, an intermediate object which offers only restricted services on the original object.

Object capabilities afford simpler and more fine-grained protection than privilege levels (as in Unix), static types, ad-hoc dynamic security managers (as in Java or JSand [1]), or state-machine-based event monitoring [2]. On the other hand, object capability systems are only secure as long as trusted capabilities (that is, trusted objects) are never leaked to untrusted code. Object capabilities have been adopted in several programming languages [26, 22, 45] and are increasingly used for the provision of security in web-programming in industry [27, 46, 39].

On a different development strand to object capabilities, and with the aim to restrict access across code, programming languages adopted features like packages and opaque types, `const` or fields, `private` and `protected` members, or `final` classes [48, 40]. More advanced features, such as ownership types [6], restrict access to different parts of the heap. Such features do not introduce new behaviour into the language, but restrict the set of legal programs, hence we call them here *restrictive*; they usually are part of statically typed languages.

The key problem with object capability programming as practiced today is that — because capabilities are just objects — code manipulating capabilities is tangled together with code supporting the functional behaviour of the program. The actual security policies enforced by a program are *implicit*, scattered throughout the program’s code. Any part of a program that uses an object may (by oversight, error, or fraud) hand that object to an untrusted part of the program, giving the untrusted code access to all the services provided by that object. This makes it difficult to determine what security properties are guaranteed by a given program, and as a result, programs are difficult to understand, validate, and maintain.

We argue that capability policies should be specified separately from the program implementing them. And that the specification of capability policies requires features that go beyond what is available in current specification languages. Namely, capability policies are *program centred*, *fine grained*, *open* in the sense that they specify aspects of the behaviour of all possible extensions of a program, and have *deny* elements, which require that certain effects may only take place if the originating code or the runtime context satisfy some conditions. In [10] the authors anticipated expressing such policies through extensions of temporal logics.

In this paper we propose that capability policies can be specified through a generalisation of Hoare triples, which we call *capability triples*. Capability triples allow explicit quantification over modules, and over code — and thus reflect the open nature of capability policies. They also allow premises to talk about properties observed in the configuration *before* as well as *after* execution of the code — and thus they reflect the deny elements of the policies. Finally, the triples not only talk about properties of the heap (*e.g.*, the balance of an account), but also about accessibility and tracing properties (*e.g.*, the current stack frame has a path to object `o` which involves only `public` fields) — thus they reflect the program centred elements of capability policies.

To make the meaning of such policies precise, we define some concepts relating to program execution, accessibility, reachability and tracing. We give the definition of their manifestation for a Joe-E/Java subset, but we believe that these concepts have a manifestation in most object-oriented programming languages. We use the Mint example [22] to illustrate our ideas, and gave precise meaning to five out of the six policies proposed informally in that paper. We were surprised by the many different interpretations we found for the policies. We then give a semi-formal argument showing that certain code adheres to the capability policy. In doing so, we make heavy use of restrictive language features; this was surprising for us, since in traditional program verification, but also in verification of refinement properties, restrictive features have played no role.

The rest of the paper is organised as follows: Sect. 2 presents the Mint [22] as an example of object-capability programming, implemented in Joe-E/Java. Based on

that example, sect. 3 distills the characteristics of capability policies. Sect 5 informally explores reasoning about capability policies, and the use restrictive language features. Sect. 6 discusses further useful policies, not listed in [22]. Section 7 surveys related work. Sect. 8 concludes.

2 Object-Capability Example

We use as running example a system for electronic money as proposed in [26]. This example allows for mints with electronic money, purses held within mints, and transfers of funds between purses. The *currency of a mint* is the sum of the balances of all purses created by that mint. Purses trust the mint to which they belong, and programs using the money system trust their purses (and thus the mint). Crucially, separate users of the money system *do not* trust each other.

The standard presentation of the mint example defines six capability policies, which we repeat here, as they were described in [26]:

- Pol.1** With two purses of the same mint, one can transfer money between them.
- Pol.2** Only someone with the mint of a given currency can violate conservation of that currency.
- Pol.3** The mint can only inflate its own currency.
- Pol.4** No one can affect the balance of a purse they don't have.
- Pol.5** Balances are always non-negative integers.
- Pol.6** A reported successful deposit can be trusted as much as one trusts the purse one is depositing into.

An immediate consequence of these policies is that the mint capability gives its holder the ability to subvert the currency system by “printing money”. This means that while purse capabilities may safely be passed around the system, the mint capability must be carefully protected.

There is also an implicit assumption that no purses are destroyed. This assumption is necessary because destruction of a purse would decrease the currency of a mint, in opposition to **Pol.3**. The implication of this assumption is that there will be no explicit destruction of purses, but also no garbage collection of purses.

Several different implementations have been proposed for the mint. Fig.2 contains an implementation in Joe-E [22], a capability-oriented subset of Java, which restricts static variables and reflection.

In the Joe-E version, the policies are adhered to through the interplay of appropriate actions in the method bodies (*e.g.* the `check` in line 17), with the use of Java's *restrictive* language features (`private` members are visible to the same class only; `final` fields cannot be changed after initialisation; and `final` classes cannot be extended). The code concerned with the functional behaviour is tangled with the code implementing the policy (*e.g.* in `deposit`, line 19 is concerned with the functionality, while line 17 is concerned with **Pol.2**). The implementation of *one* policy is scattered throughout the code, and may use explicit runtime tests, as well as restrictive elements (*e.g.* **Pol.2** is implemented through a check in line 17, the `private` and `final` annotations, and the initialisations in lines 9 and 13). Note that an apparently innocuous change to this code

```

1 public final class Mint {          } // the Mint capability
2
3 public final class Purse {
4     private final Mint mint;
5     private long balance;
6
7     public Purse(Mint mint, long balance) {
8         if (balance<0){throw new IllegalArgumentException();};
9         this.mint = mint;  this.balance = balance;
10    }
11
12    public Purse(Purse pts) {
13        mint = pts.mint;  balance = 0;
14    }
15
16    public void deposit(Purse prs, long amnt) {
17        if ( mint!=prs.mint || amnt>prs.balance || amnt+balance
18            <0 )
19            { throw new IllegalArgumentException(); };
20        prs.balance -= amnt;  balance += amnt;  }
21    }

```

— such as a public `getMint` accessor that returned a purse’s `mint` — would be enough to leak the `mint` to untrusted code, destroying the security of the whole system.

Both the Joe-E and the E versions suffer from tangling of policy with functionality, and scattering of policy implementations [10].

3 Capability Policies

We use the term *capability policy* as a description of how capabilities are intended to be used: which objects are trusted, which are untrusted, and precisely which capabilities can be accessed by which object. A key feature of capability systems is the *principle of least authority* — an object should only be able to access the capabilities (i.e. the other objects) that it needs in order to function correctly: even a trusted object should not have access to all the capabilities (objects) in the system [37, 29, 47]. A range of object capability policies are discernible from the literature [24, 26, 25].

Capability policies generally have the following characteristics:

- They are *program centred*: they talk about properties of uses of programs rather than properties of execution of protocols.
- They are *fine-grained*: they can talk about *individual objects*, while *coarse-grained* policies only talk about large components such as `System` or `DOM`.

- They are *open*. *Open* requirements must be satisfied for any use of the code *extended in any possible manner*, while *closed* requirements need only be satisfied for any use of code itself.
- They have *rely* as well as *deny* elements. Rely elements essentially promise that execution of a *given* code tin a state satisfying a given pre-condition will reach another state which satisfies some post-condition [13]. Deny elements promise that if an execution reaches a certain state, or changes state in a certain way, or accesses some program entity, then the code must satisfy some given properties. In other words, rely policies are about *sufficient* conditions, while deny policies are about *necessary* conditions.

None of the terms above are standard; we coined them to delineate our ideas. The mint's policies are capability policies; namely::

- They are program centred, since they refer to actual programs.
- They are fine grained, as they refer to the *individual* purses and mints;
- Even though not explicitly stated there, the policy in [26] is expected to be open; any expansion of the code (through dynamic loading, subclassing, mashups *e.t.c.*) should satisfy the requirements.
- They contain *rely* as well as *deny* elements:
 - **Pol.1** is a *rely* requirement, expressible through classic pre- and post- conditions: namely, execution of `deposit` in a state satisfying the pre-condition that the where the two purses belong to the same mint leads to a state satisfying the post-condition that where the money has been transferred.
 - **Pol.2** is a *deny* requirement; it says that a currency may be changed by some code only if the code contained a function call executed by the mint owning the currency. **Pol.3** is another *deny* requirement; it says that if the currency should change, then it increases. **Pol.4** is also a *deny* requirement, preventing objects that cannot access a purse from modifying its value.
 - **Pol.5** can be understood as an object invariant, requiring purses' balances to always be positive, but also as a *deny* requirement which requires that any code preserves this property.

Open policies are central for JavaScript security, which requires that in any mashup, untrusted code cannot access the trusted security-critical resources of the execution environment (*e.g.* the DOM), nor interfere with the execution of any other component [19, 15]. These works usually implement coarse-grained, fixed-in-advance policies. SecureJS [44] leverages local scoping (a restrictive feature) to prove fine-grained confinement (*e.g.* `decr` is confined), but cannot express the high-level policies (*e.g.* currency cannot be affected). JSand [1] uses Secure ECMAScript and proxies (other restrictive features) to isolate the DOM, and then ensures access to that DOM proxy are mediated by dynamically checking security policies expressed as JavaScript predicates.

Deny policies are related to *deny-guarantee* specifications [9] which can forbid given locations from being modified by the current, or the other threads. Deny policies typically apply throughout program execution, rather than during specific functions, and refer to any properties of the program (*e.g.* the currency), rather than specific locations.

Deny policies are also related to *correspondence assertions* [49, 12], which require principals reaching a certain point in a protocol to be preceded by other principals reaching corresponding points. Recently, correspondence assertions have been adapted to refer to program state, and thus can prove that the code adheres to security, authentication, and privacy policies [3]: functions are annotated by *refinement types* that require that the function is only called if its arguments satisfy the type’s conditions.

Deny policies go further than correspondence assertions in the following significant ways:

- They support *implicit* properties, *i.e.* properties that depend on state reachable from more than one object, perhaps quantifying over the complete heap, or even on the history of execution. In our example, the currency is the sum of the balance of all purses from the same mint, and therefore is an implicit property.
- They are *pervasive*, *i.e.* they are not attached to one function, and may be affected by several different methods. For example, the currency may be affected by the creation of purses and the payments.
- They are *persistent*, *i.e.* they allow the comparison of properties of the state at different times in execution. For example, **Pol.3** compares the currency between any two times in execution.

Deny policies could be transformed into equivalent refinement types; however, the transformation would not be trivial, and the resulting policies would not be open (because the refinement types cannot prevent the addition of functions which break the requirements), and less abstract (how would refinement types express that the currency can only grow?).

4 Towards Capability Policy Specification Languages

In this section we sketch a capability policy specification language. Our specification language is a generalisation of Hoare triples. We allow quantification over any code, and the extension of modules by any further modules — thus we deal with the open nature of policies. Moreover, in the precondition we allow the comparison of the state before and after execution — thus we deal with the deny part of policies, and can specify necessary, instead of sufficient, conditions.

We start by introducing the concepts necessary to give precise meaning to policies (sect 4.1), then use these to describe the meaning of the policies (sect 4.2), and then outline a capability policy specification language.

4.1 Concepts for Capability Policy Specifications

Capability policies should be expressible regardless of the particular program implementing them and the particular programming language the program was written in, provided that they (the policies) are expressed in terms of concepts which have their own individual manifestation in the particular programming language. For example, the concept of runtime configuration is manifest in programming languages as diverse as Cobol, Java, or Python, even though its manifestation differs.

In this section we introduce concepts used to define the meaning of policies. We describe their manifestation in a “capability-safe” Java-subset, which we call $\mathcal{L}ng_{cj}$. We give a precise definition of $\mathcal{L}ng_{cj}$ and all concepts in appendix A, while here we bring out the most salient issues.

Modules and Linking In order to reflect the open nature of capability policies, we need a handle on programs, and on the extension of programs (through subclasses, mashups, imports etc). For this we use *modules*, M , to describe parts of programs, and $*$ to describe the combination of two programs into one larger program.

A *module* is essentially a collection of definitions. In ML, modules could be structures and functors, in Java they could be class definitions and packages. In $\mathcal{L}ng_{cj}$, modules are class declarations (*c.f.* App. A.1). Thus, the class definitions M_{Mint} and M_{Purse} from Fig. 1 are modules.

Remember that adherence to policies often relies on the correct use of restrictive features. Therefore, in $\mathcal{L}ng_{cj}$ we support annotations like `private` and `final`, and the type rules forbid access to private fields or methods outside their classes, forbid extensions of final classes and redefinition of method declared as final in the superclass, and forbid assignment to final fields outside their contractor (*c.f.* App. A.2).

The operator $*$ is a linking operation. In general, linking performs some compatibility checks, and therefore is only partially defined. For example, because the field `balance` is private, $M_{\text{Purse}} * M'$ would be undefined, if M' contained somewhere the expression `newPurse(-, -).balance`. In App. A.3 the operation $*$ is only defined if it gives rise to a well-formed module.

Modules are not directly executable, but are necessary for the execution of *code snippets*. In $\mathcal{L}ng_{cj}$ code snippets are expressions, *c.f.* App. A.1. We use variables *code*, *code'* to range over code snippets.

Runtime Execution Context and Expression Evaluation Execution contexts, *ctxt*, stand for runtime contexts in which code snippets are executed. In the case of Java, a context is a stack frame (i.e. a mapping from variable names to object addresses or scalar values), and a heap (i.e. a mapping from object addresses to objects), *c.f.* App. A.4.

Execution of a code snippet *code* for a module M takes a context *ctxt* and returns a value v and a new context *ctxt'*. We describe this through a large step semantics, of the shape $M, ctxt, code \rightsquigarrow ctxt', v'$. We define this relation for $\mathcal{L}ng_{cj}$ in App. A.4.

Arising Runtime Configurations When considering adherence to policies, we it is essential to examine only those runtime configurations (*i.e.*, context and code pairs) which may arise through the execution of the given modules. For example, if we allowed *any* well-formed configuration (well-formed in the sense of the type system) to be examined, then we would be unable to ascertain that **Pol.5**, which guarantees that balances are always positive, is adhered to.

Thus, $\text{Arising}(M)$ is the set of runtime configurations which may occur during execution of some initial configuration $(ctxt_0, expr_0)$. These concepts are defined in App. A.6; the sets are always defined, even if $M, ctxt_0, expr_0$ does not terminate.

Objects Accessible from a Context; Objects used during execution Object capabilities draw on the requirement that some heap entities might not be accessible from all the stack frames. For example, a context containing and having access to a purse object `prs`, will also contain the mint object `mnt` which created that purse, but it need not have access to that `mnt` object, since the field `mint` is private in `Purse`.

We therefore distinguish between the sets $\mathcal{AccAll}(M, ctxt)$ resp. $\mathcal{AccPub}(M, ctxt)$ which return the set of objects which are accessible from the frame in $ctxt$ through *any* path resp. through paths which do not visit private fields, unless they belong to objects of the same class as `this`. The definitions appear in App. A.8 and A.8.

We use the notation $z :_{ctxt} c$ to indicate that z is the name of an object which exists in the heap of $ctxt$ and belongs to class c - no requirement that there should be a path from the frame to this object. The notation $z, z' :_{ctxt} c$ stands for $z :_{ctxt} c \wedge z' :_{ctxt} c$. Finally, we define $\mathcal{Used}(M, ctxt, code)$, as the set of all addresses used during execution of $code$ in the context $ctxt$ - for full definitions *c.f.* App. A.7.

4.2 Giving precise meaning to the policies

Armed with the concepts from section 4.1, we turn our attention to the precise meaning of the six policies. An important aspect of our approach is that we quantify over modules, extensions to modules, over the code being executed, and over execution contexts.

We discuss the policies in order of increasing complexity of their specification. We were surprised how many different interpretations we uncovered while developing this part of the work.

The fifth policy **Pol.5**, “Balances are always non-negative integers”, is akin to object invariants [23, 31, 42]. We can express the policy directly, by requiring that a module M satisfies **Pol.5**, if for all M' legal extensions of M , and runtime configurations $(ctxt, expr)$ arising through execution of the augmented program $M * M'$, the configuration $(ctxt, expr)$ will produce a context $ctxt'$, in which the balance is positive.

Module M satisfies policy **Pol.5**

iff

$$\forall M'. \forall (code, ctxt) \in \mathcal{Arising}(M * M'), \text{ and } prs :_{ctxt} \text{Purse}: \\ M * M', ctxt, code \rightsquigarrow ctxt', v \Rightarrow prs.balance_{ctxt'} \geq 0$$

The subscripts in the path expression `prs.balance` indicate whether the path is looked up in the old, or the new context (`prs.balancectxt` vs `prs.balancectxt'`).

Execution uses the extended module $M * M'$, where M' is universally quantified. This reflects the open nature of capability policies. It is essential to allow $M * M'$ in the execution, because this supports calling methods and accessing fields defined in M but also in M' .

We implicitly assume that quantification over linked modules only quantifies over those where the $*$ operator is defined. Thus,

$\forall M'. \forall (ctx, code) \in \mathcal{A}rising(M * M')$
 implicitly stands for
 $\forall M' s.t. M * M'$ is defined, $\forall code, ctx, s.t. (ctx, code) \in \mathcal{A}rising(M * M')$.

The third policy **Pol.3**, stating “The mint can only inflate its own currency”, could mean that the currency of a mint never decreases, or that the mint cannot affect the currency of a different mint. As the second meaning is a corollary of **Pol.2**, we consider the first meaning only.

We first define the *currency* of a mint, which is an implicit assertion in the sense from section 3, as it quantifies over all objects of the heap:

$$Currency_{ctx}(\text{mnt}) = \sum_{p \in Ps(\text{mnt})_{ctx}} p.balance_{ctx}$$

where $Ps(\text{mnt})_{ctx} = \{p \mid p :_{ctx} \text{Purse} \wedge p.mint_{ctx} = \text{mnt}_{ctx}\}$

We can now express the policy.

$$\begin{aligned} & \text{Module } M \text{ satisfies policy } \mathbf{Pol.3} \\ & \text{iff} \\ & \forall M'. \forall (ctx, code) \in \mathcal{A}rising(M * M'), \text{ and } \text{mnt} :_{ctx} \text{Mint} \\ & \quad M * M', ctx, code \rightsquigarrow ctx', v \\ & \quad \Rightarrow \\ & \quad Currency(\text{mnt})_{ctx} \leq Currency(\text{mnt})_{ctx'} \end{aligned}$$

Pol.3 describes a monotonic property, and is therefore related to history invariants [17]. However, it differs from history invariants through its open nature, and hence the quantification over M' . Note that in the conclusion we talk about the values of functions in the old context (*i.e.* $Currency(\text{mnt})_{ctx}$) as well as those in the new context (*i.e.* $Currency(\text{mnt})_{ctx'}$).

The first policy. **Pol.1** states “With two purses of the same mint, one can transfer money between them”. It can be understood to mean that if p_1 and p_2 are purses of the same mint, then the method call $p_1.deposit(p_2, m)$ will transfer the money. Therefore, it can be specified through a Hoare Logic triple as follows:

$$\begin{aligned} & \{ p_1.mint = p_2.mint \wedge p_1.amount = k_1 \wedge p_2.amount = k_2 + m \} \\ & \quad p_1.deposit(p_2, m) \\ & \{ p_1.amount = k_1 + m \wedge p_2.amount = k_2 \} \end{aligned}$$

Using a similar notation to that we used so far, we can write **Pol.1** as:

$$\begin{aligned} & \text{Module } M \text{ satisfies policy } \mathbf{Pol.1} \\ & \text{iff} \end{aligned}$$

$$\begin{aligned}
& \forall M'. \forall (ctxt, p1.deposit(p2, m)) \in \mathcal{A}rising(M * M'), \text{ with } p1, p2:_{ctxt} \text{Purse:} \\
& \quad p1.mint_{ctxt} = p2.mint_{ctxt} \wedge p2.balance_{ctxt} \geq m \\
& \quad \wedge M * M', ctxt, p1.deposit(p2) \rightsquigarrow ctxt', v \\
& \quad \Rightarrow \\
& \quad p1.balance_{ctxt'} = p1.balance_{ctxt} + m \wedge p2.balance_{ctxt'} = p2.balance_{ctxt} - m.
\end{aligned}$$

The above specification ranges over all module extensions, M' , and thus covers a larger set of runtime configurations than if it was expressed without the quantification over M' . Therefore it guarantees that the code M' can do nothing to break the behaviour of the `deposit` method from M , thus forcing the method `deposit` to be final, or a system with contracts which ensures any subclasses will satisfy the same contract.

The above is perhaps an over-specification, as it prescribes *how* the transfer is to take place by explicitly calling the `p1.deposit(p2)` method. Instead, we may want to only requiring that it should be *possible* for the transfer to take place, without discussing features of the program design. Therefore, we define a second, more general version of the policy, which only requires the existence of some code snippet that performs the transaction, and which applies existential qualification over the code:

$$\begin{aligned}
& \text{Module } M \text{ satisfies policy } \mathbf{Pol_1, vrs2} \\
& \quad \text{iff} \\
& \quad \forall (ctxt, -) \in \mathcal{A}rising(M). \forall p1, p2:_{ctxt} \text{Purse.} \\
& \quad p1.mint_{ctxt} = p2.mint_{ctxt} \wedge p2.balance \geq m \wedge p1, p2 \in \mathcal{A}ccPub(M, ctxt) \\
& \quad \Rightarrow \\
& \quad \exists \text{code, such that } M, ctxt, \text{code} \rightsquigarrow ctxt', v \wedge \\
& \quad p1.balance_{ctxt'} = p1.balance_{ctxt} + m \wedge p2.balance_{ctxt'} = p2.balance_{ctxt} - m.
\end{aligned}$$

In the above specification, the existentially quantified code only appears in the conclusion of the specification. This gives the correct, existentially quantified meaning to the selection of *code*. We require that the two purses are accessible in *ctxt* without reading private fields ($\mathcal{A}ccPub(M, ctxt)$). The specification does not range over extending modules, because the specification as given here implies the one quantifying over module extensions M' – proof is further work.

Another possible meaning of **Pol_1**, however, is that the function `deposit` may only be called if the two objects had the same mint:

$$\begin{aligned}
& \text{Module } M \text{ satisfies policy } \mathbf{Pol_1, vrs3} \\
& \quad \text{iff} \\
& \quad \forall M', ctxt, (ctxt, p1.deposit(p2)) \in \mathcal{A}rising(M * M'). \forall p1, p2:_{ctxt} \text{Purse.} \\
& \quad M * M', ctxt, p1.deposit(p2, m) \rightsquigarrow ctxt', v
\end{aligned}$$

$$\Rightarrow \\ \text{p1.mint}_{ctxt} = \text{p2.mint}_{ctxt}$$

Note that in the above specification the conclusion is only concerned with properties observable in the original context, $ctxt$, while the premise is concerned with properties observable in $ctxt$ as well as $ctxt'$. This reflects the deny nature of the policy.

A fourth possible meaning of **Pol.1** would be that a money transfer from p1 to p2 may take place only if p1 and p2 share the mint. This poses the challenge of identifying the cause of any difference in the balance; we therefore leave it for further work.

Finally, a fifth, and more straightforward meaning would mandate that the balance of a purse p1 may change, only if deposit was executed on p1 or with p1 as an argument. This can be expressed as follows:

Module M satisfies policy **Pol.1, vrs5**

$$\text{iff} \\ \forall M'. \forall (ctxt, code) \in \mathcal{A}rising(M * M'). \forall \text{p1} :_{ctxt} \text{Purse}. \\ M * M', ctxt, code \rightsquigarrow ctxt', v \wedge \text{p1.balance}_{ctxt} \neq \text{p1.balance}_{ctxt'} \\ \Rightarrow \\ \exists ctxt', \text{s.t.} \quad (ctxt', \text{p1.deposit}(-, -)) \in \mathcal{R}each(M * M', ctxt, code) \text{ or} \\ (ctxt', \text{-.deposit}(\text{p1}, -)) \in \mathcal{R}each(M * M', ctxt, code)$$

The assertion $(ctxt', \text{p1.deposit}(-, -)) \in \mathcal{R}each(M * M', ctxt, code)$ expresses that execution of the configuration $(ctxt, code)$ will reach a point where it calls the method `deposit` on the receiver p1 – c.f. App. A.5.

The fourth policy **Pol.4**, “No one can affect the balance of a purse they don’t have”, says that if some runtime configuration affects the balance of some purse `prs`, then the original runtime configuration must have had access to the `prs` itself.

Module M satisfies policy **Pol.4**

$$\text{iff} \\ \forall M', (ctxt, code) \in \mathcal{A}rising(M * M'). \forall \text{prs} :_{ctxt} \text{Purse}: \\ M * M', ctxt \rightsquigarrow ctxt', v \wedge \text{prs.balance}_{ctxt} \neq \text{prs.balance}_{ctxt'} \\ \Rightarrow \\ \text{HasAccess}(M * M', ctxt, code, \text{prs})$$

Note that we have not yet specified the meaning of $\text{HasAccess}(M * M', ctxt, code, \text{prs})$. We will discuss the possible meanings for this together with the second policy:

The second policy. **Pol.2**, stating “Only someone with the mint of a given currency can violate conservation of that currency.”, is along the pattern of **Pol.4**, in that it mandates that certain changes of state (here change in currency) may only happen if the originating context had some property (here access to the mint).

$$\begin{aligned}
 & \text{Module } M \text{ satisfies policy } \mathbf{Pol.2} \\
 & \text{iff} \\
 & \forall M', (ctxt, code) \in \text{Arising}(M * M'). \forall \text{mnt} :_{ctxt} \text{Mint}: \\
 & M * M', ctxt, code \rightsquigarrow ctxt', v \wedge \text{Currency}_{ctxt}(\text{mnt}) \neq \text{Currency}_{ctxt'}(\text{mnt}) \\
 & \Rightarrow \\
 & \text{HasAccess}(M * M', ctxt, code, \text{mnt})
 \end{aligned}$$

We now need to fix the meaning of $\text{HasAccess}(M * M', ctxt, code, \text{mnt})$. We have the following three candidates:

1. $\text{mnt} \in \text{AccAll}(M, ctxt)$, i.e. that $ctxt$ has a path from the stack frame to mnt which involves any fields.
2. $\text{mnt} \in \text{Uses}(M, ctxt, code)$, i.e. that execution of $code$ in the context of $ctxt$ will at some point use the object mnt .
3. $\text{mnt} \in \text{AccPub}(ctxt, code)$, i.e. that $ctxt$ has a path from the stack frame to mnt which involves only public fields, or private fields from the same class as the current recover.

Given lemma 1 from App. A.8, the choices 2) and 3) give stronger guarantees, and are therefore to be preferred.

The module $M_{\text{Purse}} * M_{\text{Mint}}$ satisfies **Pol.2** with meaning 2) – and therefore also with meaning 1), by application of lemma 1. A proof sketch for why $M_{\text{Purse}} * M_{\text{Mint}}$ satisfies **Pol.2** with meaning 2) appears in section 5.

Interestingly, $M_{\text{Purse}} * M_{\text{Mint}}$ does *not* satisfy **Pol.2** with meaning 3). More importantly, without the concept of package and package-local classes, or some concept of ownership, it is impossible to write an implementation for `Purse` so that it satisfies **Pol.2** with meaning 3). Namely, we can always write another class `Cheat` with a private field `myMint` of class `Mint`, and which leaks this field through a public method `leak`. Then, in a context where `x` points to a `Cheat` object, the code snippet `new Purse(x.leak(), 300)`, affects the currency of `x.myMint`, even though the initial context did not have public access to `x.myMint`. More discussion in section 6.3.

The sixth policy – “A reported successful deposit can be trusted as much as one trusts the purse one is depositing into” – relates to trust, a question left to further work.

4.3 Capability Policy Specification Triples

The meanings of policies given in the previous section vary, but they share common characteristics:

- They have the form that execution of some code under some conditions, guarantees some conclusion (except for **Pol 2,vrs2**).
- Conditions may refer to properties of the state before as well as after execution.
- The code may be universally or existentially quantified, or explicitly given.
- Conclusions may refer to properties of the state before as well as after execution.

We propose *policy triples*, consisting of three components. The first and third component stand for the conditions and conclusions mentioned above, and they describe properties of the context before as well as after the execution. In order to distinguish between the latter two, we use the subscripts OLD and NEW. The second component may fully specify some code, or quantify over all possible codes ANY, or just require that some code exists SOME.

In the syntax from below, *PL_Code* stands for programming language code, *Func* stands for user-defined functions (such as *Currency*) as well as system defined functions (such as *AccAll* or *Used*), and *Pred* stands for user or system-defined predicates (such as \in or \geq).

$$\begin{aligned}
 \textit{Policy} & ::= \textit{PolicyTriple}^* \\
 \textit{PolicyTriple} & ::= \langle \textit{Cond}, \textit{Code}, \textit{Cond} \rangle \\
 \textit{Cond} & ::= \textit{Pred}(\textit{Entity}^*) \mid \textit{Cond} \wedge \textit{Cond} \mid \textit{Cond} \vee \textit{Cond} \\
 \textit{Entity} & ::= \textit{Arg}_{\textit{Annot}} \mid \textit{Func}(\textit{Arg}^*)_{\textit{Annot}} \\
 \textit{Annot} & ::= \textit{OLD} \mid \textit{NEW} \\
 \textit{Arg} & ::= \textit{Path} \mid \textit{Value} \\
 \textit{Code} & ::= \textit{ANY} \mid \textit{SOME} \mid \textit{PL_Code}
 \end{aligned}$$

In figure 2 we express some of the policies through policy triples. Requirements such as $\textit{mnt} :_{\textit{OLD}} \textit{Mint}$ are left implicit, as they follow from the use of the identifiers, e.g. from $\textit{Currency}(\textit{mnt})_{\textit{OLD}}$.

Our triples are a generalisation of Hoare logics. Namely,

- The assumptions and the conclusions may refer to properties of the runtime context before as well as after execution, while in Hoare triples the assumptions may only refer to properties of the context before execution of the code.
- The second component may specify the code exactly (as is Hoare triples), or through existential or universal quantification.
- The properties under consideration may contain accessibility or tracing properties.
- The interpretation of the policy triples is always open, i.e. there is an implicit quantification over modules M' extending the module M expected to satisfy the policy.

A module M satisfies a triple $\langle \textit{Cond}_1, \textit{Cd}_1, \textit{Cond}_2 \rangle$ iff for all M' , execution of \textit{Cd}_1 with code $M * M'$ in a context \textit{ctxt} leading to \textit{ctxt}' , then if \textit{Cond}_1 holds on \textit{ctxt} and \textit{ctxt}' , then \textit{Cond}_2 holds on \textit{ctxt} and \textit{ctxt}' . In further work we will define a mapping from such triples to their interpretation.

5 Towards Reasoning about Capability Policies

The main challenges in reasoning about programs' adherence to capability policies are the deny elements of policies, and the combination of rely and deny steps. We have no

$$\begin{aligned}
\mathbf{Pol_1}, \mathbf{vrs2} &\equiv (\text{p1.mint}_{\text{OLD}} = \text{p2.mint}_{\text{OLD}} \wedge \text{p2.balance}_{\text{OLD}} \geq m, \\
&\quad \text{SOME}, \\
&\quad \text{p1.balance}_{\text{NEW}} = \text{p1.balance}_{\text{OLD}} + m \quad \wedge \\
&\quad \text{p2.balance}_{\text{NEW}} = \text{p2.balance}_{\text{OLD}} + m \quad) \\
\mathbf{Pol_2} &\equiv (\text{Currency}(\mathbf{mnt})_{\text{OLD}} \neq \text{Currency}(\mathbf{mnt})_{\text{NEW}}, \\
&\quad \text{ANY}, \\
&\quad \mathbf{mnt}_{\text{OLD}} \in \text{Used} \quad) \\
\mathbf{Pol_3} &\equiv (\text{true}, \\
&\quad \text{ANY}, \\
&\quad \text{Currency}(\mathbf{mnt})_{\text{OLD}} \leq \text{Currency}(\mathbf{mnt})_{\text{NEW}} \quad) \\
\mathbf{Pol_4} &\equiv (\text{prs.balance}_{\text{OLD}} \neq \text{prs.balance}_{\text{NEW}}, \\
&\quad \text{ANY}, \\
&\quad \text{prs}_{\text{OLD}} \in \text{Used} \quad) \\
\mathbf{Pol_5} &\equiv (\text{true}, \\
&\quad \text{ANY}, \\
&\quad \text{prs.balance}_{\text{NEW}} \geq 0 \quad)
\end{aligned}$$

Fig. 1. Some of the five policies expressed through policy triples

full logics yet, but have initial ideas, which we discuss in terms of our example. We consider briefly the code from Fig.1, and sketch the proof for **Pol_1**, and **Pol_2**. The use of restrictive elements is crucial in reasoning about the Joe-E code. The arguments have an flavour of *abduction*, in that we argue that if something were to happen, then the code or the context in which it happened needs to have a certain property.

Verification sketch of Pol_1 in Joe-E/Java The treatment of **Pol_1** requires nothing more than standard Hoare Logic: if prs1 and prs2 share mints, then a call of $\text{prs1.deposit}(\text{prs2}, \text{amt})$ transfers amt from prs2 to prs1 (for appropriate amounts and balances). This can be established by using the usual extension of Hoare Logics to reason about object oriented programs [30, 38]. The treatment of **Pol_1, vrs2** then follows directly, by replacing the existentially quantified *code* by $\text{prs1.deposit}(\text{prs2}, \text{amt})$.

Verification sketch of Pol_2 in Joe-E/Java In contrast, **Pol_2** requires a novel kind of reasoning, which involves the calculation of necessary, rather than sufficient conditions for certain effects to happen. Verification of **Pol_2** can be done through the following steps, annotated by (S1), ... (S11) below:

Assume that $\text{Currency}(\mathbf{mnt})_{\text{ctx}} \neq \text{Currency}(\mathbf{mnt})_{\text{ctx}'}$. (S1): Then, by the definition of *Currency*, either $\text{Ps}(\mathbf{mnt})_{\text{ctx}} \neq \text{Ps}(\mathbf{mnt})_{\text{ctx}'}$, or there exists some $\text{prs} \in \text{Ps}(\mathbf{mnt})_{\text{ctx}}$, such that $\text{prs.balance}_{\text{ctx}} \neq \text{prs.balance}_{\text{ctx}'}$.

1st Case: $\text{Ps}(\mathbf{mnt})_{\text{ctx}} \neq \text{Ps}(\mathbf{mnt})_{\text{ctx}'}$. (S2): Then, either $\exists \text{prs} \in \text{Ps}(\mathbf{mnt})_{\text{ctx}} \setminus \text{Ps}(\mathbf{mnt})_{\text{ctx}'}$, or $\exists \text{prs} \in \text{Ps}(\mathbf{mnt})_{\text{ctx}'} \setminus \text{Ps}(\mathbf{mnt})_{\text{ctx}}$

- 1.1st Case:** $\exists \text{pr} \in Ps(\text{mnt})_{ctxt} \setminus Ps(\text{mnt})_{ctxt'}$. (S3): Then, because there is no garbage collection, $\text{pr}.\text{mint}_{ctxt} \neq \text{pr}.\text{mint}_{ctxt'}$. (S4): However, this is impossible because the field `mint` is `final`.
- 1.2nd Case:** $\exists \text{pr} \in Ps(\text{mnt})_{ctxt'} \setminus Ps(\text{mnt})_{ctxt}$. (S5): If `pr` was defined in `ctxt` then $\text{pr}.\text{mint}_{ctxt} \neq \text{pr}.\text{mint}_{ctxt'}$, which is impossible as in the first case. (S6): If `pr` was undefined in `ctxt'`, then it must have been created through calling one of the two `Purse` constructors. (S7) The constructor `Purse(Purse)` does not affect the currency of the mint. (S8) The constructor `Purse(Mint, balance)` does effect the currency of the mint, but does require access to the mint object, i.e. “uses” it.
- 2nd Case:** $P_s(\text{mnt})_{ctxt} = P_s(\text{mnt})_{ctxt'}$, and $\exists \text{pr} \in P_s(\text{mnt})_{ctxt}$, such that $\text{pr}.\text{balance}_{ctxt} \neq \text{pr}.\text{balance}_{ctxt'}$. (S9): Since the class `Purse` is `final`, and the field `balance` is `private`, the only way to modify this fields is through method in `Purse`. (S10): The only method that changes a `Purse`’s balance is `deposit`. (S11) However, the method `deposit` leaves the currency of a `Mint` constant.

In summary, the argument requires eleven steps (S1)-(S11). Of these steps, three are about normal bookkeeping and set theoretic manipulations (S1,S2,S5), four are using restrictive features (S2, S3, S4,S9), one analyses the structure if the code (S6), and three analyse the semantics of the method bodies (S7,S8,S11).

6 Discussion

In formalising the Mint capability policies, we came across a number of other potential policies that were not captured explicitly in the original six policies.

6.1 Fresh Mints & Purses

Pol.7A Only fresh mint objects are returned from mints and purses.

Pol.7B Only fresh purse objects are returned from mints and purses.

By a “fresh” object, we mean a newly allocated object that has not previous been returned out of the mint-and-purse system. These policies are trivially satisfied by the code in Fig. 2: the construct that returns mints is the default constructor of the `Mint` class, and by the semantics of Java, this will always return a new object. Similarly, the only two methods that return a purse object are the two constructors of the `Purse` class.

Although not explicit, these polices seem to be required implicitly for the other policies to be useful to clients of the Mint system. Consider **Pol.2**: “*Only someone with the mint of a given currency can violate conservation of that currency*” and **Pol.4**: “*No one can affect the balance of a purse they don’t have.*” These policies talk about having a mint or a purse, but don’t say anything about how that mint or purse was obtained. These policies only make sense if the mint and purse capabilities are strictly controlled. The underlying assumption is that a capability effectively belongs to whichever client first requested it, and that only this client can distribute the capability further. **Pol.7A** and **Pol.7B** make this ownership assumption explicit.

These “fresh object” policies may appear trivial because of the simplicity and relatively small size of the mint and purse system. But even a small change to the implementation of the system can break these policies. For example, adding a public `getMint` accessor that returns a purse’s mint (useful, perhaps, for clients with purses from multiple mints) would break policy **Pol.7B** and destroy the integrity of the whole system.

6.2 Subsidiary Capabilities

The simplicity of the Mint example code obscures the need for an additional policy:

Pol.8 Objects implementing Purses and Mints must never be exposed to clients.

Consider an alternative implementation of Purses, where the Mint contained a map from every Purse to its balance. This design is in some sense the complement of Fig. 2: where that design has a `Mint` class that is a pure capability (“`class Mint() { }`”, with no state or behaviour), here the Purse objects would be pure capabilities and the state (the map from purses to their balances) and behaviour (the `deposit` method) would be within the Mint class. The risk this offers is that the original six policies do not mention the internal map object (neither do 7A & B) so an accessor that returned that map would not breach any policy, even though a client with the map capability could do essentially anything to a mint and all its purses. The issue here is another assumption: that there no subsidiary objects in the implementation. Policies such as **Pol.8** ensure that the system will remain secure even if the implementation is changed to require subsidiary objects.

6.3 Supersidiary Capabilities

The third issue we consider relates to supersidiary capabilities, that is, when a Mint or a Purse is combined into a larger structure. In a crucial sense, the Fig. 2 design already involves supersidiary capabilities: each Purse holds the capability to its mint.

The problem this raises can be seen with respect to **Pol.2**: “*Only someone with the mint of a given currency can violate conservation of that currency*”. Purses have mints. So it follows that purses may violate currency conservation. We can guard against this specific aspect with a final policy:

Pol.9 Purses objects must not manipulate their mint capabilities to inflate the currency.

but we can only know we have to write this policy because the purse objects are effectively subsidiary to the module containing all the declarations in the system. The difficulty is much harder to address when a mint, say is part of some supersidiary system — a `Country` object perhaps:

```

1 class Country {
2     private final Mint myMint = new Mint();
3     private final Purse myTreasury = new Purse();
4     public void inflate() {
5         Purse tmpPurse = new Purse(myMint, 1000000000)
6         myTreasury.deposit(tmpPurse, 1000000000)
7     }

```


A `Country` has a mint, and its treasury (a purse belonging to that mint). The `inflate` method creates a new temporary purse containing a billion dollars from thin air, and then deposits that into the treasury — Perhaps this method should have been called `quantitativeEasing`. Now consider a supersidiary client of a `Country` object — the finance minister say. The finance minister does *not* have a reference to the mint (presumably there is an independent central bank) so by **Pol.2** she should not be able to inflate the currency. If, however, the finance minister calls `myCountry.inflate` then the currency will be inflated all the same. (This is the complementary situation to the Purse holding the mint capability and not inflating; here the finance minister doesn't hold the mint capability but does inflate). Formally, these distinctions seem to depend upon the precise semantics of *HasAccess* (see section 4.2). Elucidating these dependencies in detail we leave to future work.

7 Related Work

Object-capabilities were first introduced [24] seven years ago, and many recent studies manage or verify safety or correctness of object-capability programs.

Google's Caja [27] applies sandboxes, proxies, and wrappers to limit components' access to *ambient* capabilities. Sandboxing has been validated formally: Maffeis et al. [19] develop a model of Javascript, demonstrate that it obeys two principles of object-capability systems and show how untrusted applications can be prevented from interfering with the rest of the system. Alternatively, Taly et al. [44] model JavaScript APIs in Datalog, and then carry out a Datlog search for an “attacker” from the set of all valid API calls. This search is similar to the quantification over potential code snippets in our model. Murray and Lowe [28] model object-capability programs in CSP, and use a model checker to ensure program executions do not leak information.

Karim et al. apply static analysis on Mozilla's JavaScript Jetpack extension framework [15], including pointer analyses. Bhargavan et al. x[4] extend language-based sandboxing techniques to support “defensive” components that can execute successfully in otherwise untrusted environments. Meredith et al. [21] encode policies as types in higher order reflective π -calculus.. Politz et al. [33] use a JavaScript typechecker to check properties such as “*multiple widgets on the same page cannot communicate.*” — somewhat similar in spirit to our **Pol.4**. Lerner et al. extend this system to ensure browser extensions observe “*private mode*” browsing conventions, such as that “*no private browsing history retained*” [18]. Dimoulas et al. [8] generalise the language and typechecker based approach to enforce explicit policies, that describe which components may access, or may influence the use of, particular capabilities.

The WebSand [5, 20] and Jeeves [51] projects use dynamic techniques to monitor safe execution of policies. Richards et al. [36] extended this approach by incorporating explicit dynamic ownership of objects (and thus of capabilities), and policies that may examine the history of objects' computations. While these dynamic techniques can restrict or terminate the execution of a component that breaches its security policies, they cannot guarantee in advance that such violations can never happen.

Compared with all these approaches, our work focuses on *general* techniques for specifying (and ultimately verifying) capability policies, whereas these systems are

generally much more *specific*: focusing on one (or a small number) of actual policies. A few formal verification frameworks address JavaScript’s highly dynamic, prototype-based semantics. Gardner et al. [11] developed a formalisation of JavaScript based on separation logic and verified examples. Xiong and Qin et al. [50, 34] worked on similar lines. Swamy et al. [43] recently developed a mechanised verification technique for JavaScript based on the Dijkstra Monad in the F* programming language. Finally, Jang et al. [14] developed a machine-checked proof of five important properties of a web browser — again similar to our simple deny policies — such as “*cookies may not be shared across domains*” by writing the minimal kernel of the browser in Coq.

8 Conclusions and Future Work

In paper, we have advocated that capability policies are necessary for reasoning about programs using object-capability security. We have argued that capability policies must be program centred, fine grained, open, and support both rely and deny elements.

These novel features of the policies require novel features in specifications. We have proposed capability triples to specify policies, which incorporate quantification over program code (to model open policies) predicates over pre- and post- conditions (to handle deny conditions), and by describing paths through programs, are centre upon those programs’ designs. Finally, we have sketched out informal arguments of how adherence to policies may be argued, and we have shown how efforts at specifying policies precisely can uncover additional implicit policies that can be made explicit.

The arguments we have used do not fit the Hoare Logic nor the type-inference format. Nevertheless, they reflect the way one informally reasons about code. They argue in terms of the footprint of a property, and of the set of method calls which might affect that footprint. They consider the uses of restrictive language features (*e.g.* `final`) in the program to reduce that set. They also use rely reasoning (*e.g.* calls to `deposit` or `Purse` preserve the currency in the mint).

We want to develop a formal logic to support reasoning about capability policies. Such a logic will need to combine both rely and deny steps. It will have the usual Hoare Logic rules, as well as inference rules for the calculation of footprints of properties, the effect of restrictive features, for the passing of object capabilities, for lexically scoped languages. To prove soundness of our logic [35] we will need to expand the approach to deal with the deny arguments, perhaps applying ideas from provenance [32], and considerate reasoning [41].

References

1. Pieter Agten, Steven Van Acker, Yoran Brondsema, Phu H. Phung, Lieven Desmet, and Frank Piessens. JSand: complete client-side sandboxing of third-party JavaScript without browser modifications. In *ACSAC*, 2012.
2. Lujo Bauer, Jay Ligatti, and David Walker. Composing security policies with Polymer. In *PLDI*, 2005.
3. Jesper Bengtson, Kathiekeyan Bhargavan, Cedric Fournet, Andrew Gordon, and S.Maffeis. Refinement Types for Secure Implementations. *ACM ToPLAS*, 2011.

4. Karthikeyan Bhargavan, Antoine Delignat-Lavaud, and Sergio Maffei. Language-based defenses against untrusted browser origins. In *USENIX Security*, 2013.
5. Arnar Birgisson, Alejandro Russo, and Andrei Sabelfeld. Capabilities for information flow. In *Programming Languages and Analysis for Security (PLAS)*, 2011.
6. D. G. Clarke, J. M. Potter, and J. Noble. Ownership types for flexible alias protection. In *OOPSLA*. ACM, 1998.
7. Jack B. Dennis and Earl C. Van Horn. Programming Semantics for Multiprogrammed Computations. *Comm. ACM*, 9(3), 1966.
8. Christos Dimoulas, Scott Moore, Aslan Askarov, and Stephen Chong. Declarative policies for capability control. 2013.
9. Mike Dodds, Xinyu Feng, Matthew Parkinson, and Viktor Vafeiadis. Deny-guarantee reasoning. In *ESOP*. Springer, 2009.
10. Sophia Drossopoulou and James Noble. The need for capability policies. In *(FTfJP)*, 2013.
11. Philippa Gardner, Sergio Maffei, and Gareth David Smith. Towards a program logic for JavaScript. In *POPL*, 2012.
12. Andrew D. Gordon and Alan Jeffrey. Typing correspondence assertions for communication protocols. In *MFPS*. Elsevier, ENTCS, 2001.
13. C. A. R. Hoare. Proofs of correctness of data representation. *Acta Informatica*, 1:271–281, 1972.
14. Dongseok Jang, Zachary Tatlock, and Sorin Lerner. Establishing browser security guarantees through formal shim verification. In *USENIX Security*, 2012.
15. Rezwana Karim, Mohan Dhawan, Vinod Ganapathy, and Chung-Chieh Shan. An Analysis of the Mozilla Jetpack Extension Framework. In *ECOOP*, Springer, 2012.
16. Butler W. Lampson and Howard E. Sturgis. Reflection on an Operating System Design. *Communications of the ACM*, 19(5), 1976.
17. K. Rustan M. Leino and Wolfram Schulte. Using history invariants to verify observers. In *ESOP*, 2007.
18. Benjamin S. Lerner, Liam Elberty, Neal Poole, and Shriram Krishnamurthi. Verifying web browser extensions' compliance with private-browsing mode. In *European Symposium on Research in Computer Security (ESORICS)*, September 2013.
19. S. Maffei, J.C. Mitchell, and A. Taly. Object capabilities and isolation of untrusted web applications. In *Proc of IEEE Security and Privacy*, 2010.
20. Jonas Magazinius, Alejandro Russo, and Andrei Sabelfeld. On-the-fly inlining of dynamic security monitors. *Computers & Security*, 31(7):827–843, 2012.
21. L.G. Meredith, Mike Stay, and Sophia Drossopoulou. Policy as types. arXiv:1307.7766 [cs.CR], July 2013.
22. Adrian Mettler, David Wagner, and Tyler Close. Joe-E a Security-Oriented Subset of Java. In *NDSS*, 2010.
23. B. Meyer. *Object-Oriented Software Construction*. Prentice-Hall, 1988.
24. Mark Samuel Miller. *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control*. PhD thesis, Baltimore, Maryland, 2006.
25. Mark Samuel Miller. Secure Distributed Programming with Object-capabilities in JavaScript. Talk at Vrije Universiteit Brussel, mobicrant-talks.eventbrite.com, October 2011.
26. Mark Samuel Miller, Chip Morningstar, and Bill Frantz. Capability-based Financial Instruments: From Object to Capabilities. In *Financial Cryptography*. Springer, 2000.
27. Mark Samuel Miller, Mike Samuel, Ben Laurie, Ihab Awad, and Mike Stay. Safe active content in sanitized JavaScript. code.google.com/p/google-caja/.
28. Toby Murray and Gavin Lowe. Analysing the information flow properties of object-capability patterns. In *FAST*, LNCS, 2010.
29. Roger Needham. Protection systems and protection implementations. In *Joint Computer Conference*, pages 571–578, 1972.

30. M. Parkinson and G. Bierman. Separation logic, abstraction and inheritance. In *POPL*, pages 75–86. ACM Press, 2008.
31. Matthew Parkinson. Class invariants: the end of the road? In *IWACO*, 2007.
32. Roly Perera, Umut Acar, James Cheney, and Paul Blain Levy. Functional programs that explain their work. In *ICFP*. ACM, 2012.
33. Joe Gibbs Politz, Spiridon Aristides Eliopoulos, Arjun Guha, and Shriram Krishnamurthi. Adsafety: Type-based verification of javascript sandboxing. In *USENIX Security*, 2011.
34. Shengchao Qin, Aziem Chawdhary, Wei Xiong, Malcolm Munro, Zongyan Qiu, and Huibiao Zhu. Towards an axiomatic verification system for javascript. In *TASE*, pages 133–141, 2011.
35. Azalea Raad and Sophia Drossopoulou. A Sip of the Chalice. In *FTFJP*, July 2011.
36. Gregor Richards, Christian Hammer, Francesco Zappa Nardelli, Suresh Jagannathan, and Jan Vitek. Flexible access control for JavaScript. In *OOPSLA*, pages 305–322, 2013.
37. Jerome H. Saltzer. Protection and the control of information sharing in Multics. *CACM*, 17(7):p.389ff, 1974.
38. Jan Smans, Bart Jacobs, and Frank Piessens. Implicit Dynamic Frames. *ToPLAS*, 2012.
39. Marc Stiegler. The lazy programmer’s guide to security. HP Labs, www.object-oriented-security.org.
40. Bjarne Stroustrup. *The C++ Programming Language*. Addison Wesley, 1986.
41. Alexander J. Summers and Sophia Drossopoulou. Considerate Reasoning and the Composite Pattern. In *VMCAI*, 2010.
42. Alexander J. Summers, Sophia Drossopoulou, and Peter Müller. The need for Flexible Object Invariants. In *IWACO*, ACM DL, July 2009.
43. Nikhil Swamy, Joel Weinberger, Cole Schlesinger, Juan Chen, and Benjamin Livshits. Verifying higher-order programs with the dijkstra monad. In *PLDI*, pages 387–398, 2013.
44. Ankur Taly, Ulfar Erlingsson, John C. Mitchell, Mark S. Miller, and Jasvir Nagra. Automated Analysis of Security-Critical JavaScript APIs. In *IEEE Symposium on Security and Privacy (SP)*, 2011.
45. The NewpeakTeam. Several Newspeak Documents. newspeaklanguage.org/, September 2012.
46. Tom van Cutsem. Membranes in Javascript. prog.vub.ac.be/tvcutsem/invokedynamic/js-membranes.
47. M. V. Wilkes and R. M. Needham. The Cambridge CAP computer and its operating system, 1979.
48. Niklaus Wirth. *Programming in Modula-2*. Springer, 1982.
49. T. Wood and S. Lam. A semantic model for authentication protocols. In *IEEE Computer Society Symposium on Research in Security and Privacy*, 1993.
50. Wei Xiong. *Verification and Validation of JavaScript*. PhD thesis, 2013, Durham University.
51. Jean Yang, Kuat Yessenov, and Armando Solar-Lezama. A language for automatically enforcing privacy policies. In *POPL*, 2012.

A Formal Definition of $\mathcal{L}ng_{c,j}$

A.1 Modules and Syntax

Modules map class identifiers to their superclass, field definitions and method declarations. Extensibility annotations, *ea*, are attached to classes and methods, and they forbid (*noext*) or allow (*exts*) subclasses, resp. redefinition of the method in a subclass. Privacy annotations (*pa*) are attached to fields and methods, and restrict (*priv*) or allow

(pub) access outside the defining class. Mutability annotations, *ma*, make fields immutable (fin) or mutable (vol). We expect that constructors will be defined as methods with the identifier *constr*.

$$\begin{aligned} \text{Module} &= \text{ClassId} \longrightarrow (\text{ea} \times \\ &\quad \text{ClassId} \times \\ &\quad (\text{FieldId} \longrightarrow \text{ma pa ClassId}) \times \\ &\quad (\text{MethId} \longrightarrow \text{ea pa methBody}) \quad) \\ \text{ea} &::= \text{noext} \mid \text{exts} \\ \text{pa} &::= \text{priv} \mid \text{pub} \\ \text{ma} &::= \text{fin} \mid \text{mut} \\ \text{meth} &::= \text{ClassId } m \text{ (ClassId } x \text{) } \{ e \} \\ \text{Used } e &::= e.f \mid e.f := e \mid e.m(e) \mid \text{new } c \mid x \mid \text{this} \mid \text{null} \end{aligned}$$

A.2 The Type System of $\mathcal{L}ng_{c_j}$

The lookup functions \mathcal{F} and \mathcal{M} return corresponding field and method definitions.

$$\begin{aligned} \mathcal{F}(M, c, f) &= \begin{cases} \text{undefined} & \text{if } c = \text{Object}, \\ M(c) \downarrow_3(f) & \text{if } M(c) \downarrow_3(f) \text{ is defined,} \\ \mathcal{F}(M, M(c) \downarrow_1, f) & \text{otherwise.} \end{cases} \\ \mathcal{M}(M, c, m) &= \begin{cases} \text{undefined} & \text{if } c = \text{Object}, \\ M(c) \downarrow_3(m) & \text{if } M(c) \downarrow_3(m) \text{ is defined,} \\ \mathcal{M}(M, M(c) \downarrow_1, m) & \text{otherwise.} \end{cases} \end{aligned}$$

Typing is expressed through the judgment $\Gamma \vdash e : c$, where Γ is a mapping from $\{\text{this}, x\}$ to the set of class identifiers, and from *constr* to a boolean (to indicate whether Γ belongs to a constructor's method body).

$$\begin{array}{c} \text{VarThisNull} \qquad \text{fldRd} \\ \frac{M(c) \text{ is defined}}{M, \Gamma \vdash x : \Gamma(x)} \qquad \frac{M, \Gamma \vdash e : c}{\mathcal{F}(M, c, f) = \text{ma pa } c'} \\ \frac{M, \Gamma \vdash \text{this} : \Gamma(\text{this}) M, \Gamma \vdash \text{null} : c}{\text{pa} = \text{pub} \vee \Gamma(\text{this}) = c} \\ M, \Gamma \vdash e.f : c' \\ \text{fldAss} \qquad \text{fldInitConstr} \\ \frac{M, \Gamma \vdash e : c \quad \mathcal{F}(M, c, f) = \text{ma pa } c'' \quad M, \Gamma \vdash e' : c' \quad \text{ma} = \text{mut} \wedge (\text{pa} = \text{pub} \vee \Gamma(\text{this}) = c) \quad M \vdash c' \leq c''}{M, \Gamma \vdash e.f := e' : c'} \quad \frac{\Gamma(\text{this}) = c \quad \mathcal{F}(M, c, f) = \text{immut pa } c'' \quad \Gamma(\text{constr}) = \text{true} \quad M, \Gamma \vdash e' : c' \quad M \vdash c' \leq c''}{M, \Gamma \vdash \text{this.f} := e' : c'} \\ \text{constrCall} \qquad \text{methCall} \\ \frac{\mathcal{M}(M, c, \text{constr}) = _ \text{pa } c \text{ constr}(_ _) \{ e \} \quad M, \Gamma \vdash e : c'' \quad \text{pa} = \text{pub} \vee \Gamma(\text{this}) = c \quad M \vdash c'' \leq c'}{M, \Gamma \vdash \text{new } c(e) : c} \quad \frac{M, \Gamma \vdash e_0 : c \quad M, \Gamma \vdash e_1 : c_1 \quad \text{pa} = \text{pub} \vee \Gamma(\text{this}) = c \quad \mathcal{M}(M, c_0, m) = \text{ea pa } c' m(c_3 x) \{ _ \} \quad M \vdash c_1 \leq c_3}{M, \Gamma \vdash e_0.m(e_1) : c'} \end{array}$$

A.3 Well-formed class and module, linking

well_formed_class

$$\begin{array}{l}
M(c) = ea\ c' \dots \\
c' = \text{Object} \vee (ea = \text{exts} \wedge M(c') \text{ is defined}) \\
\forall f : M(c) \downarrow_3 (f) = c' \implies M(c') \text{ is defined} \wedge \mathcal{F}(M, c', f) \text{ is not defined} \\
\forall m : M(c) \downarrow_4 (m) = ea' pa\ c' m(c''\ x) \{ e \} \implies \\
\quad M(c'), M(c'') \text{ are defined} \\
\quad \Gamma(\text{this}) = c, \Gamma(x) = c'', \Gamma(\text{constr}) = (m = \text{constr}) \\
\quad M, \Gamma \vdash e : c''' \quad M \vdash c''' \leq c' \\
\quad \mathcal{M}(M, c', m) \text{ is undefined} \vee (ea' = \text{exts} \wedge \mathcal{M}(M, c', m) = ea' pa\ c' m(c''\ x) \{ - \}) \\
\hline
M \vdash c \diamond \\
\vdash M \diamond \quad \text{iff} \quad \forall c \in \text{dom}(M) : M \vdash c \diamond
\end{array}$$

Linking of two modules M and M' is the union of their respective mappings, provided that the domains of the two modules are disjoint, and that the union of the mappings creates a well-formed module:

$$\begin{array}{l}
* : \text{Module} \times \text{Module} \longrightarrow \text{Module} \\
M * M' = \begin{cases} M *_{aux} M', & \text{if } \vdash M *_{aux} M' \diamond, \text{ and } \text{dom}(M) \cap \text{dom}(M') = \emptyset \\ \text{undefined}, & \text{otherwise.} \end{cases} \text{ where} \\
(M *_{aux} M')(c) = M(c), \text{ if } M(c) \text{ is defined, } M'(c), \text{ otherwise.}
\end{array}$$

A.4 Runtime context

The runtime context consists of a stack frame, ϕ , and a heap χ . The frame is a tuple consisting of the address belonging to the receiver (**this**), and the value for the argument (**x**). Values are addresses or null). Addresses are ranged over by ι , and they are natural numbers. The heap maps addresses to objects. Objects are tuples consisting of the class of the object, and a mapping field identifiers onto values.

$$\begin{array}{ll}
\text{ctxt} & = \text{frame} \times \text{heap} & \phi \in \text{frame} & = \text{addr} \times \text{val} \\
\chi \in \text{heap} & = \text{addr} \longrightarrow \text{object} & \mathbf{v} \in \text{val} & = \{ \text{null} \} \cup \text{addr} \\
\text{object} & = \text{ClassId} \times (\text{FieldId} \longrightarrow \text{val}) & \iota \in \text{addr} & = \mathbb{N}
\end{array}$$

The Operational Semantics of $\mathcal{L}ng_{c_j}$ Execution uses module M , and maps a runtime context $ctxt$ and expression e (code in the general case) onto a new context $ctxt'$ and a result.

$$\begin{array}{l}
\rightsquigarrow : \text{Module} \times \text{ctxt} \times \text{expr} \longrightarrow \text{ctxt} \times \text{res} \\
\text{res} = \{ \text{nullPtrExc}, \text{stuckErr} \} \cup \text{val}
\end{array}$$

Here the salient rules for execution; we omit those which raise & propagate exceptions.

$$\begin{array}{c}
\text{val} \\
\frac{}{M, \text{ctxt}, \mathbf{v} \rightsquigarrow \text{ctxt}, \mathbf{v}} \\
\text{fldRd} \\
\frac{M, \text{ctxt}, \mathbf{e} \rightsquigarrow \text{ctxt}', \iota \quad \text{ctxt} = \rightarrow, \chi' \quad \chi'(\iota) \downarrow_2 (\mathbf{f}) = \mathbf{v}}{M, \text{ctxt}, \mathbf{e.f} \rightsquigarrow \text{ctxt}', \mathbf{v}}
\end{array}
\qquad
\begin{array}{c}
\text{thisPar} \\
\frac{\text{ctxt} = (\iota, \mathbf{v}), \chi}{M, \text{ctxt}, \text{this} \rightsquigarrow \text{ctxt}, \iota} \\
\frac{}{M, \text{ctxt}, \mathbf{x} \rightsquigarrow \text{ctxt}, \mathbf{v}} \\
\text{fldAss} \\
\frac{M, \text{ctxt}, \mathbf{e} \rightsquigarrow \text{ctxt}'', \iota \quad M, \text{ctxt}'', \mathbf{e}' \rightsquigarrow \text{ctxt}''', \mathbf{v} \quad \text{ctxt}''' = \phi''', \chi''', \chi' = \chi'''[\iota \mapsto \chi'''(\iota)][\mathbf{f} \mapsto \mathbf{v}]}{M, \text{ctxt}, \mathbf{e.f} := \mathbf{e}' \rightsquigarrow (\phi''', \chi'), \mathbf{v}}
\end{array}$$

new	methCall
$M, ctxt, e \rightsquigarrow ctxt', v$ $\{f \mid \mathcal{F}(M, c, f) \text{ is defined}\} = \{f_1, \dots, f_r\}$ $\iota \text{ is new in } \chi''$ $ctxt'' = \phi, \chi''$ $\chi''' = \chi''[\iota \mapsto (c, (f_1 : \text{null}, \dots, f_r : \text{null}))]$ $\mathcal{M}(M, c, \text{constr}) = \dots (..) \{e'\}$ $ctxt''' = (\iota, v), \chi'''$ $M, ctxt''', e' \rightsquigarrow ctxt', v'$	$M, ctxt, e \rightsquigarrow ctxt'', \iota$ $M, ctxt'', e' \rightsquigarrow ctxt''', v$ $ctxt''' = (\iota, v), \chi'''$ $\chi'''(\iota) \downarrow_1 = c$ $\mathcal{M}(M, c, m) = \dots m(\dots) \{e''\}$ $ctxt'''' = (\iota, v), \chi''''$ $M, ctxt''''', e'' \rightsquigarrow ctxt', v$
$M, ctxt, \text{new } c(e) \rightsquigarrow ctxt', \iota$	$M, ctxt, e.m(e') \rightsquigarrow ctxt', v$

A.5 Reach: Reachable Configurations

Configurations consist of runtime contexts and expressions. A configuration is reachable from another configuration, if the former may be required for the evaluation of the former after any number of steps. We are only interested in the configurations which correspond to method calls, or constructor invocations. $\text{Reach}(M, ctxt, e)$ returns the sets of all configurations reachable from $ctxt, e$:

$$\text{Reach} : \text{Program} \times \text{Context} \times \text{Expr} \mapsto \mathcal{P}(\text{Context} \times \text{Expr})$$

We define the function Reach by cases on the structure of the expression, and depending on the execution of the expression. Note that the function $\text{Reach}(M, ctxt, expr)$, even when the execution should diverge. This is important, because it allows us to give meaning to capability policies without requiring termination. In case of divergence, $\text{Reach}(M, ctxt, \text{null})$ will be an infinite set.

$$\begin{aligned}
\text{Reach}(M, ctxt, \text{null}) &= \emptyset \\
\text{Reach}(M, ctxt, \text{this}) &= \emptyset \\
\text{Reach}(M, ctxt, x) &= \emptyset \\
\text{Reach}(M, ctxt, e.f) &= \text{Reach}(ctxt, e) \\
\text{Reach}(M, ctxt, e.f := e') &= \text{Reach}(M, ctxt, e) \\
&\cup \begin{cases} \text{Reach}(M, ctxt', e'), \\ \text{if } \exists ctxt', v, s.t : M, ctxt, e \rightsquigarrow ctxt', v \\ \emptyset, \text{ otherwise.} \end{cases} \\
\text{Reach}(M, ctxt, \text{new } c(e)) &= \text{Reach}(M, ctxt, e) \\
&\cup \begin{cases} \{ctxt', e'\} \\ \text{if } (A) : \exists ctxt'', v, s.t : \\ M, ctxt, e \rightsquigarrow ctxt'', v, \\ \mathcal{M}(M, c, \text{constr}) = \dots (..) \{e'\} \\ \text{and where } ctxt'' = \dots, \chi'', \text{ and } ctxt' = (\iota, v), \chi''' \\ \iota \text{ new in } \chi'', \chi''' = \chi''[\iota \mapsto (c, (f_1 : \text{null}, \dots, f_r : \text{null}))] \\ \{f \mid \mathcal{F}(M, c, f) \text{ is defined}\} = \{f_1, \dots, f_r\} \\ \emptyset, \text{ otherwise.} \end{cases} \\
&\cup \begin{cases} \text{Reach}(M, ctxt', e'), & \text{if } (A) \text{ as above.} \\ \emptyset, & \text{otherwise.} \end{cases}
\end{aligned}$$

$$\begin{aligned}
\mathcal{R}each(M, ctxt, e.m(e')) &= \mathcal{R}each(M, ctxt, e) \\
&\cup \left\{ \begin{array}{l} \mathcal{R}each(M, ctxt', e') \\ \text{if } (B) : \exists ctxt', \iota, s.t : \\ M, ctxt, e \rightsquigarrow ctxt', \iota, \\ \emptyset, \text{ otherwise.} \end{array} \right. \\
&\cup \left\{ \begin{array}{l} \{ctxt''', e''\} \\ \text{if } (B) \text{ as above, and } (C), \exists ctxt'', v, s.t : \\ M, ctxt', e' \rightsquigarrow ctxt'', v, \\ \text{and where } ctxt'' = _., \chi'', \chi''(iota) \downarrow_1 = c \\ \mathcal{M}(M, c, m) = _..m(_.) \{e''\} \\ \text{and where } ctxt''' = (\iota, v), \chi''. \end{array} \right. \\
&\cup \left\{ \begin{array}{l} \emptyset, \text{ otherwise.} \\ \mathcal{R}each(M, ctxt''', e'') \quad \text{if } (B) \text{ and } (C), \text{ as above.} \\ \emptyset, \text{ otherwise.} \end{array} \right.
\end{aligned}$$

A.6 *Arising*: configurations reachable from initial configurations

We define as initial configurations all configurations that may be encountered at the start of program execution.

$$\mathcal{I}nit : Program \mapsto \mathcal{P}(Context \times Expr)$$

Initial configuration should be as “minimal” as possible, We therefore construct a heap which has only one object. And we require that the expression be well typed under the assumption that `this` and `x` are denoting objects of class `Object`.

$$\begin{aligned}
\mathcal{I}nit(M) &= \{ (ctxt_0, e) \mid \exists c. M, \Gamma_0 \vdash e : c \} \\
&\text{where } ctxt_0 = ((\iota_0, \text{null}), \chi_0), \\
&\text{and } dom(\chi_0) = \{ \iota \}, \text{ and } \chi_0(\iota_0) = (\text{Object}, \emptyset), \\
&\text{and } \Gamma = \text{this} \mapsto \text{Object}, x \mapsto \text{Object}, \text{constr} \mapsto \text{false}.
\end{aligned}$$

The *arising* configurations are those which may be reached by executing an initial configuration:

$$\mathcal{A}rising : Program \mapsto \mathcal{P}(Context \times Expr)$$

Arising configurations they are defined as follows:

$$\mathcal{A}rising(M) = \bigcup_{(ctxt, e) \in \mathcal{I}nit(M)} \mathcal{R}each(M, ctxt, e)$$

A.7 *Used*: the objects reads by execution of a configuration

We are interested in collecting all addresses read during execution of a configuration:

$$\mathcal{U}sed : Program \times Context \times Expr \mapsto \mathcal{P}(Addr)$$

$\mathcal{U}sed$ is defined by cases on the structure of the expression, and depending on the execution of this expression. We use similar cases as those in the definition in A.5.

$$\begin{aligned}
\mathcal{U}sed(M, ctxt, \text{null}) &= \emptyset \\
\mathcal{U}sed(M, ctxt, \text{this}) &= \{ \phi(\text{this}) \} \text{ where } ctxt = \phi, _ \\
\mathcal{U}sed(M, ctxt, x) &= \{ \phi(x) \} \text{ where } ctxt = \phi, _ \\
\mathcal{U}sed(M, ctxt, e.f) &= \mathcal{U}sed(M, ctxt, e) \cup \left\{ \begin{array}{l} \{ \chi'(\iota) \downarrow (f) \} \\ \text{if } \exists ctxt', \iota, s.t : \\ M, ctxt, e \rightsquigarrow (\phi', \chi'), \iota \\ \emptyset, \text{ otherwise.} \end{array} \right.
\end{aligned}$$

$$\begin{aligned}
Used(M, ctxt, e.f := e') &= Used(M, ctxt, e.f) \\
&\cup \begin{cases} Used(M, ctxt', e'), \\ \text{if } \exists ctxt', v, s.t : \\ M, ctxt, e \rightsquigarrow ctxt', v \\ \emptyset, \text{ otherwise.} \end{cases} \\
Used(M, ctxt, \text{new } c(e)) &= Used(M, ctxt, e) \\
&\cup \begin{cases} \{\iota\} \text{ if } (A), \text{ where } (A) \text{ as defined in A.5} \\ \emptyset, \text{ otherwise.} \end{cases} \\
&\cup \begin{cases} Used(M, ctxt', e'), \text{ if } (A) \text{ as above.} \\ \emptyset, \text{ otherwise.} \end{cases} \\
Used(M, ctxt, e.m(e')) &= Used(M, ctxt, e) \\
&\cup \begin{cases} Used(M, ctxt', e') \text{ if } (B), \text{ where } (B) \text{ as defined in A.5} \\ \emptyset, \text{ otherwise.} \end{cases} \\
&\cup \begin{cases} \{ctxt''', e''\} \\ \text{if } (B) \text{ and } (C), \text{ where } (B), (C) \text{ as defined in A.5} \\ \emptyset, \text{ otherwise.} \end{cases} \\
&\cup \begin{cases} Reach(M, ctxt''', e'') \text{ if } (B) \text{ and } (C), \text{ as above.} \\ \emptyset, \text{ otherwise.} \end{cases}
\end{aligned}$$

A.8 *AccAll*: objects accessible through any paths, *AccPub*: objects accessible through public paths

We collect the addresses of all objects which are accessible from a stack frame through any paths which go through public or through private fields, through the function *AccAll*: We also collect the addresses of all objects which are accessible from a stack frame through paths which go through public fields, or through private fields which, however, belong to objects of the same class as the current receiver:

$$\begin{aligned}
AccAll &: Program \times Context \mapsto \mathcal{P}(Addr) \\
AccPub &: Program \times Context \mapsto \mathcal{P}(Addr)
\end{aligned}$$

These sets are defined as follows:

$$\begin{aligned}
AccAll(M, ctxt) &= \begin{cases} \{\iota, v\} \cup \\ \{\chi(\iota') \downarrow_2(\mathbf{f}) \mid \iota' \in AccAll(M, ctxt) \\ \wedge \chi(\iota') \downarrow_1 = c' \wedge \mathcal{F}(M, c', \mathbf{f}) = _ _ \} \\ \text{where } ctxt = ((\iota, v), \chi). \end{cases} \\
AccPub(M, ctxt) &= \begin{cases} \{\iota, v\} \cup \\ \{\chi(\iota') \downarrow_2(\mathbf{f}) \mid \iota' \in AccPub(M, ctxt) \\ \wedge \chi(\iota') \downarrow_1 = c' \wedge \mathcal{F}(M, c', \mathbf{f}) = _ pa _ \\ \wedge (pa = \text{pub} \vee c = c') \} \\ \text{where } ctxt = ((\iota, v), \chi) \text{ and } \chi(\iota) \downarrow_1 = c. \end{cases}
\end{aligned}$$

Lemma 1. For all $M, ctxt, ctxt', e, v$ s.t. $M, ctxt, e \rightsquigarrow ctxt', v$:

$$AccPub(M, ctxt) \subseteq AccAll(M, ctxt) \quad \wedge \quad Used(M, ctxt, e) \subseteq AccAll(M, ctxt)$$

The lemma is not surprising, and the full proof is further work. Note that in general, there is no subset relation between $AccPub(M, ctxt)$ and $Used(M, ctxt, e)$.