Security properties of typed applets*

Xavier Leroy François Rouaix

INRIA Rocquencourt**
Domaine de Voluceau, 78153 Le Chesnay, France.

Abstract. This paper formalizes the folklore result that strongly-typed applets are more secure than untyped ones. We formulate and prove several security properties that all well-typed applets possess, and identify sufficient conditions for the applet execution environment to be safe, such as procedural encapsulation, type abstraction, and systematic type-based placement of run-time checks. These results are a first step towards formal techniques for developing and validating safe execution environments for applets.

1 Introduction

What, exactly, makes strongly-typed applets more secure than untyped ones? Most frameworks proposed so far for safe local execution of foreign code rely on strong typing, either statically checked at the client side [19, 44], statically checked at the server side and cryptographically signed [35], or dynamically checked by the client [8]. However, the main property guaranteed by strong typing is type soundness: "well-typed programs do not go wrong", e.g. do not apply an integer as if it were a function. While violations of type soundness constitute real security threats (casting a well-chosen string to a function or object type allows arbitrary code to be executed), there are many more security concerns, such as integrity of the running site (an applet should not delete or modify arbitrary files) and confidentiality of user's data (an applet should not divulge personal information over the network). The corresponding security violations do not generally invalidate type soundness in the conventional sense.

If we examine the various security problems identified for Java applets [12], some of them do cause a violation of Java type soundness [21]; others correspond to malicious, but well-typed, uses of improperly protected functions from the applet's execution environment [7]. Another typical example is the ActiveX applet described in [10] that does a Trojan attack on the Quicken home-banking software: money gets transferred from the user's bank account to some offshore account, all in a perfectly type-safe way.

---

* Appeared in *Secure Internet Programming – Security issues for Mobile and Distributed Objects*, edited by J. Vitek and C. Jensen, pages 147–184, Springer-Verlag LNCS 1603, 1999.

** François Rouaix's current affiliation is Liquid Market Inc, 5757 West Century Bld, Los Angeles, CA 90045, USA.

On these examples, it is intuitively obvious that security properties must be enforced by the applet's execution environment. It is the environment that eventually decides which computer resources the applet can access. This is the essence of the so-called "sandbox model". Strong typing comes in the picture only to guarantee that this environment is used in accordance with its publicized interface. For instance, typing prevents an applet from jumping in the middle of the code for an environment function, or scanning the whole memory space of the browser, which would allow the applet to abuse or bypass entirely the execution environment.

The purpose of this paper is to give a formal foundation to the intuition above. We formulate and prove several security properties that all well-typed applets possess. Along the way, we identify sufficient conditions for the execution environment to be safe, such as procedural encapsulation, type abstraction, and systematic type-based placement of run-time checks. These results are a first step towards formal techniques for developing and validating safe execution environments for applets.

The remainder of this paper is organized as follows. Section 2 introduces a simple language for applets and their environment, and formalizes the security policy that they must obey. Section 3 proves a security property based on the notion of lexical scoping, then extends it to take procedural abstraction into account. In section 4, we equip our language with a simple type system, which is used in section 5 to prove three type-based security properties, two relying on run-time checks and the other on a combination of run-time checks and type abstraction. After a brief parallel with object-oriented languages in section 6, section 7 outlines the security architecture for applets in the MMM Web browser and discusses the relevance of our results to this practical example. Related work is discussed in section 8, followed by concluding remarks in section 9.

2 The language and its security policy

2.1 The language

The language we consider in this paper is a simple lambda-calculus with base types (integers, strings, ...), pairs as the only data structure, and references in the style of ML. The syntax of terms, with typical element a, is as follows:

Terms:		
	$a ::= x$	identifiers
	$\mid b$	constant of base type (integer, ...)
	$\mid \lambda x.a$	function abstraction
	$\mid a_1(a_2)$	function application
	$\mid (a_1, a_2)$	pair construction
	$\mid \mathtt{fst}(a) \mid \mathtt{snd}(a)$	pair projections
	$\mid \mathtt{ref}(a)$	reference creation
	$\mid\ !a$	dereferencing
	$\mid a_1 := a_2$	reference assignment

References are included in the language from the beginning not only to account for imperative programming (all kinds of assignment on variables and mutable data structures such as arrays can easily be modeled with references), but also to provide easily observable criteria on which we base the security policy.

2.2 The security policy

The security policy we apply to applets is based on the notion of sensitive store locations: locations of references that an applet must not modify during its execution, or more generally references that can be modified during the applet execution, but whose successive contents must always satisfy some invariant, i.e. remain within a given set of permitted values.

The first motivation for this policy is to formalize the intuitive idea that an applet must not trash the memory of the computer executing it. In particular, the internal state of the browser, the operating system, and other applications running on the machine must not be adversely affected by the applet.

This security policy can also be stretched to account for input/output behavior, notably accesses to files and simple cases of network connections. A low-level, hardware-oriented view of I/O is to consider hardware devices such as the disk controller and network interface as special locations in the store; I/O is then controlled by restricting what can be written to these locations. For a higher-level view, each file or network connection can be viewed as a reference, which can then be controlled independently of others. Here, the file system, the name service and the routing tables become dictionary-like data structures mapping file names, host names, and network addresses to the references representing files and connections.

By concentrating on writes to sensitive locations, we focus on integrity properties of the system running the applet. It is also possible to control reads from sensitive locations, thus establishing simple privacy properties. We will not do it in this paper for the sake of simplicity, but the results of section 3 also extend to controlled reads. More advanced privacy properties, as provided for instance by information flow models, are well beyond our approach, however.

2.3 The instrumented semantics

To enforce the security policy, we give a semantics to our language that monitors reference assignments, and reports run-time errors in the case of illegal writes. We use a standard big-step operational semantics in the style of [32, 39, 25]. Source terms are mapped to values, which are terms with the following syntax:

Values:	$v ::= b$	values of base types
	$\mid \lambda x.a[e]$	function closures
	$\mid (v_1, v_2)$	pairs of values
	$\mid \ell$	store locations
Results:	$r ::= v/s$	normal termination
	$\mid \mathtt{err}$	write violation detected

Normal rules:

$$\varphi, e, s \vdash x \to e(x)/s \quad (1) \qquad\qquad \varphi, e, s \vdash b \to b/s \quad (2)$$

$$\varphi, e, s \vdash \lambda x.a \to \lambda x.a[e]/s \quad (3)$$

$$\frac{\varphi, e, s \vdash a_1 \to \lambda x.a'[e']/s_1 \qquad \varphi, e, s_1 \vdash a_2 \to v_2/s_2 \qquad \varphi, e'\{x \leftarrow v_2\}, s_2 \vdash a' \to r}{\varphi, e, s \vdash a_1(a_2) \to r} \quad (4)$$

$$\frac{\varphi, e, s \vdash a_1 \to v_1/s_1 \qquad \varphi, e, s_1 \vdash a_2 \to v_2/s_2}{\varphi, e, s \vdash (a_1, a_2) \to (v_1, v_2)/s_2} \quad (5)$$

$$\frac{\varphi, e, s \vdash a \to (v_1, v_2)/s'}{\varphi, e, s \vdash \mathtt{fst}(a) \to v_1/s'} \quad (6) \qquad \frac{\varphi, e, s \vdash a \to (v_1, v_2)/s'}{\varphi, e, s \vdash \mathtt{snd}(a) \to v_2/s'} \quad (7)$$

$$\frac{\varphi, e, s \vdash a \to v/s' \qquad \ell \notin \mathrm{Dom}(s') \cup \mathrm{Dom}(\varphi)}{\varphi, e, s \vdash \mathtt{ref}(a) \to \ell/s'\{\ell \leftarrow v\}} \quad (8) \qquad \frac{\varphi, e, s \vdash a \to \ell/s'}{\varphi, e, s \vdash !a \to s'(\ell)/s'} \quad (9)$$

$$\frac{\varphi, e, s \vdash a_1 \to \ell/s_1 \qquad \varphi, e, s_1 \vdash a_2 \to v_2/s_2 \qquad \ell \notin \mathrm{Dom}(\varphi) \text{ or } v_2 \in \varphi(\ell)}{\varphi, e, s \vdash (a_1 := a_2) \to ()/s_2\{\ell \leftarrow v_2\}} \quad (10)$$

$$\frac{\varphi, e, s \vdash a_1 \to \ell/s_1 \qquad \varphi, e, s_1 \vdash a_2 \to v_2/s_2 \qquad \ell \in \mathrm{Dom}(\varphi) \text{ and } v_2 \notin \varphi(\ell)}{\varphi, e, s \vdash (a_1 := a_2) \to \mathtt{err}} \quad (11)$$

Error propagation rules:

$$\frac{\varphi, e, s \vdash a_1 \to \mathtt{err}}{\varphi, e, s \vdash a_1(a_2) \to \mathtt{err}} \quad (12) \qquad \frac{\varphi, e, s \vdash a_1 \to v_1/s_1 \qquad \varphi, e, s_1 \vdash a_2 \to \mathtt{err}}{\varphi, e, s \vdash a_1(a_2) \to \mathtt{err}} \quad (13)$$

$$\frac{\varphi, e, s \vdash a_1 \to \mathtt{err}}{\varphi, e, s \vdash (a_1, a_2) \to \mathtt{err}} \quad (14) \qquad \frac{\varphi, e, s \vdash a_1 \to v_1/s_1 \qquad \varphi, e, s_1 \vdash a_2 \to \mathtt{err}}{\varphi, e, s \vdash (a_1, a_2) \to \mathtt{err}} \quad (15)$$

$$\frac{\varphi, e, s \vdash a \to \mathtt{err}}{\varphi, e, s \vdash \mathtt{fst}(a) \to \mathtt{err}} \quad (16) \qquad \frac{\varphi, e, s \vdash a \to \mathtt{err}}{\varphi, e, s \vdash \mathtt{snd}(a) \to \mathtt{err}} \quad (17)$$

$$\frac{\varphi, e, s \vdash a \to \mathtt{err}}{\varphi, e, s \vdash \mathtt{ref}(a) \to \mathtt{err}} \quad (18) \qquad \frac{\varphi, e, s \vdash a \to \mathtt{err}}{\varphi, e, s \vdash !a \to \mathtt{err}} \quad (19)$$

$$\frac{\varphi, e, s \vdash a_1 \to \mathtt{err}}{\varphi, e, s \vdash (a_1 := a_2) \to \mathtt{err}} \quad (20) \qquad \frac{\varphi, e, s \vdash a_1 \to v_1/s_1 \qquad \varphi, e, s_1 \vdash a_2 \to \mathtt{err}}{\varphi, e, s \vdash (a_1 := a_2) \to \mathtt{err}} \quad (21)$$

Fig. 1. Evaluation rules

Environments: $e ::= [x_1 \leftarrow v_1 \ldots, x_n \leftarrow v_n]$
Stores: $s ::= [\ell_1 \leftarrow v_1 \ldots, \ell_n \leftarrow v_n]$

The evaluation relation, defined by the inference rules in Fig. 1, is written $\varphi, e, s \vdash a \rightarrow r$, meaning that in evaluation environment e, initial store s, and store control φ, the source term a evaluates to the result r, which is either `err` if an illegal write was detected or a pair v/s' of a value v for the source term and a modified store s'.

The only unusual ingredient in this semantics is the φ component, which maps store locations to sets of values: if $\varphi(\ell)$ is defined, values written to the location ℓ must belong to the set $\varphi(\ell)$, otherwise a run-time error `err` is generated; if $\varphi(\ell)$ is undefined, any value can be stored at ℓ. (See rules 10 and 11.) For instance, taking $\varphi(\ell) = \emptyset$ prevents any assignment to ℓ.

The rules for propagating the `err` result and aborting execution (rules 12–21) are the same rules as for propagating run-time type errors (`wrong`) in [39]; the only difference is that we have no rules to detect run-time type errors, thus making no difference between run-time type violations and non-terminating programs (no derivations exist in both cases): the standard type soundness theorems show that type violations cannot occur at run-time in well-typed source terms.

An unusual aspect of our formalism is that the store control φ must be given at the start of the execution. The reason is that, with big-step operational semantics, it does not suffice to perform a regular evaluation $e, s \vdash a \rightarrow v/s'$ and observe the differences between s and s' to detect illegal writes. For one thing, we would not observe temporary assignments, where a malicious applet writes illegal values to a sensitive location, then restores the original values before terminating. Also, we could not say anything about non-terminating terms: the applet could perform illegal writes, then enter an infinite loop to avoid detection. By providing the store control φ in advance, we ensure that the first write error will be detected immediately and reported as the `err` result.

Unfortunately, this provides no way to control stores to locations created during the evaluation (rule 8 chooses these locations outside of $\mathrm{Dom}(\varphi)$, meaning that writes to these locations will be free): only preexisting locations can be sensitive. (This can be viewed as an inadequacy of big-step semantics, and a small-step, reduction-based semantics would fare better here. However, several semantic features that play an important role in our study are easier to express in a big-step semantics than in a reduction semantics: the clean separation between browser-supplied environment and applet-supplied source term, and the ability to interpret abstract type names by arbitrary sets of values.)

3 Reachability-based security

3.1 Simple reachability

The first security property for our calculus formalizes the idea that an applet can only write to locations that are reachable from the initial environment in which

it executes, or that are created during the applet's execution. For instance, if the references representing files are not reachable from the execution environment given to applets, then no applet can write to a file.

Reachability, here, is to be understood in the garbage collection sense: a location is reachable if there exists a path in the memory graph from the initial environment to the location, following one or several pointers. More formally, we define the set $RL(v, s)$ of locations reachable from a value v in a store s by the following equations:

$$\begin{aligned}
RL(b, s) &= \emptyset \\
RL(\lambda x.a[e], s) &= RL(e, s) \\
RL((v_1, v_2), s) &= RL(v_1, s) \cup RL(v_2, s) \\
RL(\ell, s) &= \{\ell\} \cup RL(s(\ell), s) \\
RL(e, s) &= \bigcup_{x \in \mathrm{Dom}(e)} RL(e(x), s)
\end{aligned}$$

The definition above is not well-founded by induction on v, since in the fourth case the value $s(\ell)$ is arbitrarily large and may contain ℓ again. It should be viewed as a fixpoint equation $RL = F(RL)$, where F is an increasing operator; RL is, then, the smallest fixpoint of that operator.

If p is a set of sensitive locations, we define $Prot(p)$ as the store control that maps locations $\ell \in p$ to $\emptyset$, thus disallowing all writes on ℓ, and is undefined on locations $\ell \notin p$. We can now formulate the first security property:

Security property 1 *Let p be a set of sensitive locations. If $p \cap RL(e, s) = \emptyset$, then for all applets a, we have $Prot(p), e, s \vdash a \not\rightarrow \texttt{err}$.*

In other words, if none of the locations in p is reachable from the initial environment e and store s, then no applet a can trigger an error by writing to a location in p: the applet will either terminate normally or loop. The proof of this property is a simple inductive argument on the evaluation derivation, using the following lemma as the induction hypothesis:

Lemma 1. *If $p \cap RL(e, s) = \emptyset$ and $Prot(p), e, s \vdash a \rightarrow r$, then $r \neq \texttt{err}$. Instead, $r = v/s'$ for some v and s'. Moreover, $p \cap RL(v, s') = \emptyset$, and for all values w such that $p \cap RL(w, s) = \emptyset$, we have $p \cap RL(w, s') = \emptyset$.*

It is important to remark that Property 1 holds in practice only if the evaluation of a on an actual processor adheres to the rules given in Fig. 1. For instance, the rules do not allow the evaluation of $b := a$ where b is an integer constant, but an untyped or weakly-typed implementation might execute $b := a$ without crashing nor reporting an error in the case where b is a valid memory address; this would allow the applet to write to arbitrary locations. We rely on a being evaluated by a type-safe implementation to ensure that this cannot happen: either the evaluator must perform run-time type-checking, or a must be well-typed in some sound static type system (such as the one presented below in section 4).

Property 1, though simple and already well-known in the field of garbage collection [26], establishes a number of properties without which no security is possible at all. First, our language has safe pointers: locations cannot be forged by casting well-chosen integers. Second, automatic memory management (garbage collection) is feasible and does not weaken security: a location that becomes unreachable remains unreachable; moreover, newly allocated locations are always initialized; therefore, unreachable locations can be reused safely. Third, the language enforces lexical scoping: the execution of the applet depends only on the environment in which it proceeds; the applet does not have access to the full execution environment of the browser — as would be the case in a dynamically-scoped language, such as Emacs Lisp, or a language with special constructs to access the environment of the caller, such as Tcl.

3.2 Reachability and procedural abstraction

In defining reachable locations, we have treated closures like tuples (as garbage collectors do): the locations reachable from $\lambda x.a[e]$ are those reachable from $e(y)$ for some $y \in \mathrm{Dom}(e)$. There is, however, a big difference between closures and tuples. Tuples are passive data structures: any piece of code that has access to the tuple can then obtain pointers to the components of the tuple. Closures are active data structures: only the code part of the closure can access directly the data part of the closure (the values of the free variables); other code fragments can only apply the closure, but not access the data part directly. In other words, the code part of a closure mediates access to the data part. This property is often referred to as procedural abstraction [34].

For instance, consider the following function, similar to many Unix system calls, where `uid` is a reference holding the identity of the caller (applet or browser):

```
λx. if !uid = browser
    then do something with high privileges
    else do something with low privileges
```

Assume this function is part of the applet environment, but not the reference `uid` itself. Then, there is no way that the applet can modify the location of `uid`, even though that location is reachable from the environment.

A less obvious example, where the reference `uid` is not trivially read-only, is the following function in the style of the Unix `setuid` system call:

```
λnewid. if !uid = browser
        then uid := newid
        else raise an error
```

Assuming `uid` is not initially `browser`, an applet cannot change `uid` by calling this function.

Procedural abstraction can be viewed as the foundation for access control lists and similar programming techniques, which systematically encapsulate resources inside functions that check the identity and credentials of the caller

before granting access to the requested resources. For instance, a file opening function contains the whole data structure representing the file system in its closure, but grants access only to files with suitable permissions. Thus, while all files are reachable from the closure of the `open` function, only those that have suitable permissions can be modified by the caller.

To formalize these ideas, we set out to define the set of locations $ML(v, s)$ that are actually modifiable (not merely reachable) from a value v and a store s, and show that if a location ℓ is not in $ML(e, s)$, then any applet evaluated in the environment e does not write to ℓ. This result is stronger than Property 1 because the location ℓ that is not modifiable from e in s can still be reachable (via closures) from e in s.

For passive data structures (locations, tuples), modifiability coincides with reachability: a location is modifiable from v/s if a sequence of `fst`, `snd`, and ! operations applied to v in s evaluates to that location. The difficult case is defining modifiable locations for a function closure. The idea is to consider all possible applications of the closure to an argument: a location ℓ is considered modifiable from the closure only if one of those applications writes to the location, or causes the location to become modifiable otherwise.

More precisely, let e_{api} be the execution environment given to applets, and let $c = \lambda x.a[e]$ be one of the closures contained in e_{api}. A location ℓ is modifiable from c in store s if there exists a value v such that the following conditions hold:

Condition 1. The application of the closure to v causes ℓ to be modified, i.e. $\{\ell \mapsto \emptyset\}, e\{x \leftarrow v\}, s \vdash a \rightarrow \texttt{err}$.

Example: Let e be the environment $[r \leftarrow \ell]$. Then, ℓ is reachable from $(\lambda x.\ r := x + 1)[e]$ in any store, since any application of the closure causes ℓ to be assigned.

Condition 2. Alternatively, the application of c to v does not modify ℓ, but returns a result value and a new store from which ℓ is modifiable.

Example: With e as in the previous example, ℓ is reachable from $(\lambda x.\, (r, 1))[e]$, since any application of that closure returns a pair with ℓ as first component. An applet can thus write to ℓ by applying the closure, extracting ℓ from the returned result, and assigning ℓ directly.

Condition 3. Alternatively, the application of c to v modifies a reference accessible to the applet in such a way that ℓ becomes modifiable from that reference.

Example: Consider $c = (\lambda p.\ p := r)[e]$, where $e = [r \leftarrow \ell]$ as usual. Applying c to a location ℓ', we obtain a modified store s' such that $s'(\ell') = \ell$, i.e. ℓ is modifiable from ℓ' in s' and can be modified by the applet.

Condition 4. Alternatively, the application of c to v assigns references internal to the browser functions in such a way that ℓ becomes modifiable from the environment e_{api} in the store s' at the end of the application.

Example: Consider the following environment e_{api}:

$$e_{api}(f) = (\lambda x.\ n := !n + 1)[n \leftarrow \ell_n]$$
$$e_{api}(g) = (\lambda x.\ \texttt{if}\ !n \geq 1\ \texttt{then}\ r := 0)[n \leftarrow \ell_n; r \leftarrow \ell]$$

In a store s such that $s(\ell_n) = 0$, the location ℓ is not modifiable from $e_{api}(g)$. However, ℓ is modifiable from $e_{api}(f)$ in s, since one application of that closure returns a store s' such that $s'(\ell_n) = 1$, and in that store s', ℓ is modifiable from $e_{api}(g)$: any application of $e_{api}(g)$ with initial store s' writes to ℓ.

Condition 5. In conditions 1–4 above, it must be the case that the location ℓ found to be modifiable from the closure c is not actually modifiable from the argument v passed to the closure. Otherwise, we would not know whether the location really "comes from" the closure c, or is merely modified by the applet-provided argument v.

Example: Consider the higher-order function $c = (\lambda f.\ f(0))[\emptyset]$. If we apply c to $(\lambda n.\ r := n)[r \leftarrow \ell]$, we observe a write to location ℓ. However, ℓ should not be considered as modifiable from c, since it is also modifiable from the argument given to c.

As should now be apparent from the conditions 1–5 above, the notion of modifiability raises serious problems, both practical and technical. On the practical side, the set of modifiable locations $ML(v, s)$ is not computable from v and s: in the closure case, we must consider infinitely many possible arguments. Thus, a full mathematical proof is needed to determine $ML(v, s)$.

Moreover, modifiable locations cannot be determined locally. As condition 4 shows, the modifiable locations of a closure depend on the modifiable locations of all functions from the applet environment e_{api}. Thus, if we manage to determine $ML(e_{api}, s)$, then add one single function to the applet environment, we must not only determine the modifiable locations from the new function, but also reconsider all other functions in the environment to see whether their modifiable locations have changed. This is clearly impractical. Hence, the notion of modifiability is not effective and is interesting only from a semantic viewpoint and as a guide to derive decidable security criteria in the sequel.

On the technical side, the conditions 1–5 above do not lead to a well-founded definition of the sets of modifiable locations $ML(v, s)$. The problem is condition 5 (the requirement that the location must not be modifiable from the argument given to the closure): viewing conditions 1–4 as a fixpoint equation for some operator, that operator is not increasing because of the negation in condition 5.

In appendix B, we tackle this problem and show that non-modifiable locations are indeed never modified in the particular case where the applet's environment e_{api} is well-typed and its type E_{api} does not contain any `ref` types, so that no references are exchanged directly between the applet and its environment. In the remainder of this paper, we abandon the notion of modifiability in its full generality, and develop more effective techniques to restrict writes to reachable locations, relying on type-based instrumentation of the browser code.

4 The type system

We now equip our language with a simple type system. The type system is based on simply-typed λ-calculus, with the addition of named, user-defined types. Despite its simplicity, this type system does not restrict drastically the expressiveness of our language. In particular, recursive functions can still be defined using references [39]. The type algebra is:

Types:	$\tau ::= \iota$	base type (`int`, `string`, etc.)
	$\mid t$	named type
	$\mid \tau_1 \to \tau_2$	function type
	$\mid \tau_1 \times \tau_2$	product type
	$\mid \tau\ \texttt{ref}$	reference type

Conversely, we enrich the syntax of terms with two new constructs: explicit coercions to and from named types, and run-time validation of values of named types.

Terms:	$a ::= \ldots$	(as before)
	$\mid \tau(a)$	coercion to type τ
	$\mid OK_t(a)$	run-time validation at type t

The typing rules for the calculus are shown in Fig. 3, and the extra evaluation rules for the new constructs in Fig. 2.

$$\frac{\varphi, e, s \vdash a \to r}{\varphi, e, s \vdash \tau(a) \to r} \quad (22)$$

$$\frac{\varphi, e, s \vdash a \to v/s' \qquad t \in \mathrm{Dom}(PV) \text{ implies } v \in PV(t)}{\varphi, e, s \vdash OK_t(a) \to v/s'} \quad (23)$$

$$\frac{\varphi, e, s \vdash a \to \texttt{err}}{\varphi, e, s \vdash OK_t(a) \to \texttt{err}} \quad (24)$$

Fig. 2. Extra evaluation rules for coercions and run-time checks

4.1 Named types

The overall approach followed in section 5 is to identify groups of references having the same type, and apply a given security policy to all of them. However, types from the simply-typed λ-calculus are too coarse for this purpose: references of type `string ref` can hold many different kinds of data, such as messages, filenames, and cryptographic keys; clearly, different security restrictions must be applied to these different kinds of strings.

$$E \vdash x : E(x) \quad (25) \qquad E \vdash b : \mathit{Typeof}(b) \quad (26)$$

$$\frac{E\{x \leftarrow \tau'\} \vdash a : \tau}{E \vdash \lambda x.a : \tau' \to \tau} \ (27) \qquad \frac{E \vdash a_1 : \tau' \to \tau \qquad E \vdash a_2 : \tau'}{E \vdash a_1(a_2) : \tau} \ (28)$$

$$\frac{E \vdash a_1 : \tau_1 \qquad E \vdash a_2 : \tau_2}{E \vdash (a_1, a_2) : \tau_1 \times \tau_2} \ (29) \qquad \frac{E \vdash a : \tau_1 \times \tau_2}{E \vdash \mathtt{fst}(a) : \tau_1} \ (30) \qquad \frac{E \vdash a : \tau_1 \times \tau_2}{E \vdash \mathtt{snd}(a) : \tau_2} \ (31)$$

$$\frac{E \vdash a : \tau}{E \vdash \mathtt{ref}(a) : \tau\ \mathtt{ref}} \ (32) \qquad \frac{E \vdash a : \tau\ \mathtt{ref}}{E \vdash !a : \tau} \ (33)$$

$$\frac{E \vdash a_1 : \tau\ \mathtt{ref} \qquad E \vdash a_2 : \tau}{E \vdash (a_1 := a_2) : \mathtt{unit}} \ (34)$$

$$\frac{E \vdash a : \tau \qquad \tau = \mathit{TD}(t)}{E \vdash t(a) : t} \ (35) \qquad \frac{E \vdash a : t \qquad \tau = \mathit{TD}(t)}{E \vdash \tau(a) : \tau} \ (36)$$

$$\frac{E \vdash a : t}{E \vdash \mathit{OK}_t(a) : t} \ (37)$$

Fig. 3. Typing rules

To this end, we introduce named types t, defined by a mapping TD for type names to type expressions, stating that the type t is interconvertible with its implementation type $TD(t)$. For instance, we could introduce a type `filename`, defined to be equal to `string`. An expression e of type `string` cannot be used implicitly with type `filename`; an explicit injection into `filename`, written `filename`(e), is required. Conversely, accessing the string underlying an expression e' of type `filename` is achieved by `string`(e'). (See rules 35 and 36.) In this respect, our named types behave very much like the `is new` type definition in Ada, and unlike type abbreviations in ML. Making the coercions explicit facilitates the definition of the program transformations in section 5, ensuring in particular that each term has a unique type.

The mapping TD of type definitions is essentially global: type definitions local to an expression are not supported. Still, it is possible to type-check some terms against a set of type definitions TD' that is a strict subset of TD, thus rendering the named types not defined in TD' abstract in that term. We will use this facility in section 5.2 and 5.3 to make named types abstract in the applet.

4.2 Run-time validation of values

The other unusual feature of our type system is the family of operators OK_t (one for each named type t) used to perform run-time validation of their argument. For each named type t, we assume given a set $PV(t)$ of permitted values for type t. (We actually allow $PV(t)$ to be undefined for some types t, which we

take to mean that all values of type t are valid.) The expression $OK_t(e)$ checks whether the value of e is in $PV(t)$; if yes, it returns the value unchanged; if no, it aborts the execution of the applet and reports an error. In the evaluation rules, the "yes" case corresponds to rule 23; there is no rule for the "no" case, meaning that no evaluation derivation exists if an OK_t test fails. In effect, we do not distinguish between failure of OK_t and non-termination. At any rate, we must not return the `err` result when OK_t fails: no write violation has occurred yet.

By varying $PV(t)$, we can control precisely the values of type t that will pass run-time validation. For instance, $PV(\texttt{filename})$ could consist of all strings referencing files under the applet's temporary directory `/tmp/applet.`x, a new directory that is created empty at the beginning of the applet's execution. Combined with the techniques described in section 5, this would ensure that only files in this temporary directory can be accessed by the applet. Similarly, the set $PV(\texttt{widget})$ could consist of all GUI widget descriptors referring to widgets that are children of the applet's top widget, thus preventing the applet from interacting with widgets belonging to the browser. Other examples of run-time validation include checking cryptographic signatures on the applet code itself or on sensitive data presented by the applet.

In practice, validation $OK_t(e)$ involves not only the value of its argument e, but also external information such as the identity of the principal, extra capability arguments passed to the validation functions, and possibly user replies to dialog boxes. A typical example is the Java `SecurityManager` class, which determines the identity of the principal by inspection of the call stack [43]. For simplicity, we still write OK_t as a function of the value of its argument.

The evaluation rule for OK_t assumes of course that membership in $PV(t)$ is decidable. This raises obvious difficulties if t stands for a function type, at least if the domain type is infinite. Difficulties for defining $PV(t)$ also arise if t is a reference type: checking the current contents of the references offers no guarantees with respect to future modifications; checking the locations of the references against a fixed set of locations is very restrictive. For those reasons, we restrict ourselves to types t that are defined as algebraic datatypes: type expressions obtained by combining base types with datatype constructors such as `list` or tuples, but not with `ref` nor the function arrow.

4.3 Type soundness

To relate the typing rules to the dynamic semantics, we define a semantic typing relation $S \models v : \tau$ saying whether the value v is a semantically correct value for type τ. The S component is a store typing, associating types to store locations. We simultaneously extend the $\models$ relation to stores and store typings ($\models s : S$), and to evaluation environments and typing environments ($S \models e : E$). The definition of $\models$ is shown below, and is completely standard [39, 23].

- $S \models b : \iota$ if $\mathit{Typeof}(b) = \iota$
- $S \models v : t$ if $S \models v : TD(t)$

- $S \models \lambda x.a[e] : \tau_1 \rightarrow \tau_2$ if there exists a typing environment E such that $S \models e : E$ and $E \vdash \lambda x.a : \tau_1 \rightarrow \tau_2$
- $S \models (v_1, v_2) : \tau_1 \times \tau_2$ if $S \models v_1 : \tau_1$ and $S \models v_2 : \tau_2$
- $S \models \ell : \tau\ \texttt{ref}$ if $\tau = S(\ell)$
- $S \models e : E$ if $\mathrm{Dom}(e) = \mathrm{Dom}(E)$ and for all $x \in \mathrm{Dom}(e)$, $S \models e(x) : E(x)$
- $\models s : S$ if $\mathrm{Dom}(s) = \mathrm{Dom}(S)$ and for all $\ell \in \mathrm{Dom}(s)$, $S \models s(\ell) : S(\ell)$.

Using the semantic typing relations defined above, we then have the familiar strong soundness property below for the type system. We say that a store typing S' extends another store typing S if $\mathrm{Dom}(S') \supseteq \mathrm{Dom}(S)$, and for all $\ell \in \mathrm{Dom}(S)$, we have $S'(\ell) = S(\ell)$. Remark that semantic typing is stable under store extension: if $S \models v : \tau$ and S' extends S, we also have $S' \models v : \tau$.

Lemma 2 (Type soundness). *Assume $E \vdash a : \tau$ and $\models s : S$ and $S \models e : E$. If $\varphi, e, s \vdash a \rightarrow v/s'$, then there exists a store typing S' extending S such that $S' \models v : \tau$ and $\models s' : S'$.*

Proof. The proof is a simple inductive argument on the evaluation derivation; see [23, Prop. 3.6] for details.

5 Type-based security properties

The type-based security properties we develop in this section are based on a common idea: assuming all sensitive references have types of the form $t\ \texttt{ref}$, we instrument the functions composing the execution environment by inserting OK_t run-time checks at certain program points, in order to prevent illegal writes to references of type $t\ \texttt{ref}$. The applets themselves are not instrumented, of course, since their source code is generally unavailable. All we know about the applet is that it is well typed in a given typing environment and set of type definitions. It is the combination of this well-typing with the instrumented environment functions that guarantees security.

To illustrate the three instrumentation schemes proposed below, we use pictures in the style of Fig. 4 showing the flow of values between the applet and its execution environment. We focus on values of a named type t whose implementation type is $\tau = TD(t)$. The goal of our instrumentation schemes is to make sure that only checked values of type t (solid arrows in the pictures) can reach an assignment on a $t\ \texttt{ref}$ performed by the execution environment.

5.1 Instrumentation of writes and procedural abstraction

The first transformation we consider inserts an OK_t check before any write to a reference of type $t\ \texttt{ref}$ with $t \in \mathrm{Dom}(PV)$. This way, we are certain that if a sensitive reference has type $t\ \texttt{ref}$, the environment functions will always store in it values that belong to $PV(t)$. See Fig. 5 for a graphical illustration. Formally, we define the instrumentation scheme IW, operating on terms a^τ annotated with their type τ, as follows:

$$IW((a^{t\ \texttt{ref}} := b^t)^{\texttt{unit}}) = IW(a^{t\ \texttt{ref}}) := OK_t(IW(b^t)) \quad \text{if } t \in \mathrm{Dom}(PV)$$

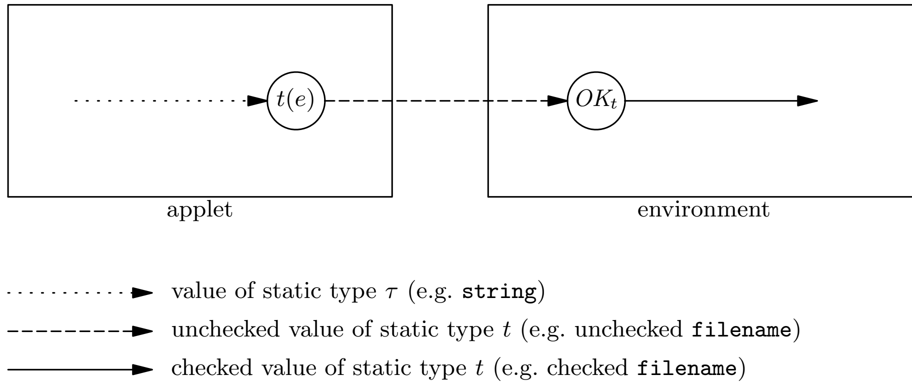

Fig. 4. Flow of values between the applet and its environment

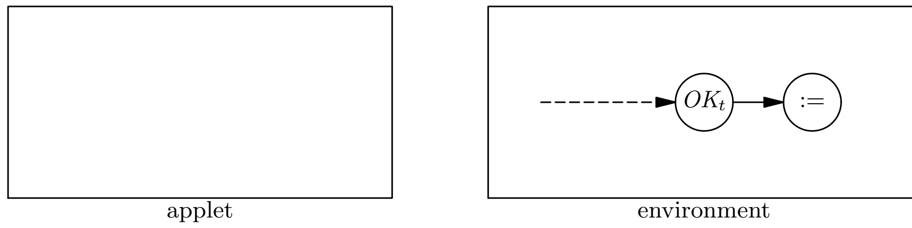

Fig. 5. Instrumentation of writes and procedural abstraction

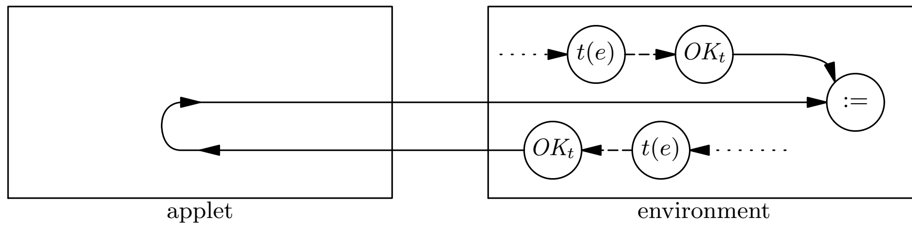

Fig. 6. Instrumentation of coercions and type abstraction

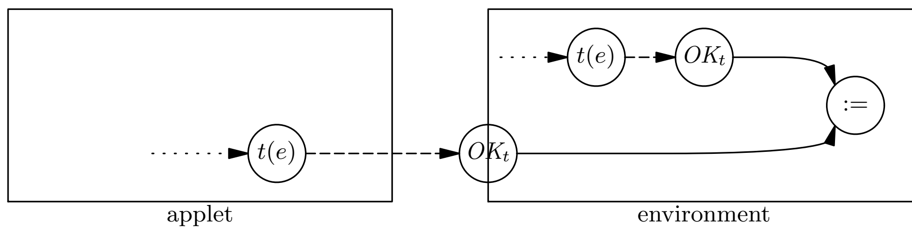

Fig. 7. Instrumenting coercions without type abstraction

On other kinds of terms, IW is a simple morphism, e.g. $IW((\lambda x.a^\tau)^{\sigma\to\tau}) = \lambda x.\, IW(a^\tau)$, etc.

It is easy to show that write errors cannot happen inside instrumented terms. Given the permitted values PV and a store typing S, we define $Prot(PV, S)$ as the store control that restricts references of type t `ref`, $t \in \mathrm{Dom}(PV)$ to values in $PV(t)$, and allows arbitrary writes to other references: $Prot(PV, S)(\ell) = PV(t)$ if $S(\ell) = t$ and $t \in \mathrm{Dom}(PV)$, and $Prot(PV, S)(\ell)$ is undefined otherwise.

Lemma 3. *Assume* $E \vdash (a^{\tau\ \mathtt{ref}} := b^\tau) : \mathtt{unit}$ *and* $S \models e : E$ *and* $\models s : S$*. Define* $\varphi = Prot(PV, S)$*. If* $\varphi, e, s \vdash IW(a^{\tau\ \mathtt{ref}}) \not\to \mathtt{err}$ *and* $\varphi, e, s \vdash IW(b^\tau) \not\to \mathtt{err}$*, then* $\varphi, e, s \vdash IW(a^{\tau\ \mathtt{ref}} := b^\tau) \not\to \mathtt{err}$*.*

Proof. If $\tau = t$ for some named type $t \in \mathrm{Dom}(PV)$, then by definition of the instrumentation scheme, the right-hand side of the assignment is of the form $OK_t(b')$ for some b', which can only evaluate to an element of $PV(t)$ by rule 23. Hence, the assignment is valid with respect to $Prot(PV, S)$. If τ is not a t type or is outside $\mathrm{Dom}(PV)$, by Lemma 2, $IW(a)$ evaluates to a location ℓ which does not have a type of the form t `ref`, $t \in \mathrm{Dom}(PV)$, and therefore such that $Prot(PV, S)(\ell)$ is undefined; hence, no write error occurs either.

Lemma 3 only provides half of the security property: it shows that writes in instrumented code are safe, but only the execution environment contains instrumented code; the applet code is not instrumented and could therefore perform illegal writes to sensitive locations, if it could access those locations. In other terms, we must make sure that all sensitive locations are encapsulated inside functions, as in section 3.2. To this end, we will restrict the type E_{api} of the applet's execution environment e_{api} to ensure that sensitive references cannot "leak" into the applet, and be assigned illegal values there. There are several ways by which a sensitive reference of type t `ref` could leak into an applet:

- The reference is exported directly in the environment, e.g. $E_{api}(x) = t$ `ref` or $E_{api}(x) = \mathtt{int} \times (t\ \mathtt{ref})$.
- The reference is returned by one of the functions of the environment, e.g. $E_{api}(f) = \mathtt{int} \to t\ \mathtt{ref}$.
- The environment contains a higher-order function such as $E_{api}(h) = (t\ \mathtt{ref} \to \mathtt{int}) \to \mathtt{int}$. The applet could get access to a sensitive reference if h passes one to its functional argument, which can be provided by the applet.
- The environment contains a function taking as argument an applet-provided reference to a t `ref`, e.g. $E_{api}(f) = t\ \mathtt{ref}\ \mathtt{ref} \to \mathtt{unit}$. The environment function f could then store a sensitive location into that t `ref ref`, from which the applet can recover the sensitive location later.

We rule out all these cases by simply requiring that no type t `ref` occurs (at any depth) in E_{api}. This leads to the following security property:

Security property 2 *Assume* $S \models e_{api} : E_{api}$ *and* $\models s : S$*. Further assume that* E_{api} *contains no occurrence of* t `ref` *for any* t*, and that all function closures*

in e and s have been instrumented with the IW scheme (that is, e and s are obtained by evaluating source terms instrumented with IW). Then, for every applet a well-typed in E, we have $Prot(PV, S), s, e \vdash a \not\rightarrow \texttt{err}$.

Proof. We consider all assignments to references that occur in the evaluation derivation for a. By Lemma 3, assignments performed by environment functions cannot cause a write error, since the right-hand side has been instrumented by IW. For assignments performed by the applet, we use a containment lemma developed in appendix A to show that the reference being assigned cannot belong to $\mathrm{Dom}(Prot(PV, S))$. More precisely, we apply the results of appendix A with T being the set of all type expressions where t `ref` occurs, L_{app} being the set of locations with types in T allocated by the applet, and L_{env} being the set of locations with types in T allocated by environment functions or initially present in e_{api}. The containment lemma (Lemma 7) then shows that the location being assigned does not belong to L_{env}. Moreover, by construction of $Prot$ and L_{env}, we have $\mathrm{Dom}(Prot(PV, S)) \subseteq L_{env}$. Thus, assignments performed by the applet do not cause write violations either. Hence, the applet cannot evaluate to `err`.

The requirement that no t `ref` occurs in E_{api} is clearly too strong: nothing wrong could happen if, for instance, one of the environment functions has type t `ref` $\rightarrow$ `unit` (the t `ref` argument is provided by the applet). We conjecture that it suffices to require that no type t `ref` occurs in E_{app} at a positive occurrence or under a `ref` constructor. However, our proof of Property 2, and in particular the crucial containment lemma, does not extend to this weaker hypothesis.

5.2 Instrumentation of coercions and type abstraction

Instead of putting run-time checks on writes to references of type t `ref`, we can ensure that all values of type $t \in \mathrm{Dom}(PV)$ that flow through the applet's execution environment always belong to $PV(t)$. This way, values stored in a t `ref` will automatically satisfy $PV(t)$ as well. This is achieved by adding checks to all creations of values of type $t \in \mathrm{Dom}(PV)$ in the execution environment, i.e. to coercions of the form $t(a)$, following the instrumentation scheme IC below:

$$IC(t(a)) = OK_t(t(IC(a))) \quad \text{if } t \in \mathrm{Dom}(PV)$$

Of course, this is not enough: the applet could forge unchecked values of type t, by direct coercion from t's implementation type, and pass them to environment functions. Hence, we also need to make the types $t \in \mathrm{Dom}(PV)$ abstract in the applet, by type-checking it with a set of type definitions TD' obtained from TD by removing the definitions of the types $t \in \mathrm{Dom}(PV)$. Then, for any $t \in \mathrm{Dom}(PV)$, the only values of type t that can be manipulated by the applet have been created and checked by the environment. This is depicted in Fig. 6.

To capture the run-time behavior of instrumented terms, we introduce a variant of the semantic typing predicate, written $PV, S \models v : \tau$, which is similar to the predicate $S \models v : \tau$ from section 4.3, with the difference that a value v belongs to a named type t only if $v \in PV(t)$ in addition to v belonging to the definition $TD(t)$ of t:

- $PV, S \models b : \iota$ if $\mathit{Typeof}(b) = \iota$
- $PV, S \models v : t$ if $PV, S \models v : \mathit{TD}(t)$ and $t \in \mathrm{Dom}(PV)$ implies $v \in PV(t)$
- $PV, S \models \lambda x.a[e] : \tau_1 \to \tau_2$ if there exists a typing environment E such that $PV, S \models e : E$ and $E \vdash \lambda x.a : \tau_1 \to \tau_2$
- $PV, S \models (v_1, v_2) : \tau_1 \times \tau_2$ if $PV, S \models v_1 : \tau_1$ and $PV, S \models v_2 : \tau_2$
- $PV, S \models \ell : \tau\ \mathtt{ref}$ if $\tau = S(\ell)$
- $PV, S \models e : E$ if $\mathrm{Dom}(e) = \mathrm{Dom}(E)$ and for all $x \in \mathrm{Dom}(e)$, $PV, S \models e(x) : E(x)$
- $PV \models s : S$ if $\mathrm{Dom}(s) = \mathrm{Dom}(S)$ and for all $\ell \in \mathrm{Dom}(s)$, $PV, S \models s(\ell) : S(\ell)$.

We then have the following characterization of the the behavior of terms instrumented with IC:

Lemma 4. *Assume $E \vdash a : \tau$ and $PV, S \models e : E$ and $PV \models s : S$. Further assume that all closures contained in e and s have function bodies instrumented with IC. If $\mathit{Prot}(PV, S), e, s \vdash IC(a) \to r$, then $r \neq \mathtt{err}$; instead, r is of the form v/s', and there exists a store typing S' extending S such that $PV, S' \models v : \tau$ and $PV \models s' : S'$.*

Compared with Lemma 2, we now use the more restrictive interpretation of named types PV, and require that all terms occurring in the evaluation derivation have been instrumented with IC. The proof is essentially identical to that of Lemma 2.

Notice that an applet a well-typed within the restricted set TD' of type definitions contains no coercions $t(a)$ for any $t \in \mathrm{Dom}(PV)$, hence is equal to $IC(a)$. Thus, Lemma 4 applies not only to the terms from the execution environment, but also to the applet itself. In a sense, this lemma can be viewed as a parametricity result, in that it shows soundness for arbitrary interpretations PV of the types left abstract in TD'. From this remark, we immediately obtain the following security property:

Security property 3 *Let e be the execution environment for applets, and s the initial store. Assume that all function closures in e and s have been instrumented with the IC scheme (that is, e and s are obtained by evaluating source terms instrumented with IC). Assume $PV \models s : S$ and $PV, S \models e : E$. Then, for every applet a well-typed in E and in the restricted set TD' of type definitions, we have $\mathit{Prot}(PV, S), s, e \vdash a \not\to \mathtt{err}$.*

Property 3 can be viewed as a formal justification for capability-based systems: by making the type of capabilities abstract to the applets, run-time security checks are necessary only at points where new capabilities are constructed and returned to the applet; capabilities presented by the applet can then be trusted without further checks.

Unlike Property 2, Property 3 does not require that types $t\ \mathtt{ref}$ do not occur in the typing environment E. Indeed, once t is made abstract, it is perfectly safe to make references of type $t\ \mathtt{ref}$ accessible to the applet: the applet can then

write to them, but only write safe values of type t. Hence, sensitive references no longer need to be systematically wrapped inside functions. As Reynolds points out [34], type abstraction and procedural abstraction are two orthogonal ways to protect data.

Another advantage of the approach described in this section over the instrumentation of writes described in section 5.1 is that it often leads to fewer run-time checks. In particular, checks at coercions can sometimes be proven redundant and therefore can be eliminated. Consider the following function that adds a `.old` suffix to a file name:

$$\lambda\texttt{f}.\,OK_{\texttt{filename}}(\texttt{filename(concat(string(f),".old"))})$$

With the definition of $PV(\texttt{filename})$ given in section 4.2, it is easy to show that if f belongs to $PV(\texttt{filename})$, then so does the concatenation of f and `.old`. Hence, the OK test can be removed.

Of course, not all run-time tests can be removed this way: consider what happens if the suffix is given as argument:

$$\lambda\texttt{f}.\lambda\texttt{s}.\,OK_{\texttt{filename}}(\texttt{filename(concat(string(f),s))})$$

and the applet passes a suffix `s` starting with “`../`”.

5.3 Instrumenting coercions without type abstraction

In some cases, the types $t \in \mathrm{Dom}(PV)$ cannot be made abstract in the applet, e.g. because it would make writing the applet too inconvenient, or entail too much run-time overhead. We can adapt the approach presented in section 5.2 to these cases, by reverting to procedural abstraction and putting checks not only at coercions, but also on all values of types $t \in \mathrm{Dom}(PV)$ that come from the applet. (This matches current practice in Unix kernels, where parameters to system calls are always checked for validity on entrance to the system call.) Figure 7 depicts this approach.

The checking of values coming from the applet is achieved by a standard wrapping scheme applied to all functions of the execution environment, inserting OK_t coercions at all negative occurrences of types $t \in \mathrm{Dom}(PV)$. For instance, if the execution environment needs to export a function $f : t \to t$, it will actually export the function $\lambda x.f(OK_t(x))$, which validates its argument before passing it to the original function.

We formalize these ideas in a slightly different way, in order to build upon the results of section 5.2. Start from an applet environment defined by top-level bindings of the form

$$\texttt{let}\ f_i : \tau_i = a_i$$

We assume given a set TD' of type definitions against which the a_i and the applets are type-checked, and a valuation PV' assigning permitted values to named types in TD'.

We first associate a new named type $\hat{t}$ to each sensitive type $t \in \mathrm{Dom}(PV')$. The type $\hat{t}$ is defined as synonymous with t, and is intended to represent those

values of type t that have passed run-time validation. We define the $\hat{t}$ types by taking

$$TD = TD' \oplus [\hat{t} \mapsto t \mid t \in \mathrm{Dom}(PV')]$$

and restrict the values they can take using the valuation PV defined by $PV(\hat{t}\,) = PV'(t)$ and PV undefined on other types.

Let Σ be the substitution $\{t \leftarrow \hat{t} \mid t \in \mathrm{Dom}(PV')\}$. We transform the bindings for the applet environment as follows:

$$\mathtt{let}\ f_i : \tau_i = W^+(IC(\Sigma(a_i)) : \tau_i)$$

That is, we rewrite the terms a_i to use the type $\hat{t}$ instead of t for all $t \in \mathrm{Dom}(PV')$; then apply the IC instrumentation scheme to it, thus adding an $OK_{\hat{t}}$ check to each coercion $\hat{t}(a)$; finally, apply the W^+ wrapping scheme to the instrumented term, in order to perform both validation and coercion from t to $\hat{t}$ on entrance, and the reverse coercion from $\hat{t}$ to t on exit. Wrapping is directed by the expected type for its result, and is contravariant with respect to function types. We thus define both a wrapping scheme W^+ for positive occurrences of types and another W^- for negative occurrences.

$$\begin{aligned}
W^+(a : \iota) = W^-(a : \iota) &= a \\
W^+(a : t) &= t(a) \quad \text{if } t \in \mathrm{Dom}(PV) \\
W^-(a : t) &= OK_{\hat{t}}(\hat{t}(a)) \quad \text{if } t \in \mathrm{Dom}(PV) \\
W^+(a : t) = W^-(a : t) &= a \quad \text{if } t \notin \mathrm{Dom}(PV) \\
W^+(a : \tau_1 \times \tau_2) &= (W^+(\mathtt{fst}(a) : \tau_1), W^+(\mathtt{snd}(a) : \tau_2) \\
W^-(a : \tau_1 \times \tau_2) &= (W^-(\mathtt{fst}(a) : \tau_1), W^-(\mathtt{snd}(a) : \tau_2) \\
W^+(a : \tau_1 \rightarrow \tau_2) &= \lambda x. W^+(a(W^-(x : \tau_1)) : \tau_2) \\
W^-(a : \tau_1 \rightarrow \tau_2) &= \lambda x. W^-(a(W^+(x : \tau_1)) : \tau_2) \\
W^+(a : \tau\ \mathtt{ref}) = W^-(a : \tau\ \mathtt{ref}) &= a \\
&\quad \text{if no } t \in \mathrm{Dom}(PV) \text{ occurs in } \tau
\end{aligned}$$

Wrapping is not defined on reference types containing a $t \in \mathrm{Dom}(PV)$ because there is no way to validate these references so that they are protected against future modifications.

It is easy to see that the transformed bindings are well-typed: $\Sigma(a_i)$ has type $\Sigma(\tau_i)$; the IC instrumentation preserves typing; the W^+ wrapping applied to a term of type $\Sigma(\tau_i)$ returns a term of type τ_i, as shown by a simple induction over τ_i.

Moreover, the right-hand sides of the bindings always perform an $OK_{\hat{t}}$ check before each coercion to a type $\hat{t}$: this is ensured by the IC instrumentation for the coercions initially in $\Sigma(a_i)$, and by definition of the wrapping scheme for the coercions introduced by the wrapping.

Finally, the applets themselves are still type-checked in the original set TD' of type definitions, in which the types $\hat{t}$ are not defined, and thus abstract for the applet.

We are therefore back to the situation studied in section 5.2: the types $\hat{t}$ are abstract in the applets and all coercions to $\hat{t}$ are instrumented in the applet environment. Thus, by Property 3, we obtain that the values of references with types $\hat{t}$ `ref` always remain within $PV(\hat{t}) = PV'(t)$ during the execution of any well-typed applet.

Security property 4 *Let e be the execution environment for applets and s the initial store. Assume that e and s are obtained by evaluating a set of transformed bindings* `let` $f_i : \tau_i = W^+(IC(\Sigma(a_i)) : \tau_i)$ *as described above. Assume $PV \models s : S$ and $PV, S \models e : E$. Then, for every applet a well-typed in E and in the initial set TD' of type definitions, we have $Prot(PV, S), s, e \vdash a \not\rightarrow$* `err`.

6 Connections with object-oriented languages

Although the language used for this work is functional, the techniques developed here translate reasonably well to object-oriented languages. Procedural abstraction as presented in section 3.2 corresponds most closely to Smalltalk-style private instance variables: just as variables in a closure environment can only be accessed by the code associated with that closure, private instance variables of a Smalltalk object can only be accessed by the methods of that object. Java does not offer a strictly equivalent mechanism: the `private` modifier makes instance variables accessible not only to the methods of the object, but also to methods of other objects of the same class. Still, the visibility rules associated with the package mechanism and the `private` and `protected` modifiers also ensure some degree of procedural abstraction, since they restrict the set of methods that can access a given instance variable.

Similarly, type abstraction as exploited in section 5.3 corresponds most closely to `final` classes in Java. A `final` class containing only `private` fields and no default constructors offers the same level of guarantees as our abstract types: the applet cannot tamper with the fields of an existing object, nor create an object with arbitrary initial values for the fields. In non-`final` classes, the visibility rules might still ensure some degree of type abstraction, though it is unclear how much is guaranteed. Subtype polymorphism in systems such as $F_{<:}$ [11] also provide some amount of type abstraction, but this is not so in Java because the actual type of an object can be tested at run-time (downcasts).

7 Application: safety in the MMM browser

MMM [35] is a Web browser with applets developed by the second author. MMM ensures safe execution of applets using various techniques similar in spirit to those formulated earlier in this paper. (Essentially identical techniques are now used in the SwitchWare project at U. Penn [4, 5], an active network infrastructure that allows safe downloading of applets on network routers.) Indeed, the formalization presented in this paper grew out of a desire to make more systematic and prove correct the security techniques used in MMM. This section

presents the main techniques used in the MMM safe execution environment for applets, and relates them with the formal results we have obtained in this paper.

7.1 Type-safe dynamic linking

The MMM browser is written in Objective Caml [24] and compiled to bytecode by the Caml bytecode compiler. The bytecode is then executed by the Caml virtual machine. Applets are also compiled to Caml bytecode, then loaded in memory and linked with the browser by the Caml dynamic linker (the `Dynlink` library).

The dynamic linker can be configured to restrict the set of external modules and C primitives that the object file can reference. MMM uses this feature to enforce lexical scoping for applets: only the modules of MMM that comprise the applet execution environment can be referenced by the applet object code; modules internal to MMM are not made accessible to the applet.

The Caml bytecode was designed before applets were fashionable, and is therefore ill-suited to JVM-style bytecode verification. In particular, a number of low-level optimizations are performed in this bytecode (such as erasing type annotations and conflating several high-level types into the same machine-level representations) that makes type reconstruction on bytecode if not impossible, but at least very difficult.

Instead of bytecode verification, the Objective Caml dynamic linker relies on type annotations on the relocation information contained in object files to guarantee type correctness. Bytecode object files contain a list of names of external modules referenced in the file, along with the type interfaces of those modules. The object file thus records which interfaces for those external modules were used for type-checking the source code of the file. The dynamic linker, then, checks that those interfaces are identical to those of the implementations of those modules provided by the MMM browser. This ensures that the applet and the browser have been type-checked against the same type specifications for the applet-browser interface, and can therefore safely be linked together.

To keep them small, object files actually do not contain the whole Caml type interfaces for their imported modules, but only MD5 checksums of those interfaces. The dynamic linker then compares interfaces by comparing their checksums. Since MD5 is a hash function of cryptographic quality, the probability that two different interfaces have the same checksum is extremely low, and thus we can safely assume that if the checksums are equal, then so are the interfaces. However, this approach forces the two interfaces (the one expected by the object file and the one provided by the MMMimplementation) to be exactly identical, while, if we kept whole interfaces, we could use a more lenient comparison such as Caml's interface matching. The latter would allow more flexibility in evolving the applet-browser API.

Of course, comparing import interfaces at link-time guarantees type safety only if the object code contained in the object file is indeed type-correct with respect to the interfaces contained in the object file. Thus, we need to make sure that this object file has been produced by a correct Caml compiler (a compiler

that performs sound static type-checking and records correctly all external modules referenced) and has not been modified afterwards (e.g. by hand-modifying some of the interfaces in the object file). This can be enforced in two ways in the context of Web applets.

The first way is to transmit applets over the network as Caml source code, which is type-checked and compiled locally (by calling the Objective Caml bytecode compiler), then dynamically linked inside the browser. Unless the local host is compromised, no tampering with the object files nor the Caml type-checker itself can happen. Unlike applet systems based on source-level interpretation, we still benefit from the efficiency of the Caml bytecode interpreter.

Transmission of applets in source form is often criticized on several grounds: the source code is larger than compiled bytecode; local compilation takes too much time; applet writers do not want to publicize their source. Our experience with Caml is that bytecode object files are about the same size as the source code (unless heavily commented); the bytecode compiler is very fast and the compilation times are small compared with Internet latencies; finally, bytecode is easy to decompile and does not offer significant protection against reverse engineering.

Another way to ensure the correctness of type annotations in bytecode object files is to rely on a cryptographic signature on the object file, which is checked locally by the MMM browser against a list of trusted signers before linking the object file in memory. Unlike Microsoft-style applet signing, this signature is not necessarily made by the author of the applet, and carries no guarantees on what the applet actually does; instead, the signature is made by the person or site who performs the compilation and type-checking of the applet, and certifies only that the applet passed type-checking by an unmodified Caml compiler and that its object file has not been tampered since.

We initially envisioned having some centralized type-checking authority: one or several reputable sites (such as INRIA) that accept source code from applet developers, type-check and compile them locally, sign the object file and return it to the applet developer. The object file can then be made available on the Web. Seeing the authority's signature, browsers can trust the type information contained in the object file and use it to validate the applet against the environment they provide. In retrospect, this approach relies too much on a centralized authority to scale up to the world-wide Web. However, it is perfectly suited to the distribution of compiled applets across a restricted network such as a corporate Intranet.

7.2 The execution environment for applets

The execution environment made available to applets is composed of a number of modules that provide a safe subset of the Caml standard library, as well as access to the CamlTk graphical user interface and to selected parts of the MMM browser internals. In particular, an MMM applet can not only interact graphically with the user and trigger navigation functions, like Java applets, but also extend the MMM browser itself by installing new navigation functions,

menu items, viewers for new types of embedded documents, display functions for new HTML tags, and decoding functions for new “content-encoding” types.

The applet environment is derived from the OCaml standard library modules and the MMM implementation modules by two major techniques: hiding of unsafe functions via module thinning, and wrapping of unsafe functions with capability checks.

Module thinning: The ML module system offers the ability to take a restricted view of an existing module via a signature constraint:

```
module RestrM = (M : RestrSig)
```

Only those components of `M` that are mentioned in the signature `RestrSig` are visible in `RestrM`, and they have the types specified in `RestrSig`, which may be less precise (more abstract) than their original types in `M`. This module thinning mechanism thus supports both hiding components (functions, variables, types, exceptions, sub-modules) of a module and making some type components abstract. No code duplication occurs during thinning: the functions in `RestrM` share their code with the corresponding functions in `M`.

Large parts of the applet environment are obtained by thinning existing OCaml library modules. In the OCaml standard library, we hide by thinning all file input-output functions, as well as related system interface functions (such as reading environment variables and executing shell commands), and of course all type-unsafe operations (such as array accesses without bound checks and functions that operate on the low-level representation of data structures). For good measure, we also make abstract a few data types such as lexer buffers to hide their internal structure.

In the CamlTk GUI toolkit, we hide all functions that return widgets that may belong to the browser or to other applets. Such functions include finding the parent of a widget, finding a widget by its name, finding which widget owns the focus or the selection, etc. This is less restrictive than checking that widgets manipulated by the applet are children of its top-level widget: the applet can still open new windows and populate them with widgets unrelated with its initial top-level widget. Other functions that might affect browser widgets (such as binding events on all widgets of a given class or tag) are also removed.

Capabilities: MMM uses capabilities to control potentially harmful functionalities, such as file and network input/output, as well as extending the browser. Individual capabilities include:

- Reading files whose names match a given regular expression.
- Writing files whose names match a given regular expression.
- Accessing the contents of Web documents (e.g. the pages being displayed).
- Installing extension functions in the browser.

Applets are given an initial set of capabilities at load time, which is empty for applets loaded through the network, but allows document access and browser

extension for applets loaded from the local disk. Capabilities are of course represented by an abstract data type, to prevent an applet from forging extra capabilities.

All input/output operations as well as registration of browser extensions check the capabilities presented by the applet. If the applet does not possess the capability to perform the requested operation, the browser prompts the user via a pop-up window. The user can then refuse the operation (aborting the execution of the applet), grant permission for this particular operation, or grant permission for further operations of the same kind as well. In the latter case, the browser extends in place the capabilities of the applet. To minimize the number of times the user is prompted, an applet can also request in advance the capabilities it needs later.

Hiding capability arguments: A well-known problem with capability-based security for applets is that all sensitive functions require an extra capability argument [42]. This clutters the source code for the applet and makes it difficult to convert a piece of stand-alone code into an applet: capability arguments must be added at many different places in the source code.

In functional languages, partial application can be used to pass the capability argument only once per sensitive function. Assume defined the functions `open_in_cap : capability → string → in_channel` and `open_out_cap : capability → string → out_channel` as capability-checking variants of the normal file opening functions `open_in : string → in_channel` and `open_out : string → out_channel`. Then, an applet can begin with the following definitions:

```
let mycapa = Capabilities.get()
let open_in = open_in_capa mycapa
let open_out = open_out_capa mycapa
... code using open_in and open_out as usual...
```

The remainder of the applet can then use the functions `open_in` and `open_out` thus obtained by partial application as if they were the normal file opening functions.

MMM makes this approach more convenient by using ML functors to perform the partial applications. Functors are parameterized modules, presented as functions from modules to modules. They provide a convenient mechanism for parameterization *en masse*: rather than partially applying n functions to m parameters, a structure containing the m parameters is passed once to a functor that return a structure containing the n functions already partially applied to the parameters. In the case of MMM, we thus have a number of functors that take the applet's capability as argument and return structures defining capability-enabled variants of the standard I/O and browser interface functions, with the same interface as for standalone programs. For instance, to perform file I/O, an MMM applet does the following:

```
module MyCapa =
  struct
    let capabilities = Capabilities.get()
  end
module IO = Safeio(MyCapa)
open IO
```
... code using `open_in` *and other I/O operations as usual...*

7.3 Assessment

The MMM security architecture relies heavily on the three basic ingredients that we considered in our formalization:

- Lexical scoping ensures that applets cannot access all the modules composing the browser and the Caml standard library, but only those "safe" modules made available to the applet during linking.
- Type abstraction prevents the applet from forging or tampering with its capability list. It also ensures that the applet cannot forge file descriptors or GUI widgets, but has to go through the safe libraries to create them.
- Procedural abstraction is used systematically to wrap sensitive functions (such as opening files or network connections, as well as installing a browser extension) with capability checks.

In particular, capability checks inside environment functions are placed essentially as suggested by the wrapping scheme W^+ of section 5.3. In the MMM implementation, those capability checks were placed by hand and the same type was used for checked and unchecked data, instead of exploiting two different types as in section 5.3. One of the motivations for our formalization was to systematize the placement of those capability checks, based on typing information.

Some aspects of the MMM security architecture are not accounted for by our formalism. The first one is the necessity of taking copies of mutable objects before validating them. A prime example is character strings, which, in Caml, can be modified in place like character arrays. An applet could pass a string containing a valid file name to the file opening function, then modify the string in place between the time it is validated and the time the file is actually opened. This concurrent modification can be achieved via multiple threads or via GUI callback functions.

To avoid this attack, the execution environment must first take private copies of all strings provided by the applet, then validate and utilize the private copies of those strings. Symmetrically, the execution environment must not return strings shared with its internal data structures to the applet, but copies of those strings. Inserting those string copy operations is tedious and error-prone. From the standpoint of applet security, it would be much better to have immutable strings in the language. The guarantees offered by immutable strings could possibly be achieved just by removing all string modification primitives from the applet execution environment, but we have not investigated this approach yet.

Another issue not addressed by our framework is validation of functional values, such as document decoders and HTML tag handling functions installed by the applet. The main reason our framework does not handle function validation is that membership of a functional value in a set $PV(t)$ is in general undecidable. In the particular case of MMM, this problem does not arise: depending on the capabilities given to the applet, either all functions presented by the applet are rejected, or they are all accepted. For more complex situations, we could also move the validation inside the applet-provided function, by wrapping this function with checks on its inputs or outputs.

Resources that can be explicitly deallocated and reassigned later raise interesting problems. A prime example is Unix file descriptors, which are small integers. The type of file descriptors is abstract, and an applet cannot forge a file descriptor itself. However, it could open a file descriptor on a permitted file, close it immediately, keep the abstract value representing the descriptor, and wait until the browser opens a file or network connection and receives the same file descriptor in return. Then, the applet can do input/output directly on the file descriptor, thus accessing unauthorized files or network connections. This "reuse" attack is of course possible only with explicitly-deallocated resources: with implicit deallocation, the resources cannot be deallocated and reallocated as long as the applet keeps a handle on the resource. MMM addresses this problem by wrapping Unix file descriptors in an opaque data structure containing a "valid" bit that is set to false when the file descriptor is closed, and checked before every input/output operations.

The last MMM security feature not addressed by our framework is confidentiality. Our framework focuses on ensuring the integrity of the browser and of the host machine. The MMM applet environment contains security restrictions to ensure integrity (such as controlled write access to files), but also restrictions intended to protect the confidentiality of user data (such as controlled read accesses to files and restricted access to the network). Some restrictions address both integrity and confidentiality problems: for instance, a malicious document decoder could both distort the document as displayed by the browser (an integrity threat) and leak confidential information contained in the document to a third party (a confidentiality threat).

However, the confidentiality policy enforced by the MMM applet environment is simple and does not go beyond traditional access control. While the formal framework presented in this paper focuses on integrity via access control on writes, we believe that parts of this framework (procedural encapsulation and type abstraction) can be extended to deal with access control on reads, thus ensuring simple confidentiality properties adequate for MMM. More advanced confidentiality policies (such as allowing an applet to read local files or open network connections, but not both) are definitely beyond the scope of our framework, however.

8 Related work

8.1 Type systems for security

The work most closely related to ours is the recent formulations of Denning's information flow approach to security [13, 14] as non-standard type systems by Palsberg and Ørbaek [31], Volpano and Smith [41, 40], and Heintze and Riecke [20]. (Abadi *et al.* [2] reformulate some of those type systems in terms of a more basic calculus of dependency.) The main points of comparison with our work are listed below.

Information flow vs. integrity: The type systems developed in previous works all focus on secrecy properties, following the information flow approach. In particular, they allow high-security data to be exposed as long as no low-security code uses this data. Our work focuses on more basic integrity properties via access control. We view those integrity guarantees as a prerequisite to establishing meaningful confidentiality properties.

Imperative vs. purely functional programs: [31] and [20] consider purely functional languages in the style of the λ-calculus. This makes formulating the security properties delicate: [31] proves no security property properly speaking, only a subject reduction property that shows the internal consistency of the calculus, but not its relevance to security; [20] does show a non-interference property (that the value of a low-security expression is independent of the values of high-security parameters), but it is not obvious how this result applies to actual applet/browser interactions, especially input/output. Instead, we have followed [41, 40] and formulated our security policy in terms of in-place modifications on a store, which provides a simple and intuitive notion of security violation.

Run-time validation of data: Only [31] and our work consider the possibility of checking low-security data at run-time and promoting them to high security. In [41, 40, 20], once some data is labeled "low security", it remains so throughout the program and causes all data it comes in contact with to be marked "low security" as well. We believe that, in a typical applet/browser interaction, this policy leads to rejecting almost all applets as insecure. Run-time validation of untrusted data is essential in practice to allow a reasonable range of applets to run.

Subtyping vs. named types and coercions: All previous works consider type systems with subtyping, which provides a good match for the flow analysis approach they follow [30]. In contrast, we only use type synonyms with possibly checked coercions between a named type and its implementation type. However, the connections between subtyping and explicit coercions are well known [9], and we do not think this makes a major difference.

8.2 Security of applets

Concerning the security issues raised by applets in general, we are aware of case studies of security flaws [12], as well as informal descriptions of current and

proposed security architectures for applets [16, 44, 35, 42, 18]. Our work seems to be one of the first formal studies of applet security.

On the Java side, considerable effort has been expended in proving the soundness of the Java type system [15, 37, 29] and of the JVM bytecode verifier [33, 36, 17]. Other aspects of Java that are equally important for security, such as formalizing the visibility rules and the encapsulation guarantees they provide, have only recently started to receive attention [22]. We are not aware of any formal description of security policies for Java applets.

Proof-carrying code [28, 27] provides an elegant framework to establish the safety of mobile code, but requires general proof from the applet's developer. Our approach lies at the other end of the complexity spectrum: all we require from the applet is that it is well typed in a simple, standard type system.

9 Concluding remarks

We have identified three basic techniques for enforcing a fairly realistic security policy for applets: lexical scoping, procedural abstraction, and type abstraction. These programming techniques are of course well known, but we believe that this work is the first to characterize precisely their implications for program security.

The techniques proposed here seem to match relatively well current practice in the area of Web applets. In particular, they account fairly well for Rouaix's implementation of safe libraries in the MMM browser.

Our techniques put almost no constraints on the applets, except being well-typed in a simple, completely standard type system. The security effort is concentrated on the execution environment provided by the browser. Typing the applets in a richer type system, such as the type systems for information flow of [31, 41, 40, 20] or the effect and region system of [38], could provide more information on the behavior of the applet and enable more flexible security policies in the execution environment. However, it is impractical to rely on rich type systems for applets, because these type systems are not likely to be widely accepted by applet developers. Whether these rich type systems can be applied to the execution environment only, while still using a standard type system for the applets, is an interesting open question.

On the technical side, the proofs of the type-based security properties are variants of usual type soundness proofs. It would be interesting to investigate the security content of other classical semantic results such as representation independence and logical relations. Given the importance of communications between the applet and its environment, it could be worthwhile to reformulate our security results for a calculus of communicating processes [6, 3, 1].

Acknowledgements

This work has been partially supported by GIE Dyade under the "Verified Internet Protocols" project. A preliminary version of this paper appeared in the

proceedings of the 25th ACM symposium on Principles of Programming Languages.

References

1. M. Abadi. Secrecy by typing in security protocols. In *Theoretical Aspects of Computer Software '97*, volume 1281 of *Lecture Notes in Computer Science*, pages 611–638. Springer-Verlag, Sept. 1997.
2. M. Abadi, A. Banerjee, N. Heintze, and J. G. Riecke. A core calculus of dependency. In *26th symposium Principles of Programming Languages*, pages 147–160. ACM Press, 1999.
3. M. Abadi and A. D. Gordon. Reasoning about cryptographic protocols in the Spi calculus. In *CONCUR'97: Concurrency Theory*, volume 1243 of *Lecture Notes in Computer Science*, pages 59–73. Springer-Verlag, July 1997.
4. D. S. Alexander, W. A. Arbaugh, M. W. Hicks, P. Kakkar, A. D. Keromytis, J. T. Moore, C. A. Gunter, S. M. Nettles, and J. M. Smith. The SwitchWare active network architecture. *IEEE Network*, 12(3):29–36, 1998.
5. D. S. Alexander, W. A. Arbaugh, A. D. Keromytis, and J. M. Smith. Security in active networks. In *Secure Internet Programming*, Lecture Notes in Computer Science, pages ??–?? Springer-Verlag Inc., New York, NY, USA, 1999.
6. J.-P. Banâtre and C. Bryce. A security proof system for networks of communicating processes. Research report 2042, INRIA, Sept. 1993.
7. J.-P. Billon. Security breaches in the JDK 1.1 beta2 security API. Dyade, `http://www.dyade.fr/fr/actions/VIP/SecHole.html`, Jan. 1997.
8. N. S. Borenstein. Email with a mind of its own: the Safe-Tcl language for enabled mail. In *IFIP International Working Conference on Upper Layer Protocols, Architectures and Applications*, 1994.
9. V. Breazu-Tannen, T. Coquand, C. A. Gunter, and A. Scedrov. Inheritance as implicit coercion. *Information and Computation*, 93(1):172–221, 1991.
10. K. Brunnstein. Hostile ActiveX control demonstrated. *RISKS Forum*, 18(82), Feb. 1997.
11. L. Cardelli, S. Martini, J. C. Mitchell, and A. Scedrov. An extension of system F with subtyping. *Information and Computation*, 109(1–2):4–56, 1994.
12. D. Dean, E. W. Felten, D. S. Wallach, and D. Balfanz. Java security: Web browsers and beyond. In D. E. Denning and P. J. Denning, editors, *Internet Besieged: Countering Cyberspace Scofflaws*, pages 241–269. ACM Press, 1997.
13. D. E. Denning. A lattice model of secure information flow. *Commun. ACM*, 19(5):236–242, 1976.
14. D. E. Denning and P. J. Denning. Certification of programs for secure information flow. *Commun. ACM*, 20(7):504–513, 1977.
15. S. Drossopoulou and S. Eisenbach. Java is type safe – probably. In *Proc. 11th European Conference on Object Oriented Programming*, volume 1241 of *Lecture Notes in Computer Science*, pages 389–418. Springer-Verlag, June 1997.
16. M. Erdos, B. Hartman, and M. Mueller. Security reference model for the Java Developer's Kit 1.0.2. JavaSoft, `http://java.sun.com/security/SRM.html`, Nov. 1996.
17. S. N. Freund and J. C. Mitchell. A type system for object initialization in the Java bytecode language. In *Object-Oriented Programming Systems, Languages and Applications 1998*, pages 310–327. ACM Press, 1998.

18. L. Gong. Java security architecture (JDK1.2). JavaSoft, `http://java.sun.com/products/jdk/1.2/docs/guide/security/spec/security-spec.doc.html`, Oct. 1998.
19. J. Gosling and H. McGilton. The Java language environment – a white paper. JavaSoft, `http://java.sun.com/docs/white/langenv`, May 1996.
20. N. Heintze and J. G. Riecke. The SLam calculus: programming with secrecy and integrity. In *25th symposium Principles of Programming Languages*, pages 365–377. ACM Press, 1998.
21. D. Hopwood. Java security bug (applets can load native methods). *RISKS Forum*, 17(83), Mar. 1996.
22. T. Jensen, D. Le Métayer, and T. Thorn. Security and dynamic class loading in Java: A formalisation. In *International Conference on Computer Languages 1998*, pages 4–15. IEEE Computer Society Press, 1998.
23. X. Leroy. Polymorphic typing of an algorithmic language. Research report 1778, INRIA, 1992.
24. X. Leroy, J. Vouillon, D. Doligez, et al. The Objective Caml system. Software and documentation available on the Web, `http://caml.inria.fr/ocaml/`, 1996.
25. R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The definition of Standard ML (revised)*. The MIT Press, 1997.
26. G. Morrisett, M. Felleisen, and R. Harper. Abstract models of memory management. In *Functional Programming Languages and Computer Architecture 1995*, pages 66–77. ACM Press, 1995.
27. G. C. Necula. Proof-carrying code. In *24th symposium Principles of Programming Languages*, pages 106–119. ACM Press, 1997.
28. G. C. Necula and P. Lee. Safe kernel extensions without run-time checking. In *Proc. Symp. Operating Systems Design and Implementation*, pages 229–243. Usenix association, 1996.
29. T. Nipkow and D. von Oheimb. JavaLight is type-safe — definitely. In *25th symposium Principles of Programming Languages*, pages 161–170. ACM Press, 1998.
30. J. Palsberg and P. O'Keefe. A type system equivalent to flow analysis. *ACM Trans. Prog. Lang. Syst.*, 17(4):576–599, 1995.
31. J. Palsberg and P. Ørbaek. Trust in the λ-calculus. *Journal of Functional Programming*, 7(6):557–591, 1997.
32. G. D. Plotkin. A structural approach to operational semantics. Technical Report DAIMI FN-19, Aarhus University, 1981.
33. Z. Qian. A formal specification of a large subset of Java Virtual Machine instructions. In J. Alves-Foss, editor, *Formal Syntax and Semantics of Java*, Lecture Notes in Computer Science. Springer-Verlag, 1998. To appear.
34. J. C. Reynolds. User-defined types and procedural data structures as complementary approaches to data abstraction. In C. Gunter and J. Mitchell, editors, *Theoretical aspects of object-oriented programming*, pages 13–23. MIT Press, 1994.
35. F. Rouaix. A Web navigator with applets in Caml. In *Proceedings of the 5th International World Wide Web Conference, Computer Networks and Telecommunications Networking*, volume 28, pages 1365–1371. Elsevier, May 1996.
36. R. Stata and M. Abadi. A type system for Java bytecode subroutines. In *25th symposium Principles of Programming Languages*, pages 149–160. ACM Press, 1998.
37. D. Syme. Proving JavaS type soundness. Technical Report 427, University of Cambridge Computer Laboratory, June 1997.
38. J.-P. Talpin and P. Jouvelot. The type and effect discipline. *Information and Computation*, 111(2):245–296, 1994.

39. M. Tofte. Type inference for polymorphic references. *Information and Computation*, 89(1), 1990.
40. D. Volpano and G. Smith. A type-based approach to program security. In *Proceedings of TAPSOFT'97, Colloquium on Formal Approaches in Software Engineering*, volume 1214 of *Lecture Notes in Computer Science*, pages 607–621. Springer-Verlag, 1997.
41. D. Volpano, G. Smith, and C. Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3):1–21, 1996.
42. D. S. Wallach, D. Balfanz, D. Dean, and E. W. Felten. Extensible security architectures for Java. Technical report 546-97, Department of Computer Science, Princeton University, Apr. 1997.
43. D. S. Wallach and E. W. Felten. Understanding Java stack inspection. In *Proceedings of the 1998 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 1998.
44. F. Yellin. Low level security in Java. In *Proceedings of the Fourth International World Wide Web Conference*, pages 369–379. O'Reilly, 1995.

A Appendix: the containment lemma

In this appendix, we formalize the intuition that if the type of the applet environment does not contain certain `ref` types, then references of those types cannot be exchanged between the applet and the environment, and remain "contained" in one of them.

We annotate each source-language term a as coming either from the execution environment (a_{env}) or from the applet (a_{app}). We let m, n range over the two "worlds" *env* and *app*, and write $\overline{m}$ for the complement of m, i.e. $\overline{env} = app$ and $\overline{app} = env$.

Let T be a set of type expressions satisfying the following closure property: if $\tau \in T$, then all types τ' that contain τ as a sub-term also belong to T. (In section 5.1, we take T to be the set of all types containing an occurrence of t `ref`; in appendix B, T is the set of all types containing an occurrence of any `ref` type). We partition the set of locations into three countable sets:

- L_{app} is the set of locations with type τ `ref` $\in T$ that have been allocated by the applet;
- L_{env} is the set of locations with type τ `ref` $\in T$ that have been allocated by the environment (i.e. either initially present in the applet environment, or allocated by environment functions);
- L_{shared} is the set of locations whose type τ `ref` does not belong to T.

To ensure that locations allocated during evaluation are drawn from the correct set, we assume all source terms a^{τ}_{m} annotated with their static type τ and their world m and replace the evaluation rule for reference creation (rule 8) by the following rule:

$$\frac{\begin{array}{c}\varphi, e, s \vdash a_m^\tau \to v/s' \qquad \ell \notin \mathrm{Dom}(s') \cup \mathrm{Dom}(\varphi) \\ \ell \in L_{env} \text{ if } \tau\ \mathtt{ref} \in T \text{ and } m = \mathit{env} \\ \ell \in L_{app} \text{ if } \tau\ \mathtt{ref} \in T \text{ and } m = \mathit{app} \\ \ell \in L_{shared} \text{ if } \tau\ \mathtt{ref} \notin T\end{array}}{\varphi, e, s \vdash \mathtt{ref}(a)_m^{\tau\ \mathtt{ref}} \to \ell/s'\{\ell \leftarrow v\}} \quad (8')$$

It is easy to see that the modified rule produces the same evaluation derivations as the initial rule, up to a renaming of fresh locations: if $\varphi, e, s \vdash a \to r$ with the initial rule, then $\varphi, e, s \vdash a \to r'$ with the modified rule, where r' is identical to r up to a renaming of locations not in $\mathrm{Dom}(s) \cup \mathrm{Dom}(\varphi)$.

We say that a value v in a store s is contained in world m, and we write $C_m(v, s)$, if any source term operating on v in s can directly access only locations that are in L_m or L_{shared}, but not locations in $L_{\overline{m}}$. Formally, $C_m(v, s)$ is defined by case analysis on v, as follows:

- $C_m(b, s)$ is always true
- $C_m(\lambda x.a_n[e], s)$ if $C_n(e, s)$
- $C_m((v_1, v_2), s)$ if $C_m(v_1, s)$ and $C_m(v_2, s)$
- $C_m(\ell, s)$ if $l \notin L_{\overline{m}}$ and $C_m(s(\ell), s)$
- $C_m(e, s)$ if $C_m(e(x), s)$ for all $x \in \mathrm{Dom}(e)$.

Notice that for closures, it's the world n of the function body a_n that determines the containment of the closure environment, not the world m in which the closure value is being used. The reason is that the environment values can only be accessed by the function body, not arbitrary terms of world m. This is characteristic of procedural abstraction in the sense of section 3.2.

As usual, the definition of C above is not well-founded by induction on v. We view the equations above as fixpoint equations for an operator, which is increasing, and define C as the greatest fixpoint of that operator. (The smallest fixpoint is always false on value/store pairs that contain a cycle, which is not what we want; it's the greatest fixpoint that gives the expected behavior for C. See [39] for detailed explanations.)

Here are two important lemmas on the C predicate. First, assigning a value contained in world m to a location which is not in $L_{\overline{m}}$ preserves the containment of all other values.

Lemma 5. *Assume either $\ell \in L_m$ and $C_m(v, s)$, or $\ell \in L_{shared}$ and $C_{app}(v, s)$ and $C_{env}(v, s)$. Then, for all values w and worlds n, $C_n(w, s)$ implies $C_n(w, s\{\ell \leftarrow v\})$.*

Proof. The proof is a standard argument by coinduction, close to the proof of Lemma 4.6 in [39]. The only non-trival case is $w = \ell$ (the modified location). By hypothesis $C_n(w, s)$, we have $\ell \notin L_{\overline{n}}$ and $C_n(s(\ell), s)$. Write $s' = s\{\ell \leftarrow v\}$. Notice that $s'(\ell) = v$. If $n = m$, we have $C_n(v, s)$ by assumption, hence $C_n(v, s')$ by the coinduction hypothesis, from which it follows that $C_n(\ell, s')$. If $n = \overline{m}$, since $\ell \notin L_{\overline{n}}$, we have $\ell \notin L_m$ and thus it must be the case that $\ell \in L_{shared}$ and $C_{app}(v, s)$ and $C_{env}(v, s)$ to comply with the assumptions of the lemma. Thus, we have $C_n(v, s)$ and we conclude as in the previous case.

Second, values that belong to a type $\tau \notin T$ can be exchanged between the *app* and *env* worlds without breaking containment.

Lemma 6. *Assume* $S \models v : \tau$ *and* $\models s : S$*. If* $\tau \notin T$*, then* $C_m(v, s)$ *implies* $C_{\overline{m}}(v, s)$ *for all worlds* m*.*

Proof. By structural induction on τ. If τ is a base type or a function type, then v is a base value or a closure, and the containment of v is independent of the world m. If τ is a product type $\tau_1 \times \tau_2$, the closure condition over T guarantees that $\tau_1 \notin T$ and $\tau_2 \notin T$; the result follows from the induction hypothesis. Finally, if τ is a `ref` type, we have $v = \ell$ and $\tau = S(\ell)$ `ref`. Since $\tau \notin T$, it follows that ℓ belongs to L_{shared}, but neither to L_{app} nor L_{env}. Hence, $l \notin L_m$. Moreover, $S(\ell) \notin T$ by the closure condition on T, hence $C_{\overline{m}}(s(\ell), s)$ by application of the induction hypothesis. It follows that $C_{\overline{m}}(\ell, s)$.

We can now show that containment is preserved at each evaluation step:

Lemma 7 (Containment lemma). *Let* E_{api} *be the type of the execution environment for applets. Assume* $E_{api}(x) \notin T$ *for all* $x \in \mathrm{Dom}(E_{api})$*. Further assume* $E \vdash a_m : \tau$ *and* $S \models e : E$ *and* $\models s : S$ *and* $C_m(e, s)$*. If* $\varphi, e, s \vdash a_m \to v/s'$*, then* $C_m(v, s')$*, and for all values* w *and worlds* n *such that* $C_n(w, s)$*, we have* $C_n(w, s')$*.*

Proof. The proof is by induction on the evaluation derivation and case analysis on a. Notice that by Lemma 2, we have the additional result that there exists a store typing S' extending S such that $S' \models v : \tau$ and $\models s' : S'$. This makes the semantic typing hypotheses go through the induction. The interesting cases are assignment and function application; the other cases are straightforward.

Assignment: a is $(a_1^{\sigma\ \mathtt{ref}} := a_2^{\sigma})$. We apply the induction hypothesis twice, obtaining

$$\begin{gathered}
\varphi, e, s \vdash a_1 \to v_1/s_1 \qquad \varphi, e, s_1 \vdash a_2 \to v_2/s_2 \\
C_m(v_1, s_2) \qquad C_m(v_2, s_2) \\
C_n(w, s) \text{ implies } C_n(w, s_2) \text{ for all } n, w \\
S_2 \models v_1 : \sigma\ \mathtt{ref} \qquad S_2 \models v_2 : \sigma \qquad \models s_2 : S_2.
\end{gathered}$$

Hence, v_1 is a location ℓ, and since v_1 is contained in m, we have $\ell \notin L_{\overline{m}}$. Therefore, either $\ell \in L_m$ or $\ell \in L_{shared}$. But in the latter case, $\sigma \notin T$ by construction of L_{shared}, hence $C_m(v_2, s_2)$ implies $C_{\overline{m}}(v_2, s_2)$ as well by Lemma 6. In both cases, the hypotheses of Lemma 5 are met. Writing $s' = s_2\{\ell \leftarrow v_2\}$, we obtain the expected result: $C_n(w, s)$ implies $C_n(w, s')$ for all n, w. The other expected result, $C_m((), s')$, is trivial.

Application: a is $a_1^{\sigma \to \tau}(a_2^{\sigma})$. By applying the induction hypothesis twice, we obtain

$$\varphi, e, s \vdash a_1 \to v_1/s_1 \qquad \varphi, e, s_1 \vdash a_2 \to v_2/s_2$$
$$C_m(v_1, s_2) \qquad C_m(v_2, s_2)$$
$$C_n(w, s) \text{ implies } C_n(w, s_2) \text{ for all } n, w$$
$$S_2 \models v_1 : \sigma \to \tau \qquad S_2 \models v_2 : \sigma \qquad \models s_2 : S_2.$$

Hence, v_1 is a closure $\lambda x.a'_n[e']$, and the last evaluation rule used is rule 4.

If $n = m$ (intra-world call), we have $C_m(e'\{x \leftarrow v_2\}, s_2)$ as a consequence of $C_m(v_1, s_2)$ and $C_m(v_2, s_2)$, and the two conclusions follows easily from the induction hypothesis applied to the evaluation of a'_n.

If $n = \overline{m}$ (cross-world call), then the type $\sigma \to \tau$ of the function must occur as a sub-term of the typing environment E_{api}. Hence, $\sigma \notin T$ and $\tau \notin T$. Since the type of v_2 is not in T, by Lemma 6 it follows that $C_n(v_2, s_2)$. Hence, $C_n(e'\{x \leftarrow v_2\}, s_2)$, and we can apply the induction hypothesis to the evaluation of a'_n: $\varphi, e'\{x \leftarrow v_2\}, s_2 \vdash a'_n \to v/s'$. The resulting value v is contained in world n and has type τ; applying again Lemma 6, we get $C_m(v, s')$, which is the expected result.

B Appendix: modifiable locations in the case of systematic procedural encapsulation

In this appendix, we formalize the notion of modifiable locations, as introduced in section 3.2, in the particular case where the applet environment e_{api} is well-typed and its type E_{api} does not contain any `ref` types. This is not to say that e_{api} is purely functional: the functions it provide may very well have side-effects and use references internally. Only, all those internal references must be encapsulated inside functions and never handed to the applet directly. This systematic procedural encapsulation does not reduce expressiveness in any significant way, and at any rate is a reasonable thing to do given that our goal is to characterize semantically procedural encapsulation.

From the technical side, requiring that no `ref` types occur in the applet's interface E_{api} has an interesting consequence: only source terms from the browser can operate directly on locations allocated by the browser, and only terms from the applet can operate directly on locations allocated by the applet. Thus, all references are contained (in the sense of appendix A, taking T to be the set of all types where the `ref` constructor occurs) either in the browser or in the applet, but never go from one world to the other.

We then rely on the notion of containment to define the set $ML(v, s)$ of locations modifiable from value v in store s as the smallest fixpoint of the following equations:

$$ML(b, s) = \emptyset$$
$$ML((v_1, v_2), s) = ML(v_1, s) \cup ML(v_2, s)$$

$$ML(\lambda x.a[e], s) = \{\ell \mid \text{there exists } v_1, s_1, v_2, s_2 \text{ such that } C_{app}(v_1, s_1) \text{ and } s_1(\ell) = s(\ell) \text{ for all } \ell \in L_{env}, \text{ and either } \{\ell \leftarrow \emptyset\}, e\{x \leftarrow v_1\}, s_1 \vdash a \rightarrow \texttt{err} \text{ or } \emptyset, e\{x \leftarrow v_1\}, s \vdash a \rightarrow v_2/s_2 \text{ and } \ell \in ML(v_2, s_2) \cup ML(e_{api}, s_2)\}$$

$$ML(e, s) = \bigcup_{x \in \mathrm{Dom}(e)} ML(e(x), s)$$

The case for closures follows conditions 1 and 4 in the informal discussion from section 3.2. For condition 5, we use the condition $C_{app}(v_1, s_1)$ instead of the more natural $\ell \notin ML(v_1, s_1)$, so that the equations remain increasing in ML and the existence of the smallest fixpoint is guaranteed. The typing hypothesis (that E_{api} contains no `ref` types) renders condition 3 vacuous, and also dispenses us with defining ML over locations.

The following lemma show that modifiable locations are indeed the only locations modified during the application of a closure.

Lemma 8. *Let p be a set of locations, with $p \subseteq L_{env}$. Let $\lambda x.a_{env}[e]$ be a closure of an environment function. Assume $p \cap ML(\lambda x.a[e], s) = \emptyset$ and $C_{app}(v, s)$. If $Prot(p), e\{x \leftarrow v\}, s \vdash a \rightarrow r$, then $r \neq$ `err`. Instead, $r = v'/s'$, and moreover $p \cap ML(v', s') = \emptyset$ and $p \cap ML(e_{api}, s') = \emptyset$.*

Proof. Assume, by way of contradiction, that $r =$ `err`. Given the evaluation rules, there must exist $\ell \in p$ such that $\{\ell \leftarrow \emptyset\}, e\{x \leftarrow v\}, s \vdash a \rightarrow$ `err`. By definition of ML, this means that $\ell \in ML(\lambda x.a[e], s)$. This contradicts the hypothesis $p \cap ML(\lambda x.a[e], s) = \emptyset$. Hence, $r \neq$ `err`. We therefore have $Prot(p), e\{x \leftarrow v\}, s \vdash a \rightarrow v'/s'$ for some v' and s'. Given the evaluation rules, this implies that we can also derive $\emptyset, e\{x \leftarrow v\}, s \vdash a \rightarrow v'/s'$. Hence, by definition of ML, any ℓ belonging to $ML(v', s')$ or $ML(e_{api}, s')$ also belongs to $ML(\lambda x.a[e], s)$. It follows that $p \cap ML(v', s') = \emptyset$ and $p \cap ML(e_{api}, s') = \emptyset$.

Using the notion of modifiable locations instead of reachable locations, we finally obtain a security property similar to Property 1: the execution of an applet cannot write to locations that are not modifiable from the initial execution environment.

Security property 5 *Assume $S \models e_{api} : E_{api}$ and $\models s : S$. Further assume that E_{api} contains no `ref` types. If $p \subseteq \mathrm{Dom}(s)$ and $p \cap ML(e_{api}, s) = \emptyset$, then for all applets a, we have $Prot(p), e_{api}, s \vdash a \not\rightarrow$ `err`.*

The property follows from the inductive lemma below.

Lemma 9. *Assume $S \models e_{api} : E_{api}$ and $\models s : S$ and E_{api} contains no `ref` types. Further assume $p \subseteq L_{env}$ and $p \cap ML(e_{api}, s) = \emptyset$. Assume $E \vdash a_{app} : \tau$ and $S \models e : E$ and $C_{app}(e, s)$. If $Prot(p), e, s \vdash a_{app} \rightarrow r$, then $r \neq$ `err`. Instead, $r = v'/s'$ and we have $p \cap ML(e_{api}, s') = \emptyset$*

Proof. The proof is by induction on the evaluation derivation. We rely on Lemmas 2 and 7 to ensure that the semantic typing and containment hypotheses go through the induction. The two non-obvious cases are assignment and application of a closure of a browser function. We write $\varphi = \mathit{Prot}(p)$.

Assignment: a is $a_1 := a_2$. Applying the induction hypothesis twice, we obtain $\varphi, e, s \vdash a_1 \rightarrow v_1/s_1$ and $\varphi, e, s_1 \vdash a_2 \rightarrow v_2/s_2$, with $C_{app}(v_1, s_2)$ and $C_{app}(v_2, s_2)$ and $S_2 \models v_1 : \sigma\ \texttt{ref}$ and $S_2 \models v_2 : \sigma$ and $p \cap \mathit{ML}(e_{api}, s_2) = \emptyset$. Hence, v_1 is a location ℓ, and since v_1 is contained in *app*, we have $\ell \in L_{app}$. Thus, $\ell \notin p$ and the evaluation of the assignment does not result in `err`. Moreover, by definition of *ML*, we have $\mathit{ML}(e_{api}, s) = \mathit{ML}(e_{api}, s')$ if $s(\ell) = s'(l)$ for all $\ell \in L_{env}$. In the present case, the store s' at the end of the evaluation is $s_2\{\ell \leftarrow v_2\}$, with $\ell \notin L_{env}$ by containment, hence $\mathit{ML}(e_{api}, s') = \mathit{ML}(e_{api}, s_2)$ and thus $p \cap \mathit{ML}(e_{api}, s') = \emptyset$ as expected.

Application: a is $a_1(a_2)$. By induction hypothesis, we have $\varphi, e, s \vdash a_1 \rightarrow v_1/s_1$ and $\varphi, e, s_1 \vdash a_2 \rightarrow v_2/s_2$, with $C_{app}(v_1, s_2)$ and $C_{app}(v_2, s_2)$ and $S_2 \models v_1 : \sigma \rightarrow \tau$ and $S_2 \models v_2 : \sigma$ and $p \cap \mathit{ML}(e_{api}, s_2) = \emptyset$. Hence, v_1 is a closure $\lambda x.a'_m[e']$. If $m = \mathit{app}$ (the function comes from the applet), the result follows from the induction hypothesis applied to the evaluation of a'. If $m = \mathit{env}$ (the function comes from the applet environment), the closure $\lambda x.a'_m[e']$ can only be obtained by looking up a variable bound in e_{api}, then possibly performing some function applications or `fst` and `snd` operations. Thus, we have $\mathit{ML}(\lambda x.a'_m[e'], s_2) \subseteq \mathit{ML}(e_{api}, s_2)$, and the result follows by Lemma 8.