Flow Locks: Towards a Core Calculus for Dynamic Flow Policies

Niklas Broberg and David Sands

Chalmers University of Technology and Göteborg University

Abstract. Security is rarely a static notion. What is considered to be confidential or untrusted data varies over time according to changing events and states. The static verification of secure information flow has been a popular theme in recent programming language research, but information flow policies considered are based on multilevel security which presents a static view of security levels. In this paper we introduce a very simple mechanism for specifying dynamic information flow policies, flow locks, which specify conditions under which data may be read by a certain actor. The interface between the policy and the code is via instructions which open and close flow locks. We present a type and effect system for an ML-like language with references which permits the completely static verification of flow lock policies, and prove that the system satisfies a semantic security property generalising noninterference. We show that this simple mechanism can represent a number of recently proposed information flow paradigms for declassification.

1 Introduction

Unlike access control policies, enforcing an information flow policy at run time is difficult because information flow is not a runtime property; we cannot in general characterise when an information leak is about to take place by simply observing the actions of a running system. From this perspective, statically determining the information-flow properties of a program is an appealing approach to ensuring secure information flow. However, security *policies*, in practice, are rarely static: a piece of data might only be untrusted until its signature has been verified; an activation key might be secret only until it has been paid for.

This paper introduces a simple policy specification mechanism based on the idea that the reading of storage location ℓ by certain actors (principals, levels) is guarded by boolean flags, which we call *flow locks*. For example, the policy $\ell_{\{High; paid \Rightarrow Low\}}$ says that ℓ can always be read by an actor with a high clearance level, and also by an actor with a low clearance level providing the “paid” lock is open.

The interface between the flow lock policies and the security relevant parts of the program is provided by simple instructions for opening and closing locks. The program itself does not depend on the lock state, and the intention is that by statically verifying that the dynamic flow policy will not be violated, the lock state does not need to be computed at run time.¹

¹ The term *dynamic* flow policy could have different interpretations. We use it in the sense that the flow policies vary over time, but they are still statically known at compile time.

In addition to the introduction of flow locks, the main contributions of this paper are:

- The definition of a type system for an ML-like language with references which permits the completely static verification of flow lock policies.
- A formulation of the semantics of secure information flow for flow locks, and a proof that well typed programs are flow-lock secure (the reader is referred to the extended version of this article for the details).
- The demonstration that flow lock policies can represent a number of recently proposed information flow paradigms.

Regarding the last point, the work presented here can be viewed as a study of *declassification* mechanisms. In a recent study by Sabelfeld and Sands [18], declassification mechanisms are classified along four dimensions: *what* information is released, *who* releases information, *where* in the system information is released, and *when* information can be released. One of the key challenges stated in that work is to *combine* these dimensions. In fact, combination is perhaps not difficult; the real challenge is to combine these dimensions without simply amassing the combined complexities of the contributing approaches. Later in this paper we argue that flow locks can encode a number of recently proposed “declassification” paradigms, including the lexically scoped flow policies introduced by Almeida Matos and Boudol [2], Chong and Myers’ notion of *noninterference until declassification* [5], and Zdancewic and Myers *robust declassification* [22, 13]. These examples, represent the “where”, “when” and “who” dimensions of declassification, respectively, suggesting that flow locks have the potential to provide a core calculus of dynamic information flow policies.

The remainder of the paper is organised as follows. Section 2 gives an informal introduction to flow locks by showing a few motivating examples. In Section 3 we then present the system formally, and outline a semantic security condition in Section 4. Section 5 discusses related systems, with an emphasis on how we can use flow locks to encode them. Finally Section 6 concludes.

2 Motivating Examples

First let us assume we have a simple imperative language without any security control mechanisms of any kind. Borrowing an example from Chong and Myers [5], suppose we want to implement a system for online auctions with hidden bids in this language. We could write part of this system as the code on the right.

This surely works, but there is nothing in the language that prevents us from committing a serious security error. We could for instance accidentally switch the lines 2 and 3, resulting in *A*’s bid being made public before *B* places her bid, giving *B* the chance to tailor her bid after *A*’s.

```

1 int aBid = getABid();
2 int bBid = getBBid();
3 makePublic(aBid);
4 makePublic(bBid);
5 ...decide winner + sell item
```

Flow locks are a mechanism to ensure that these and other kinds of programming errors are caught and reported in a static check of the code.

The basic idea is very similar to what many other systems offer. To deny the flow of data to places where it was not meant to go, we annotate variables with policies that govern how the data held by those variables may be used. Looking back on our example, a proper policy annotation on the variable `aBid` could be $\{A; \text{BBid} \Rightarrow B\}$. The intuitive interpretation of this policy is that the data held by variable `aBid` may always be accessed by *A*, and may also be accessed by *B* whenever the condition `BBid`, that *B* has placed a bid, is fulfilled. `BBid` here is a *flow lock* — only if the lock is *open* can the data held by this variable flow to *B*. To know whether the lock is open or not we must look at how the functions for getting the bids could be implemented.

The function shown on the right first fetches the bid sent by *A*. We model the incoming channel as a global variable that can be read from, one with the same policy as `aBid`. When the bid has been read, the function signals this by opening the `ABid` lock—*A* has now placed a bid and the program can act accordingly. The implementation of `getBBid` follows the same pattern, and will result in `BBid` being open.

```
function getABid() {
  int {A;BBid ⇒ B} x
    = bidChanFromA;
  open ABid;
  return x;
}
```

```
function makePublic(bid) {
  publicChannel = bid;
}
```

Now both bids have been placed and can thus be released. The `makePublic` function would be implemented as shown on the left. The outgoing `publicChannel` is also modelled as a global

variable that can be written to. This one has the policy $\{A; B\}$ attached to it, denoting that both *A* and *B* will be able to access any data written into it. At the points in the program where `makePublic` is applied, both *A* and *B* will have placed their bids, the locks `ABid` and `BBid` will both be open, and the flows to the public channel will both be allowed. However, if the lines 2 and 3 were now accidentally switched, it would be a different story. Then we would attempt to release *A*'s bid, guarded by the policy $\{A; \text{BBid} \Rightarrow B\}$, onto the public channel with policy $\{A; B\}$. Since the flow lock `BBid` will then not yet be opened, this flow is illegal and the program can be rejected.

Taking the example one step further, assume that we have two items up for auction, one after the other. We can implement this rather naively as the program to the right. The locks `ABid` and `BBid` will both be opened on the first calls to the `getXBid` functions. But unless we have some means to reset them, there is again nothing to stop us from accidentally switching lines to make our program insecure, this time lines 9 and 10. The same problem could also be seen from a different angle: what if the locks were already open when we got to this part of the program? Clearly we need a closing mechanism to go with the open.

The function `auctionItem` could then be implemented as shown here. By closing the locks when an auction is initiated, we can rest assured that both *A* and *B* must place new bids for the new item before either bid is made public.

```
1 auctionItem(firstItem);
2 aBid = getABid();
3 bBid = getBBid();
4 makePublic(aBid);
5 makePublic(bBid);
6 ... decide winner + sell item
7 auctionItem(secondItem);
8 aBid = getABid();
9 bBid = getBBid();
10 makePublic(aBid);
11 makePublic(bBid);
12 ... decide winner + sell item
```

```
function auctionItem(item){
  close ABid, BBid;
  ... present item ... }
```

It should be fairly easy to see that what we have here is a kind of state machine. The state at any program point is the set of locks that are open at that point, and the open and close state-

ments form the state transitions. A clause $\sigma \Rightarrow A$ in a policy means that A may access any data guarded by that policy in any state where σ is open.

Our lock-based policies also give us an easy way to separate truly secret data from data that is currently secret, but that may be released to other actors under certain circumstances. Assume for instance that payment for auctioned items is done by credit card, and that the server stores credit card numbers in memory locations `aCCNum` and `bCCNum` respectively. Assume further that the line `aBId := aCCNum;` is inserted, either by sheer mistake or through malicious injection, just before where `aBId` is made public. This would release A 's credit card number to B , however, the natural policy on `aCCNum` would be $\{A\}$, meaning only A may view this data, ever. Thus when we attempt the assignment above, it will be statically rejected since the policy on `aBId` is too permissive.

All the above are examples of policies to track confidentiality. The dual of confidentiality is integrity, i.e. deciding to what extent data can be trusted, and it should come as no surprise that flow locks can handle both kinds.

Returning to the example with the credit card, we assume that when A gives her credit card number, it must be validated (in some unspecified way) before we can trust it. To this end we introduce a “pseudo” actor T (for “trusted”) who should only be allowed to read data that is fully trusted. We then use an intermediate location `tmpACCNum` to hold the credit card number when it is submitted by A . This location is given the policy $\{A; \text{ACCVal} \Rightarrow T\}$, stating that this data is trusted only if the lock `ACCVal` is open, which is done when the submitted number has been validated. Once validated we can transfer the value to `aCCNum`, which now has the policy $\{A; T\}$ stating that this data is trusted.²

3 A Secure Type and Effect System

In the previous section we used a simple imperative language to give an easy introduction to the concept of flow locks. In this section we define the type system for flow locks in the more general context of an ML-like language with recursion and references (but without polymorphism).

3.1 The Language λ_{FL}

The terms and types of our language, dubbed λ_{FL} , are listed in Figure 1.

The policy language is worth some extra attention. The flow lock policies with which we work assumes a set of *actors* (or *levels*, *principals*) ranged over by A, B , and a set of flow locks ranged over by σ , with Σ for sets of locks. Both actors and flow locks are global in a program. A *policy* is a set of *clauses*, where each clause of the form

² In order to prevent overwriting this data with a new number that hasn't been validated, we should also be sure to close the lock `ACCVal` once the assignment is done.

Policies:	$p ::= \{c_1; \dots; c_n\} \quad c ::= \{\sigma_1, \dots, \sigma_k\} \Rightarrow A$
Values and types:	$v ::= n \mid b \mid () \mid \lambda x.M \mid \ell_{p,\tau}$ $\tau ::= int \mid bool \mid unit \mid (\tau, p) \xrightarrow{\Sigma, p, p, \Sigma} \tau \mid ref_p \tau$
Terms:	$M ::= v \mid x \mid MM \mid \text{if } M \text{ then } M \text{ else } M \mid \text{rec } x.M$ $\mid ref_{p,\tau} M \mid !M \mid M := M \mid \text{open } \sigma \mid \text{close } \sigma$
Derived forms:	$\text{let } x = M_1 \text{ in } M_2 \equiv (\lambda x.M_2)M_1 \quad M_1; M_2 \equiv (\lambda_.M_2)M_1$

Fig. 1. The λ_{FL} language

$\Sigma \Rightarrow A$ states the circumstances (Σ) under which A may view the data governed by this policy. Σ is a set of locks which we name the *guard* of the clause, and interpret it as a conjunction. Thus for the guard to be fulfilled, all the locks in Σ must be open. We can however have more than one clause for the same A , in which case the separate clauses also form a conjunction — A may read the data if either of the guards are fulfilled. In the special case where the guard contains no locks, signifying that the corresponding actor A may always view the data, we write the clause as only A instead of $\{\} \Rightarrow A$. From a logical perspective a policy is just a conjunction of definite Horn clauses, i.e. $\bigwedge_i \{\sigma_{i1} \wedge \dots \wedge \sigma_{in} \Rightarrow A_i\}$. We implicitly identify policies up to logical equivalence.³

Now we can continue with the language itself. Apart from the terms from standard λ calculus with recursion, λ_{FL} has constructs for creating (ref), dereferencing (!) and assigning to ($:=$) memory locations ($\ell_{p,\tau}$) through references. In addition to the core terms, we can also derive a few useful language constructs as is also shown in Figure 1.

The reference creation construct takes an extra parameter p which is the policy that the contents should be governed by. The same parameter also shows up on the memory locations themselves, together with the base type τ of the contents. In many cases this τ is irrelevant, or clear from the context, and in those cases we omit it and just write ℓ_p . Function types are annotated with read and write policies, and start and end states, and arguments are annotated with a reading policy. We discuss the meaning of these when we define the type system. There are also the open and close terms for manipulation flow locks, thereby changing the state of the program.

The semantics of the language is standard, but apart from the term M and a memory μ , the configurations include the current state Σ . This state is the set of currently open locks, which are effected by the execution of **open** and **close** expressions. The small-step semantics of these are simply:

$$\langle \Sigma, \text{open } \sigma, \mu \rangle \rightarrow \langle \Sigma \cup \{\sigma\}, (), \mu \rangle \quad \langle \Sigma, \text{close } \sigma, \mu \rangle \rightarrow \langle \Sigma \setminus \{\sigma\}, (), \mu \rangle$$

It is important to note that the only interaction between a program and the lock state is via the open and close instructions. This is because we are aiming for a completely static verification — we include the lock state in the semantics only to be able to prove

³ It is worth noting that we do not allow negative flow policies. Our policy language is monotonic, i.e. the more locks that are open, the more flows are allowed.

properties about flows, but the state is not actually represented at runtime. For this reason we also do not need to consider potential covert channels introduced by the flow lock state.

3.2 Some Intuitions About Flow-Lock Security

Before we define our type system, it is useful to get some intuitions about which programs we deem secure/insecure. At this point we only concern ourselves with information leaks arising from direct or indirect data flows. In particular we will not consider timing or termination sensitivity.

A few small example programs are presented on the right. All of these contain insecure direct data flows, except (3). In (1) the contents of $m_{\{B\}}$ may only be read by B, but we are attempting to leak them into a location readable by A. Same thing goes for (2) — even though B can read the contents of the target location, we are still leaking the contents of $m_{\{B\}}$ to A. The simple pattern is that we may not write data to a memory location if that location may be read by someone who cannot already access the data. What's more, this should hold for future time as well. Thus if a reader could access the data from the location we are writing to in some future state, that reader must also have access to the data that is being written, in that same state. Thus the example $m_{\{\sigma \Rightarrow A\}} := !\ell_{\{\sigma \Rightarrow A\}}$ is secure while program (4) is not. In program (5) we attempt to take data not yet readable by A, and put it in a location where A could read it right away. This should clearly not be allowed for the same reasons as for (4).

- (1) $\ell_{\{A\}} := !m_{\{B\}}$
- (2) $\ell_{\{A;B\}} := !m_{\{B\}}$
- (3) $\ell_{\{A\}} := !m_{\{A;B\}}$
- (4) $\ell_{\{\sigma \Rightarrow A;B\}} := !m_{\{B\}}$
- (5) $\ell_{\{A\}} := !m_{\{\sigma \Rightarrow A\}}$

The lock state in effect at the point of the assignment determines its validity, so the programs (6) and (7) are secure. However, we also want a program like (8) below to be considered

- (6) **open** σ ; $\ell_{\{A\}} := !m_{\{\sigma \Rightarrow A\}}$
- (7) $\ell_{\{A\}} := (\mathbf{open} \ \sigma; !m_{\{\sigma \Rightarrow A\}})$

secure, so we should take the policy of data read from some memory location to be the policy on the location, but taking into account the current state.

- (8) $\ell_{\{A\}} := \mathbf{let} \ x = (\mathbf{open} \ \sigma; !m_{\{\sigma \Rightarrow A\}}) \ \mathbf{in} \ (\mathbf{close} \ \sigma; x)$

In program (8) above, the data read from the reference will thus have the policy $\{A\}$ and not $\{\sigma \Rightarrow A\}$, since it is read in a state where σ is open.

Putting all this slightly more formally, data may be written to a memory location if and only if the policy on the location is at least as restrictive as the one on the data, with respect to the state in effect at the point of the assignment. We give a formal definition of this in the next section.

We must also handle indirect flows that arise from various branching situations. A very simple example program containing an invalid indirect flow is

- (9) **if** $!\ell_{\{A\}}$ **then** $m_{\{B\}} := \mathbf{true}$ **else** $m_{\{B\}} := \mathbf{false}$

This program is obviously insecure since it will leak the value of $\ell_{\{A\}}$ into $m_{\{B\}}$, but for some programs it is not so easy to tell. Consider the three programs

- (10) **if** $!\ell_{\{\sigma \Rightarrow A\}}$ **then** (**open** σ ; $m_{\{A\}} := \text{true}$) **else** (**open** σ ; $m_{\{A\}} := \text{false}$)
 (11) **if** $!\ell_{\{\sigma \Rightarrow A\}}$ **then** (**open** σ ; $m_{\{A\}} := \text{true}$; **close** σ) **else** ()
 (12) **if** (**open** σ ; $!\ell_{\{\sigma \Rightarrow A\}}$) **then** (**close** σ ; $m_{\{A\}} := \text{true}$) **else** ()

Program (10) could be argued correct since at the points where we leak the information to A , i.e. the assignments, the state allows A to access the result of the branching conditional directly, and hence the leak is secure.

However, as program (11) shows it is not that simple. If the second branch in (11) is chosen, the value of the condition is still leaked to A by the absence of a write, but at no point does the state allow the flow. The leaks come from knowing which of the two branches is taken, which suggests that the leak actually occurs at the branch point. Thus it is the policy of the condition, taken in the state in effect at the branch point, that decides what writes the branches may perform. This means that (9), (10) and (11) are all insecure, while (12) is secure even though the lock is closed again before the write.

Another possible source of indirect leaks is function application. If the function itself is secret, an attacker could still get information about what that function is by observing its effects, just like he could know which branch was taken by observing the effects of a conditional expression. Thus in a sense we can view function application as a kind of branching.

Consider the programs (13) – (19). In the program (13) we must ensure that the function read from the reference does not write to locations visible by anyone other than A , otherwise we could leak information about which function that was used. As an example, if the function read from $\ell_{\{A\}}$ in (13) is $(\lambda x. m_{\{B\}} := 1)$ or $(\lambda x. m_{\{B\}} := 2)$, B can determine which of the two that was used by reading $m_{\{B\}}$. We treat the application point in the same way as the branch point of a conditional, so in program (14) the body of the function must not write to a location directly visible to A , even if it first opens σ . However, since we have a call-by-value semantics, in program (15) the function body may perform writes to locations directly visible to A , even if it first closes σ , since σ will be open at the application point.

- (13) $(!\ell_{\{A\}}) ()$
 (14) $(!\ell_{\{\sigma \Rightarrow A\}}) ()$
 (15) $(!\ell_{\{\sigma \Rightarrow A\}}) (\text{open } \sigma; ())$
 (16) $(!\ell_{\{A\}}) := 0$
 (17) $(!\ell_{\{\sigma \Rightarrow A\}}) := (\text{open } \sigma; 0)$
 (18) $(\lambda x. \ell_{\{B\}} := x) (!\ell_{\{A\}})$
 (19) $(\lambda x. \ell_{\{B\}} := 0) (!\ell_{\{A\}})$

A similar situation is assignment to a reference that in turn has been read from a reference, as illustrated in program (16) which should be disallowed if the reference read from $\ell_{\{A\}}$ is visible to anyone other than A . In particular, the contents of $\ell_{\{A\}}$ could be $m_{\{B\}}$ or $n_{\{B\}}$, in which case B can determine the contents of $\ell_{\{A\}}$ by checking which of the two latter locations that contain the value 0. However, just as for application, program (17) is secure if the reference assigned to has policy $\{A\}$, or any policy that is more restrictive than $\{A\}$, since σ is opened before the assignment takes place.

We also need to look at how functions handle the values passed to them as arguments. Clearly we want to rule out a direct leak in the function body, as the one in example (18). One solution attempt could be to rule out all functions that write to “low” memory, i.e. locations with less restrictive policies than the one placed on the argument. But this also rules out perfectly secure programs such as (19) which in particular would

mean that we could not derive a sequential composition form as in figure 1 without placing too heavy restrictions on the writing capabilities of the second sub-program. Thus we want our type system to treat these two programs differently — (18) should be deemed insecure, but not (19).

Other issues such as whether our system is termination sensitive or timing sensitive (see [16] for an overview of these concepts) are orthogonal to the above discussion. We choose to develop a type system and semantics for termination and timing insensitive security. Termination insensitivity makes the type system simpler but the semantics more complex.

3.3 The Type System

Now we have all the intuition needed to construct the type system. We choose to model our system as a type and effect system in the style of Almeida Matos and Boudol [2]. This means in particular that all expressions will be given a *reading effect* and a *writing effect*. In our system the reading effect of an expression is a policy which states who may read the result of that expression, and in what lock states they may do so. The writing effect is also a policy, which records which actors and in what lock states they can see the memory effect of the expression's execution. Type judgments then have the form

$$\Gamma; \Sigma \vdash M : \tau, (r, w) \Rightarrow \Sigma'$$

- Γ is a typing environment for variables giving a type and policy for each variable.
- Σ is the state, i.e. the set of locks currently open.
- τ is the type of the term.
- (r, w) are the reading and writing effects of the term, both on the form of policies
- Σ' is the state the program will be in after evaluating the term

First we need to define a few operators on policies that we will use in the typing rules. The aforementioned ordering of how restrictive policies are is defined as

$$p_1 \preceq p_2 \equiv \forall (\Sigma_2 \Rightarrow A) \in p_2. \exists (\Sigma_1 \Rightarrow A) \in p_1. \Sigma_1 \subseteq \Sigma_2$$

Read out, we say that p_1 is less restrictive than p_2 if and only if every clause in p_2 is matched by a clause in p_1 for the same A with a less restrictive guard (one with no additional locks). From the logical perspective, this ordering corresponds directly to implication. The most restrictive policy is $\{\}$, also written $\top$, and data with this policy can never be accessed by anyone. On the other end of the spectrum is $\perp$, defined as the set of all actors in the system. In other words, data marked with $\perp$ can be read by everyone at all times.

To join two policies means combining their respective clauses, thereby forming the logical disjunction. We define

$$p_1 \sqcup p_2 \equiv \{ \Sigma_1 \cup \Sigma_2 \Rightarrow A \mid \Sigma_1 \Rightarrow A \in p_1, \Sigma_2 \Rightarrow A \in p_2 \}$$

It should be intuitively clear that the join of two policies is at least as restrictive as each of the two operands, i.e. $p \preceq p \sqcup p'$ for all p, p' . In contrast, forming the union of two policies, i.e. the meet, corresponding to $\sqcap$ or logical conjunction, makes the result less

$$\begin{array}{c}
\frac{}{\Gamma; \Sigma \vdash n : \text{int}, (\perp, \top) \Rightarrow \Sigma} \quad \frac{}{\Gamma; \Sigma \vdash b : \text{bool}, (\perp, \top) \Rightarrow \Sigma} \\
\\
\frac{}{\Gamma; \Sigma \vdash u_{p,\tau} : \text{ref}_p \tau, (\perp, \top) \Rightarrow \Sigma} \quad \frac{}{\Gamma; \Sigma \vdash () : \text{unit}, (\perp, \top) \Rightarrow \Sigma} \\
\\
\frac{\Gamma, x : (\tau, r_\alpha); \Delta \vdash M : \tau', (r, w) \Rightarrow \Delta'}{\Gamma; \Sigma \vdash \lambda x. M : (\tau, r_\alpha) \xrightarrow{\Delta, r, w, \Delta'} \tau', (\perp, \top) \Rightarrow \Sigma} \quad \frac{x : (\tau, r) \in \Gamma}{\Gamma; \Sigma \vdash x : \tau, (r(\Sigma), \top) \Rightarrow \Sigma} \\
\\
\frac{}{\Gamma; \Sigma \vdash \mathbf{open} \sigma : \text{unit}, (\perp, \top) \Rightarrow \Sigma \cup \{\sigma\}} \quad \frac{}{\Gamma; \Sigma \vdash \mathbf{close} \sigma : \text{unit}, (\perp, \top) \Rightarrow \Sigma \setminus \{\sigma\}} \\
\\
\frac{\Gamma, x : (\tau, r); \Sigma \vdash M : \tau, (r, w) \Rightarrow \Sigma}{\Gamma; \Sigma \vdash \mathbf{rec} x. M : \tau, (r, w) \Rightarrow \Sigma} \\
\\
\frac{\Gamma; \Sigma \vdash M : \tau, (r, w) \Rightarrow \Sigma'}{\Gamma; \Sigma \vdash \text{ref}_p M : \text{ref}_p \tau, (r, w) \Rightarrow \Sigma'} \quad \frac{\Gamma; \Sigma \vdash M : \text{ref}_p \tau, (r, w) \Rightarrow \Sigma'}{\Gamma; \Sigma \vdash !M : \tau, (r \sqcup p(\Sigma'), w) \Rightarrow \Sigma'} \\
\\
\frac{\Gamma; \Sigma \vdash M_1 : \text{ref}_p \tau, (r_1, w_1) \Rightarrow \Sigma' \quad \Gamma; \Sigma' \vdash M_2 : \tau, (r_2, w_2) \Rightarrow \Sigma''}{\Gamma; \Sigma \vdash M_1 := M_2 : \text{unit}, (\perp, w_1 \sqcap w_2 \sqcap p) \Rightarrow \Sigma''} \quad r_1(\Sigma'') \sqcup r_2(\Sigma'') \preceq p \\
\\
\frac{\Gamma; \Sigma \vdash M_0 : \text{bool}, (r_0, w_0) \Rightarrow \Sigma' \quad \Gamma; \Sigma' \vdash M_i : \tau, (r_i, w_i) \Rightarrow \Sigma_i \quad r_0(\Sigma') \preceq w_1 \sqcap w_2}{\Gamma; \Sigma \vdash \mathbf{if} M_0 \mathbf{then} M_1 \mathbf{else} M_2 : \tau, (r_0 \sqcup r_1 \sqcup r_2, w_0 \sqcap w_1 \sqcap w_2) \Rightarrow \Sigma_1 \sqcap \Sigma_2} \\
\\
\frac{\Gamma; \Sigma \vdash M_1 : (\tau, r_2) \xrightarrow{\Sigma_2, r_f, w_f, \Sigma_3} \tau', (r_1, w_1) \Rightarrow \Sigma_1 \quad \Gamma; \Sigma_1 \vdash M_2 : \tau, (r_2, w_2) \Rightarrow \Sigma_2}{\Gamma; \Sigma \vdash M_1 M_2 : \tau', (r_1 \sqcup r_f, w_1 \sqcap w_2 \sqcap w_f) \Rightarrow \Sigma_3} \quad r_1(\Sigma_2) \preceq w_f
\end{array}$$

Fig. 2. Type and Effect system

restrictive, so we have $p \sqcap p' \preceq p$ for all p, p' . Both $\sqcap$ and $\sqcup$ are clearly commutative and associative.

Finally we need to define using a policy with respect to a particular state, or normalising to a state. We say that policy p normalised at state Σ is

$$p(\Sigma) \equiv \{\Sigma' \setminus \Sigma \Rightarrow A \mid \Sigma' \Rightarrow A \in p\}$$

Informally, we remove all open locks from all guards in p , since these no longer restrict data governed by p . This function is antimonotonic, so $\Sigma \subseteq \Sigma' \implies p(\Sigma') \preceq p(\Sigma)$, and in particular $p(\Sigma) \preceq p$ for all Σ . Logically this operation is a partial evaluation, where all variables (locks) that appear in Σ are set to *true* in p .

The type and effect system is presented in Figure 2. The rules for literal values are straight-forward, giving all such values the reading effect bottom. However, from the variable rule we see that variables are given a reading policy. This is used to keep track of the reading policies of function arguments, as can be seen from the rules for abstraction and application, and the purpose is to disallow programs like (18) while still allowing (19). It is important to note that we do *not* check that $r_2(\Sigma_2) \preceq w_f$ in the application rule, since doing so would invalidate program (19). Instead we rely on the

type checking of the body of the function to find any leaks inside it, with the help of the annotation on its parameter.

In the rule for abstractions, we annotate the function arrow with the latent read and write effects that will be accurate for the function body once it is applied. We also annotate the arrow with the state that the program will be in at the application point, and the state the program will be in after evaluating the body. The interpretation of a function with type $(\tau, r_\alpha) \xrightarrow{\Delta, r, w, \Delta'} \tau'$ is thus that when applied in state Δ on an argument of type τ and with reading policy r_α , it will produce a result of type τ' with reading policy r . The writing policy w states who could see that the function has been applied, and the whole program will be in state Δ' afterwards. This is all mirrored by the appropriate states in the application rule.

Direct leaks, like the ones in programs (1), (2), (4) and (5), are handled by the check $r_2(\Sigma'') \preceq p$ in the rule for assignment. Since we normalise the policy r_2 of the assignee to the state in effect at the point of the assignment, program (5) would be secure if run in a state where σ is open, which is exactly what happens in programs (6) and (7). Also the normalisation to the current state in the dereferencing rule, i.e. $p(\Sigma')$ in the reading effect of the conclusion, means that program (8) will be deemed secure. The same kind of normalisation also appears in the variable rule.

The check $r_0(\Sigma') \preceq w_1 \sqcap w_2$ in the conditional rule will ensure that an indirect leak like the one in (9) will not be allowed. The normalisation of r_0 to Σ' means that it is the state at the branch point that is important, which disallows (10) and (11) but lets (12) through. The branches may open and close different locks, so the end states can differ. Since policies are monotonic, we can use the intersection of the end states as a safe approximation for the following program.

The checks $r_1(\Sigma'') \preceq p$ in the assignment rule, and the corresponding $r_1(\Sigma_2) \preceq w_f$ in the application rule handle indirect flows like in (13), (14) and (16), but allow (15) and (17).

In the assignment rule, the reading effect in the conclusion is $\perp$. The reason is that the result of an assignment is always $()$, independent of the result values of the two expressions M_1 and M_2 , so no information is leaked by making the $()$ result public. For similar reasons, r_2 does not show up in the reading effect in the conclusion of the application rule. Since function arguments are annotated with their reading effects, if the result of M_2 has any effect on the result of the whole application expression, this fact will be seen through r_f .⁴

4 Semantic Security Properties

For reasons of space this section gives only a brief outline of the semantic definitions and results about flow-lock security. For details the reader is referred to the full version of the paper. The main development is the definition of a notion of *flow lock security* which

⁴ The rules involving functions are fairly restrictive as they are formulated here. One could easily imagine various forms of subsumption, both for lock states and argument policies, that would make the system less restrictive. However, adding subsumption would complicate the overall formulation of the type system, so we leave it for the full version of the paper.

- generalises a standard notion of noninterference, since amongst other things it guarantees that a noninterference property holds for computation between any changes in the flow lock state.
- holds whenever a flow lock program is well-typed – i.e. well-typed programs are flow-lock secure.

The two main challenges in generalising the notion of noninterference to the flow lock setting are (i) dealing with policy change, in particular when a policy become less liberal (i.e. when locks are closed), and (ii) coping with the “latency” of a language with higher-order functions and state.

Before we can deal with policy change we must understand the underlying notion of noninterference that we build upon in the definition of flow lock security. Suppose in a computation that the set of locks Σ are open. This means that an actor A is permitted, at that point, to read the contents of a location with policy c providing $(\Delta \Rightarrow A) \in c$ for some $\Delta \subseteq \Sigma$. If the set of open locks Σ does not change, then the basic noninterference property that we expect is the following: given two memories which agree on all A -readable locations, the results of computing with these two respective memories should also agree on A -readable locations. This means that the actor A does not learn anything about the memory locations that were not visible initially.

Now to deal with change in the set of open locks we follow the “self-bisimulation” approach from [17], whereby security is characterised by a more general property of two programs being bisimilar with respect to the observable parts of memory. One particular feature of the definition from [17] is that the bisimulation is defined over programs and not configurations (program-memory pairs). The idea is that at each step of the bisimulation the pair of programs under comparison are inspected in all pairs of memory states which are indistinguishable to the attacker. This very strong requirement was needed to make the definition of security compositional with respect to parallel composition. But this approach of “resetting” the store at each step has another very useful property: it enables us to reset the state in the event of a policy change. For example, one particular difficulty is that when the current policy becomes *more* restrictive — in our case when locks are closed — then we need a way to reestablish a stronger security requirement at that point in the execution. It is notable that two previous semantic accounts of temporary policy weakening mechanisms, Mantel and Sands’s language based intransitive noninterference condition [8], and Almeida Matos and Boudol’s *nondisclosure* policy [2], both rely on such a “resetting” bisimulation not only to deal with threads, but more importantly to provide a semantics to local policy change mechanisms. Of these two earlier definitions, our definition is close in spirit to Almeida Matos and Boudol’s — although we refer to the full paper for details.

A straightforward “resetting” bisimulation is not enough to define flow lock security; it is not enough to consider just the locations which are currently visible to an actor A . Consider a program such as $\ell_{\{\sigma \Rightarrow A\}} := !m_{\{\sigma' \Rightarrow A\}}$ in a state where neither σ nor σ' are open. Since this assignment deals with locations not currently visible to A then a simple resetting bisimulation would allow it. However it is clearly insecure with respect to possible *future* states which may open σ . In order to detect the insecure flow that might be revealed at some future time we must check the equivalence of the two memories in a state where σ is open but σ' is not. More generally, the definition

of flow-lock security therefore takes into account all possible (more permissive) future lock states.

For the full details of these developments the reader is referred to the extended version of this paper.

5 Relating to Other Systems and Idioms

Standard Noninterference. As a first example of the expressiveness of our system, consider a standard termination insensitive noninterference property for a lattice-based security model in the standard Denning style [6].

In this setting we have a lattice of security levels $\langle \mathcal{L}, \sqsubseteq, \sqcup \rangle$, and a policy level: $\text{Loc} \rightarrow \mathcal{L}$ that fixes the intended security level of the storage locations in the program. Given such a policy we can define noninterference. For simplicity we consider closed programs of unit type which do not perform any storage allocation or lock open/close operations. In what follows let metavariables P and Q range over such programs.

Definition 1 (Noninterference). *Given two stores μ and ν , and a level $k \in \mathcal{L}$, define μ and ν to be indistinguishable at level k , written $\mu =_k \nu$, iff for all ℓ such that $\text{level}(\ell) \sqsubseteq k$ we have $\mu(\ell) = \nu(\ell)$.*

Then we say that P is noninterfering if for all k , whenever $\langle P, \mu \rangle \rightarrow^ \langle (), \mu' \rangle$ and $\langle P, \nu \rangle \rightarrow^* \langle (), \nu' \rangle$, then $\mu =_k \nu$ implies $\mu' =_k \nu'$.*

To represent a lattice policy we do not need any locks; we represent the reading level of a variable by the set of levels at which it may be read. Thus the policy for a storage location ℓ is the upwards closure of its lattice level, written $\uparrow \text{level}(\ell)$, where $\uparrow j = \{\{\} \Rightarrow k \mid k \sqsupseteq j\}$. Given this, we have the following:

Theorem 1. *If P is flow lock secure then P is noninterfering.*

Thus whenever we show that P is secure in the flow lock setting then it is also noninterfering. But it is perhaps not too surprising that our security specification is stronger than standard noninterference. A reasonable concern might be that the definition, or indeed the type system, is too strong to be useful. Here we show that despite being stronger, we are still able to type just as much as “typical” systems for regular noninterference.

Figure 3 presents a simple type system for a while language which can be seen as a straightforward reformulation of the typing system presented by Volpano, Irvine and Smith [21].

$$\begin{array}{c}
 \frac{p = \bigsqcup_{\ell \in E} \text{level}(\ell).}{\vdash_{NI} E : p} \quad \frac{\vdash_{NI} E : q \quad p \sqcup q \sqsubseteq \text{level}(\ell)}{p \vdash_{NI} u := E} \quad \frac{p \vdash_{NI} C_1 \quad p \vdash_{NI} C_2}{p \vdash_{NI} C_1; C_2} \\
 \frac{\vdash_{NI} E : q \quad p \sqcup q \vdash_{NI} C_i \quad i = 1, 2}{p \vdash_{NI} \text{if } E \text{ then } C_1 \text{ else } C_2} \quad \frac{\vdash_{NI} E : q \quad p \sqcup q \vdash_{NI} C}{p \vdash_{NI} \text{while } (E) C}
 \end{array}$$

Fig. 3. Standard Noninterference Type System

Define the following translation $\lceil \cdot \rceil$ from terms in the while language to λ_{FL} :

$$\begin{aligned}
\lceil \text{while } (E) C \rceil &= \text{rec } x. \text{if } \lceil E \rceil \text{ then } \lceil C \rceil; x \text{ else } () \\
\lceil \text{if } E \text{ then } C_1 \text{ else } C_2 \rceil &= \text{if } \lceil E \rceil \text{ then } \lceil C_1 \rceil \text{ else } \lceil C_2 \rceil \\
\lceil C_1; C_2 \rceil &= \lceil C_1 \rceil; \lceil C_2 \rceil \\
\lceil \ell := E \rceil &= \ell_p := \lceil E \rceil \quad \text{where } p = \uparrow \text{level}(\ell) \\
\lceil E \rceil &= E' \quad \text{where } E' \text{ is the result of replacing} \\
&\quad \text{each location } \ell \text{ in } E \text{ with } \ell_{\uparrow \text{level}(\ell)}.
\end{aligned}$$

To make our formulations easier, let us restrict the language of expressions to booleans (so we do not have to consider typing issues). Now we can state that whenever something is typeable in the simple noninterference system, a corresponding derivation holds for the flow locks system:

Theorem 2. *Let Γ_0 be the type environment that maps every storage location to *bool*. Then*

1. *If $\vdash_{NI} E : k$ then $\Gamma_0; \emptyset \vdash \lceil E \rceil : \text{bool}, (r, \top) \Rightarrow \emptyset$ where $r = \uparrow k$*
2. *If $pc \vdash_{NI} C$ then $\Gamma_0; \emptyset \vdash \lceil C \rceil : \text{unit}, (r, w) \Rightarrow \emptyset$ where $w \subseteq \uparrow pc$*

We also expect that a similar theorem holds for some suitable termination-insensitive version of DCC [1], although we have not attempted to show this formally.

Simple Declassification. We can encode a simple declassification mechanism in the same Denning-style setting as used in the previous example. The needed extra step is to extend all policies with clauses to allow declassification. For each level j not in the policy already, we introduce a flow lock σ_j representing a declassification to that level. The new policies then look like

$$\{k \mid k \sqsupseteq \text{level}(\ell)\} \cup \{\sigma_j \Rightarrow k \mid j \not\sqsupseteq \text{level}(\ell), k \sqsupseteq j\}$$

We can now define a declassification operator to level j as

$$\text{declassify}_j \equiv (\lambda v. \text{let } x = (\text{open } \sigma_j; v) \text{ in } (\text{close } \sigma_j; x))$$

It is easy to verify from the type system that the only effect of applying this function to some value is that the value will then be readable also at level j , as was our intention.

Lexically Scoped Flows. In the setting of a multilevel security model, Almeida Matos and Boudol describe how to build a system with lexically scoped dynamic flow policies [2]. They start from a λ -calculus with recursion and references like we do, and introduce a construct “*flow $\alpha \prec \beta$ in M* ” that extends the current global flow policy to also allow flows from level α to β in the scope of M . These flows are transitive, so if the current policy already allows flows from say β to γ , flows from α to γ would also be allowed in M .

Modelling scoped flows using flow locks is easy, but the global nature of policies in Almeida Matos and Boudol’s system, as opposed to our local policies on memory locations, needs special treatment. We introduce a lock $\sigma_{\alpha \prec \beta}$ for each pair of levels α and β that data could flow between. Each policy on some data must record the fact that a future flow declaration could allow that data to flow to many new locations due to the transitive nature of flows. Thus if a location in Almeida Matos and Boudol’s system would have level A , we could represent that as

$$A \cup \{ \sigma_{\alpha \prec \beta_0}, \sigma_{\beta_0 \prec \beta_1}, \dots, \sigma_{\beta_{k-1} \prec \beta_k} \Rightarrow \beta_k \mid \alpha \in A, \beta_i \notin A \}$$

where the $\notin$ is taken with respect to some universal set of levels. In effect, each location records all possible future transitive flows from it. We then derive our representation of the “flow” construct that opens a lock in the scope of some subprogram:

$$\text{flow } \sigma \text{ in } M \equiv \text{let } x = (\text{open } \sigma; M) \text{ in } (\text{close } \sigma; x)$$

Almeida Matos and Boudol also include parallel execution in their system, and as a consequence make their type system and semantic security definition, called *non-disclosure*, sensitive to possible non-termination. Our system has no parallel execution so we cannot model their full system, only the sequential subset.

Intransitive Noninterference. Flow locks represent a lower level abstraction than lattice-based information flow models in the sense that the lattice ordering is not “built in” but must be represented explicitly. One advantage of such a lower level view is that it can also represent *intransitive noninterference* policies [15, 14] — i.e. ones in which the flow relation is intentionally not transitive. Since intransitive policies are the default case for flow locks, it is straightforward to represent simple language-based intransitive policies such as the one described by Mantel and Sands [8].

Noninterference Until Declassification. Chong and Myers’ [5] introduce a class of temporal declassification policies. This is achieved by annotating variables with types of the form $k_0 \xrightarrow{c_1} \dots \xrightarrow{c_n} k_n$, which intuitively means that a variable with such an annotation may be successively declassified to the levels $k_1, \dots, k_n$, and that the conditions $c_1, \dots, c_n$ will hold at the execution of the corresponding declassification points. The exact nature of the conditions are left unspecified, and it is assumed in the type system that these conditions are verified at certain key program points by some external tool.

We can achieve a similar effect fairly naturally using flow locks, where we would use a distinct lock C_i for each condition c_i . One should then insert **open** C_i constructs in the program at points where the intended declassification takes place, and verify (with an external tool) that the corresponding condition c_i does indeed hold at these points, and that lock C_{i-1} has been opened (we assume that locks are never closed in this encoding). The policy above could then be represented as

$$\{k_0; \{C_1\} \Rightarrow k_1; \dots; \{C_1, \dots, C_n\} \Rightarrow k_n\}.$$

Robust Declassification. Information flow may be used to verify integrity properties, to ensure that untrusted (low integrity) data does not influence the values of trusted

(high integrity) data. Since flow lock policies are neutral with respect to whether we are dealing with confidentiality or integrity properties it is no problem to add such integrity policies to data, and we can easily have clauses for integrity and confidentiality in the same policy. The interesting case, however, is the interaction between confidentiality and integrity in the presence of dynamic policies.

Zdancewic and Myers [22] introduced the concept of *robust declassification* to characterise the property that an attacker (who controls low integrity data) cannot influence what is declassified. This guarantees that the attacker cannot manipulate the amount of information which is released through declassification.

In the setting of flow lock policies, “declassification” can be thought of as the process of opening locks, since whenever a lock is opened more flows are enabled. Thus we can interpret robust declassification as the question of whether low integrity data can influence the decision to open locks.⁵

One possible way of enforcing robust declassification using flow locks is to observe the following: since we cannot perform any computation with locks, the only way that an open operation can be influenced by low integrity data is via indirect information flow from low integrity data. Suppose that our policies use an indexed set of locks $\sigma_i, i \in I$ to control confidentiality. These are unguarded (i.e. we ignore *endorsement*). Let us assume that in addition to the actors of the system we have the pseudo-actor *trusted* used to track integrity information, just as we did in Section 2.

In order to prevent indirect flow from low integrity data to the opening of locks, we will log each use of an open operation by writing to a variable *log*. An obvious way to enforce this is to define a “robust” version of open:

$$\mathbf{ropen} \sigma_i \equiv \mathbf{open} \sigma_i; \mathit{log} := i$$

Now we give *log* the policy $\{\mathit{trusted}\}$. This ensures that the assignment is always safe from a confidentiality perspective (since normal actors can never read it anyway), and that the open operation can never have taken place in a low integrity context (since otherwise the assignment would cause information to flow from untrusted to trusted data). Finally, to additionally prevent the declassification of low integrity data we can syntactically enforce that lock-guarded policies are only used on high integrity data.

The Decentralized Label Model. In the Decentralized Label Model (DLM) [10, 11, 12], data is said to be *owned* by a set of principals. These principals may allow other principals to read the data, and the effective reader set is those principals that all owners agree may read the data. Allowing a new reader roughly corresponds to declassification, and we can model it similarly. The DLM also defines a global principal hierarchy, where one principal may allow another principal to *act for* it, which means it may read all the same things. This is very similar in spirit to introducing a new flow in the system by Almeida Matos and Boudol, including transitivity, and we can model it in the same way. Apart from clauses for declassification and hierarchic flows, the policies must also include clauses for the combination of the two, e.g. *A* can read the data if *B* owns it, has declassified it for *C* to read it, and *A* acts for *C*.

⁵ If we also take the view from [13], then we extend this concept with the requirement that we should not be able to declassify low integrity data.

A common extension of the DLM [22, 20, 19] deals with integrity and trust. The interesting part for us is the integration with the principal hierarchy, where if A trusts some data and A acts for B , then B also trusts that data. This can be modelled as the reverse of the normal clauses for transitive flows, and the clauses will be very similar to those for forward flows.

The complete general policy for a DLM variable encoded with flow locks would be fairly large and awkward, so we do not show it here.

Other Related Work. The JFlow language [9], as well as several recent papers [19, 23, 7], supports runtime mechanisms to enforce security in situations where this cannot be determined statically, e.g. permissions on a file that cannot be known at compile time. Our flow locks is a static, compile-time mechanism only, and thus cannot handle these issues.

Banerjee and Naumann [4] describe a combination of stack-based access control and information flow types to allow the static checking of policies such as “the method returns a result at level L unless the caller has permission p ”. It may be possible to encode these kinds of policies in a straightforward way using flow locks, but this remains a topic for future work.

6 Conclusions and Future Work

Flow locks are a very simple mechanism that generalises many existing systems and idioms for dynamic information flow policies. We have only just started looking at flow locks however, and much remains to be done.

To really establish flow locks as a core calculus, we need to show more formally how to embed other systems and idioms, and prove that our semantic condition is sufficiently strong compared to the semantic conditions of these other systems. It would also be worthwhile to look at extensions of our core system, in order to handle systems that we definitely cannot model at this point. Examples of such systems include the parallel execution of Almeida Matos and Boudol [2], and also systems that use various runtime mechanisms [19, 23, 7].

Furthermore, we would need to investigate how to implement the flow locks system as a programming language, and to determine what kinds of inference would be needed for policies and locks. Also, flow locks are fairly low-level in nature, being a raw mechanism for controlling data flows in a program. As such it is nontrivial to write and maintain correct flow lock programs. It would therefore be useful to look at what higher-level abstractions and design patterns that could be used together with flow locks. There exists some work specifically targeting the question of patterns, for instance the *seal* pattern by Askarov and Sabelfeld [3].

Acknowledgements. Thanks to Ulf Norell and our colleagues in the ProSec group for helpful comments, and to the anonymous referees for numerous helpful comments and suggestions. This work was partly supported by the Swedish research agencies SSF, VR and Vinnova, and by the Information Society Technologies programme of the European Commission, Future and Emerging Technologies under the IST-2005-015905 MOBIUS project.

References

1. M. Abadi, A. Banerjee, N. Heintze, and J. Riecke. A core calculus of dependency. In *Proc. ACM Symp. on Principles of Programming Languages*, pages 147–160, Jan. 1999.
2. A. Almeida Matos and G. Boudol. On declassification and the non-disclosure policy. In *Proc. IEEE Computer Security Foundations Workshop*, June 2005.
3. A. Askarov and A. Sabelfeld. Security-typed languages for implementation of cryptographic protocols: A case study. In *Proc. European Symp. on Research in Computer Security*, volume 3679 of *LNCS*, 2005.
4. A. Banerjee and D. A. Naumann. Stack-based access control and secure information flow. *Journal of Functional Programming*, 15(2):131–177, Mar. 2005.
5. S. Chong and A. C. Myers. Security policies for downgrading. In *ACM Conference on Computer and Communications Security*, pages 198–209, Oct. 2004.
6. D. E. Denning and P. J. Denning. Certification of programs for secure information flow. *Comm. of the ACM*, 20(7):504–513, July 1977.
7. M. Hicks, S. Tse, B. Hicks, and S. Zdancewic. Dynamic updating of information-flow policies. In *Proc. Foundations of Computer Security Workshop*, 2005.
8. H. Mantel and D. Sands. Controlled downgrading based on intransitive (non)interference. In *Proc. Asian Symp. on Programming Languages and Systems*, volume 3302 of *LNCS*, pages 129–145. Springer-Verlag, Nov. 2004.
9. A. C. Myers. JFlow: Practical mostly-static information flow control. In *Proc. ACM Symp. on Principles of Programming Languages*, pages 228–241, Jan. 1999.
10. A. C. Myers and B. Liskov. A decentralized model for information flow control. In *Proc. ACM Symp. on Operating System Principles*, pages 129–142, Oct. 1997.
11. A. C. Myers and B. Liskov. Complete, safe information flow with decentralized labels. In *Proc. IEEE Symp. on Security and Privacy*, pages 186–197, May 1998.
12. A. C. Myers and B. Liskov. Protecting privacy using the decentralized label model. *ACM Transactions on Software Engineering and Methodology*, 9(4):410–442, 2000.
13. A. C. Myers, A. Sabelfeld, and S. Zdancewic. Enforcing robust declassification. In *Proc. IEEE Computer Security Foundations Workshop*, pages 172–186, June 2004.
14. S. Pinsky. Absorbing covers and intransitive non-interference. In *Proc. IEEE Symp. on Security and Privacy*, pages 102–113, May 1995.
15. J. M. Rushby. Noninterference, transitivity, and channel-control security policies. Technical Report CSL-92-02, SRI International, 1992.
16. A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE J. Selected Areas in Communications*, 21(1):5–19, Jan. 2003.
17. A. Sabelfeld and D. Sands. Probabilistic noninterference for multi-threaded programs. In *Proc. IEEE Computer Security Foundations Workshop*, pages 200–214, July 2000.
18. A. Sabelfeld and D. Sands. Dimensions and principles of declassification. In *Proc. IEEE Computer Security Foundations Workshop*, 2005.
19. S. Tse and S. Zdancewic. Run-time principals in information-flow type systems. In *Proc. Symposium on Security and Privacy*, 2004.
20. S. Tse and S. Zdancewic. Designing a security-typed language with certificate-based declassification. In *Proc. European Symp. on Programming*, volume 3444 of *LNCS*, pages 279–294. Springer-Verlag, Apr. 2005.
21. D. Volpano, G. Smith, and C. Irvine. A sound type system for secure flow analysis. *J. Computer Security*, 4(3):167–187, 1996.
22. S. Zdancewic and A. C. Myers. Robust declassification. In *Proc. IEEE Computer Security Foundations Workshop*, pages 15–23, June 2001.
23. L. Zheng and A. Myers. Dynamic security labels and noninterference. In *Proc. Workshop on Formal Aspects in Security and Trust*, 2004.