

Dynamic Updating of Information-Flow Policies

Michael Hicks* Stephen Tse[†] Boniface Hicks[†] Steve Zdancewic[‡]
* University of Maryland [†] Pennsylvania State University [‡] University of Pennsylvania

Abstract

Applications that manipulate sensitive information should ensure *end-to-end* security by satisfying two properties: *sound execution* and some form of *noninterference*. By the former, we mean the program should always perform actions in keeping with its current policy, and by the latter we mean that these actions should never cause high-security information to be visible to a low-security observer. Over the last decade, security-typed languages have been developed that exhibit these properties, increasingly improving so as to model important features of real programs. No current security-typed language, however, permits general changes to security policies in use by running programs. This paper presents a simple information flow type system for that allows for dynamic security policy updates while ensuring sound execution and a relaxed form of noninterference we term *noninterference between updates*. We see this work as an important step toward using language-based techniques to ensure end-to-end security for realistic applications.

1 Introduction

Increasingly, personal and business information is being made available via networked infrastructures, so the need to protect the confidentiality of that information is becoming more urgent. A typical approach to enforcing data confidentiality is via access control. Unfortunately, access control only governs the release of information, not its propagation. Once a principal (e.g., a user, process, party, etc.) legally reads some data, he can freely share it, whether purposefully or inadvertently, despite the possible wishes of its owner. Instead, we would prefer applications to enforce *end-to-end* security by governing *information flow*: a principal should not, through error or malice, be permitted to transmit confidential information to an unauthorized party.

An information flow control system typically aims to enforce two properties: *noninterference* and *sound execution*. Given a principal hierarchy that defines the relative security levels of various principals, a program satisfies noninterference when it ensures that high security data is never visible, whether directly or indirectly, to low security observers. A program satisfies sound execution if it does not generate errors at run time. A typical way to satisfy these properties is to use a *security-typed language* [14] wherein the standard types on program variables include annotations to specify which principals are allowed to read. If a program type checks under a principal hierarchy, then, it is guaranteed that the program is noninterfering and sound with respect to the hierarchy. Security-typed languages are appealing because these properties are proven in advance of actual execution.

Typical security-typed languages assume that the principal hierarchy remains fixed during program execution. For long-running programs, such an assumption is unrealistic, as policies often change over time, e.g., to perform revocations [6, 2]. On the other hand, simply allowing the principal hierarchy to change at runtime could violate both the soundness and noninterference properties of running programs.

This paper presents a new security-typed language that allows dynamic updating of information-flow policies, particularly the delegation relations in the principal hierarchy. Section 2 defines a typical security-typed source language λ_{\leq}^{Π} , and then Section 4 defines λ_{tag}^{Π} as an extension of λ_{\leq}^{Π} with *tags* and the ability to accommodate updates to the principal hierarchy. We prove that λ_{\leq}^{Π} programs can be compiled to λ_{tag}^{Π} programs automatically, that these programs are sound, and that they respect a flavor of noninterference we dub *noninterference between*

$$\begin{array}{lll}
p ::= X \mid p, p & \ell ::= p : p \mid \ell, \ell & \mathcal{E} ::= \cdot \mid \mathcal{E} m \mid v \mathcal{E} \mid \text{if } \mathcal{E} m m \\
\Pi ::= \cdot \mid \Pi, p \leq p & u ::= \text{bool}_\ell \mid u \rightarrow u & \\
m ::= \text{true}_\ell \mid \text{false}_\ell \mid x \mid \lambda x : u. m \mid m m \mid \text{if } m m m \mid \text{if } (p \leq p) m m
\end{array}$$

Figure 1: Syntax of λ_{\leq}^{Π} : principals p , labels ℓ , permission Π , types u , terms m , and holes \mathcal{E} .

updates. To our knowledge, ours is the first system to safely permit general updates to the principal hierarchy, including revocations, in security-typed languages. Our discussion of the meaning of noninterference in the presence of revocation, and the definition of the term *noninterference between updates*, is also new. We believe our approach is an important step to making security-typed languages expressive enough to be used in real systems.

2 A Simple Security-Typed Language, λ_{\leq}^{Π}

To make our discussion of policy updates more concrete, we introduce a calculus λ_{\leq}^{Π} , a formalization of the *decentralized label model* [10] (DLM) based on the simply-typed lambda calculus. We present a discussion on policy updates, their challenges, and our solution to making them sound in the following two sections.

Figure 1 presents the syntax of λ_{\leq}^{Π} . Security policies specify confidentiality policies, defining which principals are allowed to read which data. Policies are specified in two parts. First, types and simple values are annotated with *labels* ℓ that consist of one or more pairs $(p_1 : p_2)$, where p_1 is the policy owner specifying p_2 as the reader. Principals p can be either literals X or *principal sets* (p_1, \dots, p_n) . A label with multiple pairs, written (ℓ_1, ℓ_2) can be used to specify more restrictive policies: a potential reader must satisfy all of the label restrictions. The second part of a DLM security policy is the principal hierarchy, or a permission context, Π is represented as a list of delegations between principals $p_1 \leq p_2$.

Terms m and types u are largely standard. We write evaluation as $\Pi \vdash m_1 \longrightarrow m_2$, which states that m_1 evaluates in a single step to become m_2 under runtime principal hierarchy Π . The principal hierarchy can be accessed dynamically using the run-time test of principal delegation $\text{if } (p_1 \leq p_2) e_1 e_2$ [17]:

$$\frac{\Pi \vdash p_1 \leq p_2}{\Pi \vdash \text{if } (p_1 \leq p_2) m_1 m_2 \longrightarrow m_1} \quad \frac{\Pi \not\vdash p_1 \leq p_2}{\Pi \vdash \text{if } (p_1 \leq p_2) m_1 m_2 \longrightarrow m_2}$$

The typing judgment has the form $\Pi; \Gamma \vdash m : \tau$, where Γ tracks the types of bound variables as usual, and permission context Π statically tracks knowledge of the principal hierarchy. Most rules are standard [17], as shown in Figure 2. The judgment $\text{lab}(u) = \ell$ returns the label of the type:

$$\text{lab}(\text{bool}_\ell) = \ell \quad \frac{\text{lab}(u_2) = \ell}{\text{lab}(u_1 \rightarrow u_2) = \ell}$$

There are two additional typing rules. The first one type-checks the run-time test of principal delegation. If the principal delegation test succeeds, the first branch can statically assume that $\Pi \vdash p_1 \leq p_2$ holds inside by adding¹ $p_1 \leq p_2$ to the permission context Π .

$$\frac{\Pi, p_1 \leq p_2; \Gamma \vdash m_1 : u \quad \Pi; \Gamma \vdash m_2 : u}{\Pi; \Gamma \vdash \text{if } (p_1 \leq p_2) m_1 m_2 : u} \quad \frac{\Pi; \Gamma \vdash m : u_1 \quad \Pi \vdash u_1 \preceq u_2}{\Pi; \Gamma \vdash m : u_2}$$

¹The second branch *cannot* assume $\Pi \vdash p_1 \leq p_2$; hence, there is no addition of constraints to the context for the second branch in the typing rule. Adding negative constraints $(p_1 \not\leq p_2)$ to the context is unnecessary, because subtyping can be decided with positive constraints.

$$\begin{array}{c}
\Pi; \Gamma, x : u \vdash x : u \qquad \frac{\Pi; \Gamma \vdash m_1 : \text{bool}_\ell \quad \Pi; \Gamma \vdash m_2 : u \quad \Pi; \Gamma \vdash m_3 : u \quad \text{lab}(u) = \ell}{\Pi; \Gamma \vdash \text{if } m_1 \ m_2 \ m_3 : u} \\
\\
\Pi; \Gamma \vdash \text{true}_\ell : \text{bool}_\ell \qquad \frac{\Pi; \Gamma, x : u_1 \vdash m : u_2}{\Pi; \Gamma \vdash \lambda x : u_1. m : u_1 \rightarrow u_2} \\
\\
\Pi; \Gamma \vdash \text{false}_\ell : \text{bool}_\ell \qquad \frac{\Pi; \Gamma \vdash m_1 : u_1 \rightarrow u_2 \quad \Pi; \Gamma \vdash m_2 : u_1}{\Pi; \Gamma \vdash m_1 \ m_2 : u_2}
\end{array}$$

Figure 2: Typing rules of λ_{\leq}^{Π} : $\Pi; \Gamma \vdash m : t$ (under hierarchy Π and context Γ , the term m has type t).

The second one is the subsumption rule that allows the flexibility of implicitly appealing to principal delegations in the permission context during typing. There exist straightforward and efficient algorithms [7, 17] for context subtyping $\Pi_1 \leq \Pi_2$ (meaning Π_1 is more permissive than Π_2), label subtyping $\Pi \vdash \ell_1 \sqsubseteq \ell_2$ (meaning ℓ_1 is less restrictive than ℓ_2 under Π), principal subtyping $\Pi \vdash p_1 \leq p_2$ (meaning principal p_1 is delegating to p_2 under Π), and type subtyping $\Pi \vdash u_1 \leq u_2$ (meaning u_1 is a subtype of u_2 under Π); we leave their formal specification to our companion technical report.

We have proved the desired properties of sound execution and noninterference for λ_{\leq}^{Π} by a sound translation to a target language, which is to be described in Section 4.

3 The Meaning of Policy Updates

Now we consider the means and the meaning of security policy updates in λ_{\leq}^{Π} (and similar languages). An information-flow security policy can change in two ways. First, the label on a particular piece of data might be altered, thereby making access to it more restricted or less restricted. The former case is permitted automatically by the subsumption rule: it is always safe to treat a piece of data more restrictively. The latter case is potentially dangerous, as the relabeling might expose sensitive information, so it is typically allowed only by an explicit *declassify* operation. A second way of changing the information-flow policy is to alter the principal hierarchy, which in turn alters the relative ordering between labels. Here again there are two kinds of changes: one can add a (directed) edge between two principals (e.g., the delegation $p_1 \leq p_2$), which corresponds to increasing the privileges of an observer. This kind of policy change is like a global form of declassification. One can also remove edges, corresponding to revocation, strengthening the security policy and decreasing the set of permissible information flows.

In this paper, we address the second style of policy update in which the principal hierarchy can change. To determine how and when the hierarchy should be permitted to change, we must consider the impact of policy updates on a program's security properties: sound execution and noninterference.

3.1 Models of updating

It is easy to see that naively allowing arbitrary policy updates could make evaluation *unsound*: the program could act in a way not consistent with its current policy. As an example, consider this simple program (extending the syntax presented earlier with `let` and integers):

$$\text{let } x : \text{int}_{p_2} \quad = \quad (\text{if } (p_1 \leq p_2) \ 1_{p_1} \ 3_{p_2}) \quad \text{in} \quad \text{if } (p_2 \leq p_3) \ x \ 2_{p_3}$$

If we evaluate this program under $\Pi = p_1 \leq p_2$, after one step the `if` branch succeeds yielding

$$\text{let } x : \text{int}_{p_2} = l_{p_1} \quad \text{in} \quad \text{if } (p_2 \leq p_3) \ x \ 2_{p_3} \tag{Ex 1}$$

Now say we wish to change the principal hierarchy to be $\Pi' = p_2 \leq p_3$. If we allow this update to occur, then the program's next evaluation step will be unsound. It will allow the data l_{p_1} to flow to variable x , whose label p_2 is not equal or higher under Π' . Clearly, this update should not be permitted.

Noninterference is more subtle because it is a global property: it ratifies all information flows that might occur during a program's *entire* evaluation relative to a single, fixed principal hierarchy. When the hierarchy can change, this definition no longer makes sense. One alternative is *noninterference between updates*, meaning that when a policy update occurs, the history of how some data received a certain label is forgotten, and the question of noninterference is reconsidered for the current program at the current policy. As motivation for this definition, imagine some data (a file, say) labeled as p_1 : under $\Pi'' = p_1 \leq p_2, p_1 \leq p_3$, meaning principals p_2 and p_3 are allowed to read it. If principal p_3 is fired and p_4 is hired, we might like to change the principal hierarchy to be $p_1 \leq p_2, p_1 \leq p_4$; i.e. to revoke the assertion $p_1 \leq p_3$ and add the assertion $p_1 \leq p_4$. From the point of view of the file labeled p_1 : and the original hierarchy Π'' , the changed hierarchy both rejects flows previously allowed (p_3 can no longer read the file) and permits flows not previously admitted (p_4 can read the file, but could not before). Thus, at least some portion of the file's information flow history must be forgotten to permit this intuitively reasonable policy change.

On the other hand, noninterference between updates can permit unintuitive, perhaps unintended flows. For example, say the program (Ex 1) takes an additional step under Π yielding

$$\text{if } (p_2 \leq p_3) \ l_{p_1} \ 2_{p_3}.$$

Under the initial principal hierarchy Π , the program would terminate with result 2_{p_3} . However, if we were to change the hierarchy to $\Pi''' = p_1 \leq p_2, p_2 \leq p_3$, then the `if` branch would be taken, and we would terminate with result l_{p_1} . The evaluation is sound and noninterfering under Π''' from the point of the change to termination. However, from the point of view of the entire program evaluation, the observed flow would have been disallowed under Π , and thus violates noninterference when considered relative to Π . Moreover, a program that satisfies noninterference between updates says nothing about the security ramifications of updates themselves. For example, one risk is that if an attacker can observe when a policy update occurs and what it consists of, he may be able to deduce the values of private data in the program.

An ideal security property would both permit policy updates to selectively “forget the past” and also reason about some flows across updates. It is an open question as to what information flow systems should enforce even when policy updates are disallowed—noninterference, though commonly supported, is too restrictive in practice. As a first step, for this paper, we ensure that program execution is sound and respects noninterference between updates, recognizing and expecting that a better property is needed. We plan to investigate stronger, adequately expressive security properties in future work.

3.2 Overview of approach

At first glance, defining principal hierarchy updates for λ_{\leq}^{Π} that ensure soundness and noninterference between updates may seem straightforward. In particular, we can show noninterference between updates by proving the standard notion of noninterference: the type system is parameterized by a fixed principal hierarchy that is enforced as usual, since it implies that as long as the policy does not change, the program is noninterfering.

Proving soundness would seem equally simple: to update the principal hierarchy Π to Π' while running program m requires that we simply type check m under Π' : if type checking succeeds then we permit the update. While this

approach is sound (by definition), it is overly restrictive. Consider our example program again. Say the program evaluates under Π and becomes as in Ex 1. Then say we wish to change the hierarchy to be Π' . This program will not type check since the expression 1_{p_1} cannot be given type int_{p_2} under Π' . But conceptually it should be legal, because x (which we have substituted for here with 1_{p_1}) should be treated as having type int_{p_2} , as defined in the original program. This fact is not revealed, however, in the run-time representation of the current state. That is, an important fact of the past (changing the type of 1_{p_1} to int_{p_2}) has been forgotten.

We can solve this problem by moving away from the view of *subtyping as subset* to the view of *subtyping as coercion* for evaluation. Rather than viewing data 1_{p_1} as having type int_{p_2} under Π , we say that we can *coerce* 1_{p_1} to a value that has type int_{p_2} ; that is, we can coerce it to 1_{p_2} . To do this, we extend λ_{\leq}^{Π} with *permission tags* that act as coercion functions. In particular, the expression $[\ell \sqsubseteq \ell'] 1_{\ell}$ will evaluate to $1_{\ell'}$. With this change, our original program becomes

$$\text{let } x : \text{int}_{p_2} = \text{if } (p_1 \leq p_2) ([p_1 \sqsubseteq p_2] 1_{p_1}) \ 3_{p_2} \quad \text{in} \quad \text{if } (p_2 \leq p_3) ([p_2 \sqsubseteq p_3] x) \ 2_{p_3}$$

That is, the uses of subsumption are made explicit as tags. Then the program will evaluate under Π to become

$$\text{if } (p_2 \leq p_3) ([p_2 \sqsubseteq p_3] 1_{p_2}) \ 2_{p_3}$$

Now we can see that changing to Π' will be legal, as 1 has a label that can be properly typed in the new policy. At the same time, we still prevent illegal updates to the policy. In the more general case that 1_{p_1} were some expression m_{p_1} , it would be unsound for the policy to change until m is a base value. Thus, while m is being evaluated (i.e., in the context of the *if* expression), it is guarded with the tag $[p_2 \sqsubseteq p_3]$. An update that violated this constraint would not be allowed (as desired).

In addition to providing a more flexible coercion semantics, it turns out that permission tags can also lead to a more efficient implementation. In particular, rather than having to type check the entire program body at each proposed change in policy, we only need to look at the tags, which succinctly capture how the current policy is being used. Section 4 presents a dynamic traversal that discovers these tags at update points without having to consider function bodies. We conjecture that with only a little more work, we can adjust the evaluation semantics to keep track of the current “tag context”. This would allow us to replace the traversal with a simple check.

3.3 Example

To show how these issues might arise in practice, we conclude this section with an example. Figure 3 shows a class for accessing the records of a company database, written in a Java-like syntax. This class defines two run-time principals *mgr*, which is a division manager, and *div*, which represents a division of company employees.² Lines 5 through 8 define some utility functions getting query inputs from the system user, processing them, creating summaries, and displaying information to the user. The policies on these methods establish that queries and the resulting processed data are owned by the *mgr* principal and readable by all principals in the group *div*, but that the results of auditing a query are only readable by *mgr*. These policies are explicit in the program: for example, the label on line 5 indicates that the result of *get_query* is owned by the principal *mgr* and readable by (principals in) the group *div*; similarly the *audit* method takes data readable by *div* and returns data only readable *mgr* (owners are implicitly considered to be readers in our model).

The method *access_records* is parameterized by a principal *emp* (employee), which is the current user of the database system. Line 15 dynamically checks that *emp* is a member of the division *div*, whose data is stored encapsulated in this database object. This line results in a runtime check of the principal hierarchy and succeeds only if $\text{div} \leq \text{emp}$ is true at the time when the check is made. Assuming that check succeeds, employee queries are received, processed, and displayed to the user until the user quits. In this scenario, the program also audits the

²Run-time principals represent principals as run-time entities, and could readily be added to our system [17].

```

01.     class Database {
02.         principal div;      /* division group */
03.         principal mgr;      /* manager for the division */
04.
05.         Query{mgr:div} get_query() {...}
06.         Data{mgr:div} process_query(Query{mgr:div} q) {...}
07.         Data{mgr:} audit(Data{mgr:div} d) {...}
08.         void display(principal p, Data{mgr:p} {...}
09.
10.     void access_records(principal emp) {
11.         Query{mgr:div} query;
12.         Data{mgr:emp} result;
13.         Data{mgr:} summary;
14.
15.         if (div < emp) { /* employee is a member of the division */
16.             while (true) {
17.                 query = get_query();
18.                 if (query == Quit) break;
19.                 result = process_query(query);
20.                 summary = audit(result);
21.                 display(emp, result);
22.
23.                 if (mgr < emp) { /* employee is a manager */
24.                     display(emp, summary);
25.                 }
26.                 ... /* log audit information */
27.             }
28.         } else { abort(); }
29.     }
30. }

```

Figure 3: Information-flow in a database system with principal delegations.

employee queries, perhaps to generate some statistics useful for making management decisions. The results of the audit process are readable only by managers (i.e. those principals p for which $\text{mgr} \leq p$). For convenience, if the user of the system *is* a manager, the results of the audit are displayed immediately—the dynamic check on line 23 ensures that only managers receive this sensitive data. Presumably the program would also log the auditing information for later inspection by a manager; in this case, the current user is not able to see that data.

The code makes an important assumption: though it checks $\text{div} \leq \text{emp}$ only once, it assumes that this relationship holds for the entire execution of the `while` loop. A problem arises if this relationship is revoked while the loop executes, say if the employee is fired or just moved to a different division. In this case, an employee who no longer belonged to a particular division would still have access to its files. Even worse, if the employee were made a manager (i.e., introducing $\text{mgr} \leq \text{emp}$ into the principal hierarchy) in a new division, he would suddenly have privileges not allowed under either policy—he could read files belonging to his original division. These scenarios reveal how policy changes can violate both sound execution and our intuitive notion of noninterference.

Introducing permission tags solves these problems. In particular, to store the returned value of `process_query` into `result`, the label of the returned value must be coerced from `mgr:div` to `mgr:emp`. This will be witnessed by a coercion $[\text{mgr:div} \sqsubseteq \text{mgr:emp}]$ on `process_query(query)`, which in turn will prevent the revocation of the

$$\begin{aligned}
t &::= \text{bool}_\ell \mid t \rightarrow t \\
v &::= \text{true}_\ell \mid \text{false}_\ell \mid \lambda[\Pi]x:u. e \quad \mathcal{E} ::= \cdot \mid \mathcal{E} e \mid v \mathcal{E} \mid \text{if } \mathcal{E} e e \mid [\ell \sqsubseteq \ell] \mathcal{E} \\
e &::= \text{true}_\ell \mid \text{false}_\ell \mid x \mid \lambda[\Pi]x:t. e \mid e e \mid \text{if } e e e \mid \text{if } (p \leq p) e e \mid [\ell \sqsubseteq \ell] e
\end{aligned}$$

Figure 4: Syntax of λ_{tag}^Π : types t , values v , terms e , holes \mathcal{E} .

edge $\text{div} \leq \text{emp}$ from the principal hierarchy.³

4 A language with dynamic policy and tagging, λ_{tag}^Π

This section formally describes an extension to λ_{\leq}^Π , called λ_{tag}^Π , that permits dynamic updates to the principal hierarchy. As just described, we use permission tags of the form $[\ell_1 \sqsubseteq \ell_2]$ to prevent illegal updates of the principal hierarchy during execution. We prove that λ_{tag}^Π enjoys the security properties of sound execution and noninterference described earlier, even as policies change at run-time. Permission tag annotations need not burden the programmer; they can be automatically inserted by the compiler. At the end of this section, we present an automatic translation from the source calculus presented earlier to the target calculus here, based on the standard formulation of *subtyping as coercions*, and prove it sound.

The syntax of λ_{tag}^Π is presented in Figure 4, and closely matches the source calculus, λ_{\leq}^Π in Section 2, except the addition of permission tags. The typing rules are the same, with one exception: the subsumption rule for subtyping is now eliminated, effectively replaced by the new typing rule for tags:

$$\frac{\Pi \vdash \ell_1 \sqsubseteq \ell_2 \quad \Pi; \Gamma \vdash e : \text{bool}_{\ell_1}}{\Pi; \Gamma \vdash [\ell_1 \sqsubseteq \ell_2] e : \text{bool}_{\ell_2}}$$

We maintain the invariant that the current principal hierarchy Π always respects the permission tag $[\ell_1 \sqsubseteq \ell_2]$ around the term e , as shown in the judgment $\Pi \vdash \ell_1 \sqsubseteq \ell_2$. Another invariant, which is enforced during the translation in Section 4.2, is that only boolean values are tagged. This permits the following evaluation rules:

$$\Pi \vdash [\ell_1 \sqsubseteq \ell_2] \text{true}_{\ell_1} \longrightarrow \text{true}_{\ell_2} \quad \Pi \vdash [\ell_1 \sqsubseteq \ell_2] \text{false}_{\ell_1} \longrightarrow \text{false}_{\ell_2}$$

Functions are not directly tagged: they contain future computations in the body that, unlike booleans, cannot be coerced to work under different security policies. Our strategy is to extend function terms with the permission context, $\lambda[\Pi']x:t. e$, such that the context Π' of the function body can be *summarized* to guard against illegal updates. The typing rules for functions and applications are:

$$\frac{\Pi \leq \Pi' \quad \Pi'; \Gamma, x:t_1 \vdash e : t_2}{\Pi; \Gamma \vdash \lambda[\Pi']x:t_1. e : t_1 \rightarrow t_2} \quad \frac{\Pi; \Gamma \vdash e_1 : u_1 \rightarrow u_2 \quad \Pi; \Gamma \vdash e_2 : u_1}{\Pi; \Gamma \vdash e_1 e_2 : u_2}$$

Given these tags, we can soundly and efficiently check if a policy update is legal during execution. We introduce *dynamic tag checking* $\Pi \vdash e$, as shown in Figure 5, for ensuring that principal hierarchy Π is valid with respect to the running program e . Any hierarchy is valid against boolean values. Section 4.1 proves that that the dynamic tag checking soundly approximates the static type checking with respect to the validity of policy updates.

At last, we formalize an update to the principal hierarchy of an evaluating program by defining a top-level evaluation relation. Under hierarchy Π , a program can either take a small evaluation step, or change to use the

³In a formal operational semantics, loops are implemented by expanding each iteration of the loop into to a fresh version of the original. Each fresh version would contain this tag, preventing any update that would violate it for the duration of the loop's execution.

$$\begin{array}{c}
\frac{\Pi \leq \Pi'}{\Pi \vdash \lambda[\Pi']x : \mathbf{t}. e} \quad \frac{\Pi \vdash e_1 \quad \Pi \vdash e_2}{\Pi \vdash e_1 e_2} \quad \frac{\Pi, p_1 \leq p_2 \vdash e_1 \quad \Pi \vdash e_2}{\Pi \vdash \text{if } (p_1 \leq p_2) e_1 e_2} \quad \frac{\Pi \vdash \ell_1 \sqsubseteq \ell_2 \quad \Pi \vdash e}{\Pi \vdash [\ell_1 \sqsubseteq \ell_2]e} \\
\\
\Pi \vdash \text{true}_\ell \quad \Pi \vdash \text{false}_\ell \quad \frac{\Pi \vdash e_1 \quad \Pi \vdash e_2 \quad \Pi \vdash e_3}{\Pi \vdash \text{if } e_1 e_2 e_3}
\end{array}$$

Figure 5: Tag checking $\Pi \vdash e$ (determines whether principal hierarchy Π is legal for the running program e).

pending hierarchy Π' . The latter step is only permitted if the new hierarchy is legal with respect to the current program, that is, if dynamic tag checking $\Pi \vdash e$ succeeds:

$$\frac{\Pi \vdash e \longrightarrow e'}{(\Pi; e) | \Pi' \longrightarrow (\Pi; e')} \quad \frac{\Pi' \vdash e}{(\Pi; e) | \Pi' \longrightarrow (\Pi'; e)}$$

Note that dynamic tag checking is meant to approximate an implementation. That is, while our formulation requires a traversal over the active part of the program (i.e., the part without functional terms), this traversal could be avoided by statically gathering the set of tags S, S' that appear in the body e_1, e_2 , respectively, of each $\text{if } (p_1 \leq p_2) e_1 e_2$ expression, and then annotating the if with a *tag constraint* $(p_1 \leq p_2 \Rightarrow S) \cup S'$ (similar to conditional types [1]). These tag constraints can be maintained to form a *tag context* at run-time, so that dynamic tag checking merely considers the current tag context, rather than the active part of the program.

4.1 Security theorems

To show that the execution of a program written in our calculus is sound, we prove that any well-typed, closed term runs without any error. To show that the information flow satisfies end-to-end security, we prove that any well-typed low-security term is noninterfering by the high-security data. These *type safety* and the *noninterference* properties are formally stated as follows.

Here \Downarrow is the top-level evaluation for the whole program, ignoring the number of policy updates, while \longrightarrow^* is the transitive-closure of the non-updating evaluations. Therefore, type safety is guaranteed during the evaluation of the whole program, but noninterference is guaranteed between updates (as discussed in Section 3.1).

Theorem 1 (Security of dynamic policy updating)

1. *Type safety during execution:* If $\Pi; \cdot \vdash e : \mathbf{t}$, then $(\Pi, e) \Downarrow (\Pi', v)$.
2. *Noninterference between updates:* If (1) $\Pi; x : \text{bool}_{\ell_1} \vdash e : \text{bool}_{\ell_2}$, and (2) $\Pi; \cdot \vdash v_1 : \text{bool}_{\ell_1}$, and (3) $\Pi; \cdot \vdash v_2 : \text{bool}_{\ell_1}$, and (4) $\Pi \vdash \ell_1 \not\leq \ell_2$, then $\Pi \vdash e\{v_1/x\} \longrightarrow^* v$ iff $\Pi \vdash e\{v_2/x\} \longrightarrow^* v$.

The proof for *type-safety* uses the standard technique of combining the progress and the preservation of a well-typed term. Some important lemmas for showing the soundness of tagging and tag checking are below. The first lemma states a well-typed *value* can also be well-typed under the empty principal hierarchy $\Pi = \cdot$, which is critical in the substitution lemma. The second states that the evaluation rule $\Pi \vdash \text{if } (p_1 \leq p_2) m_1 m_2 \longrightarrow m_1$ in Section 2 is type-preserving. The last lemma below shows that dynamic checking $\Pi \vdash e$ is a sound approximation of static type checking $\Pi; \Gamma \vdash e : \mathbf{t}$.

$$\begin{aligned}
\left[\frac{\Pi; \Gamma, \mathbf{x} : \mathbf{u}_1 \vdash \mathbf{m} : \mathbf{u}_2}{\Pi; \Gamma \vdash \lambda \mathbf{x} : \mathbf{u}_1. \mathbf{m} : \mathbf{u}_1 \rightarrow \mathbf{u}_2} \right] &= \lambda[\Pi] \mathbf{x} : \llbracket \mathbf{u}_1 \rrbracket. \llbracket \Pi; \Gamma, \mathbf{x} : \mathbf{u}_1 \vdash \mathbf{m} : \mathbf{u}_2 \rrbracket \\
\left[\frac{\Pi; \Gamma \vdash \mathbf{m} : \mathbf{u}_1 \quad \Pi \vdash \mathbf{u}_1 \preceq \mathbf{u}_2}{\Pi; \Gamma \vdash \mathbf{m} : \mathbf{u}_2} \right] &= \llbracket \Pi \vdash \mathbf{u}_1 \preceq \mathbf{u}_2 \rrbracket \llbracket \Pi; \Gamma \vdash \mathbf{m} : \mathbf{u}_1 \rrbracket \\
\llbracket \Pi \vdash \text{bool}_{\ell_1} \preceq \text{bool}_{\ell_2} \rrbracket &= \lambda[\Pi] \mathbf{x} : \text{bool}_{\ell_1}. [\ell_1 \sqsubseteq \ell_2] \mathbf{x} \quad (\text{fresh } \mathbf{x}) \\
\left[\frac{\Pi \vdash \mathbf{u}_3 \preceq \mathbf{u}_1 \quad \Pi \vdash \mathbf{u}_2 \preceq \mathbf{u}_4}{\Pi \vdash \mathbf{u}_1 \rightarrow \mathbf{u}_2 \preceq \mathbf{u}_3 \rightarrow \mathbf{u}_4} \right] &= \lambda[\Pi] \mathbf{x}_1 : \llbracket \mathbf{u}_1 \rrbracket. \lambda[\Pi] \mathbf{x}_2 : \llbracket \mathbf{u}_3 \rrbracket. \llbracket \Pi \vdash \mathbf{u}_2 \preceq \mathbf{u}_4 \rrbracket (\mathbf{x}_1 (\llbracket \Pi \vdash \mathbf{u}_3 \preceq \mathbf{u}_1 \rrbracket \mathbf{x}_2)) \quad (\text{fresh } \mathbf{x}_1, \mathbf{x}_2)
\end{aligned}$$

Figure 6: Translating principal delegations to permission taggings.

Lemma 2 (Soundness of dynamic tag checking)

1. If $\Pi; \Gamma \vdash \mathbf{v} : \mathbf{t}$, then $\cdot; \Gamma \vdash \mathbf{v} : \mathbf{t}$.
2. If $\Pi, \mathbf{p}_1 \leq \mathbf{p}_2; \Gamma \vdash \mathbf{e} : \mathbf{t}$ and $\Pi \vdash \mathbf{p}_1 \leq \mathbf{p}_2$, then $\Pi; \Gamma \vdash \mathbf{e} : \mathbf{t}$.
3. If $\Pi; \cdot \vdash \mathbf{e} : \mathbf{t}$, then $\Pi \vdash \mathbf{e}$. Moreover, if $\Pi; \cdot \vdash \mathbf{e} : \mathbf{t}$ and $\Pi' \vdash \mathbf{e}$, then $\Pi'; \cdot \vdash \mathbf{e} : \mathbf{t}$.

The proof for *noninterference* uses a logical relation for modeling the observable equivalence of a well-typed term with respect to an external observer, and shows that the substitutions preserve the equivalence [17]. Space precludes a formal development of the proofs here. Our companion technical report contains the complete rules of our calculus and the full proofs of both the type-safety and the noninterference properties. In addition, the type-safety property of the target language as well as the soundness of the translation in the next subsection are formally specified and mechanically verified⁴ in Twelf (a logical framework).

4.2 Translation from λ_{\preceq}^{Π} to λ_{tag}^{Π}

Figure 6 shows the translation rules from the typing derivation of a λ_{\preceq}^{Π} term \mathbf{m} to a typing derivation of a λ_{tag}^{Π} term \mathbf{e} . The main work is in the translation of the subsumption rule which takes the subtyping derivation of the source types and produces a well-typed *coercion function* in the target language. Breazu-Tannen et al. propose [3] such coercion semantics for the subtyping between types in the simply-typed lambda calculus. Our translation slightly extends the semantics for types with labels and permission tags. Our translation is sound as follows:

Theorem 3 (Soundness of permission tagging)

1. *Typing:* If $\llbracket \Pi; \Gamma \vdash \mathbf{m} : \mathbf{u} \rrbracket = \mathbf{e}$, then $\Pi; \llbracket \Gamma \rrbracket \vdash \mathbf{e} : \llbracket \mathbf{u} \rrbracket$.
2. *Subtyping:* If $\llbracket \Pi \vdash \mathbf{u}_1 \preceq \mathbf{u}_2 \rrbracket = \mathbf{e}$, then $\Pi; \Gamma \vdash \mathbf{e} : \llbracket \mathbf{u}_1 \rrbracket \rightarrow \llbracket \mathbf{u}_2 \rrbracket$.

We conjecture that the translation can be made *coherent* [3], meaning that target terms translated from different typing and subtyping derivations of the *same* source term have the same evaluation behavior. In particular, the

⁴We do not use the higher-order abstract syntax for encoding variable bindings. We have not performed the *totality check* which ensures that all proof cases have been completed — this property is verified externally by hand.

tag checking $\Pi \vdash e$ performs the same checks whether we tag the function or the argument of an application, hence *coherent* in allowing the same set of legal policy updates. To achieve such coherent translation, *algorithmic subtyping* must be used, instead of *declarative subtyping* as presented in this paper. The conversions and theories between these variants of subtyping are standard [12].

4.3 Discussion

As mentioned in Section 3.1, the fact that a program is noninterfering between updates says nothing of possible information flows across updates. Indeed, in the system described in this section, if an attacker p_3 can observe when updates occur, and what they consist of, it is possible for the timing of an update to communicate a secret value. Consider the following program:

```
let x = (if bp1: (λx:bool. truep1:) (λ[p2 ≤ p1]x:bool. [p2 :⊆ p1 :]truep2:)) in
let y = (if (p2 ≤ p1) truep3: falsep3:) in
let z = ... use x ... in y
```

Suppose that the program begins evaluating with principal hierarchy $\Pi = p_2 \leq p_1$ and that an update $\Pi' = \emptyset$ becomes available just after x has been computed (call this program p). In the case that b was true_{p_1} , then p would be

```
let x = λx:bool. truep1: in
let y = (if (p2 ≤ p1) truep3: falsep3:) in
let z = ... use x ... in y
```

Thus, the policy update succeeds and false_{p_3} is returned. On the other hand, if b was false_{p_1} , then p is

```
let x = λ[p2 ≤ p1]x:bool. [p2 :⊆ p1 :]truep2: in
let y = (if (p2 ≤ p1) truep3: falsep3:) in
let z = ... use x ... in y
```

Thus, the policy update is delayed due to the annotation $p_2 \leq p_1$ on the function x until z has been evaluated, meaning that true_{p_1} is returned. Hence, p_3 is able to observe b_{p_1} even though this is allowed by neither Π or Π' .

This particular example is an artifact of our dynamic tag checking algorithm, since it treats each branch of the initial `if` independently, once evaluated. A more static checking system, suggested earlier, would impose the same constraint on updates whichever function was chosen for x , and eliminate this flow. Nonetheless, the noninterference between updates property is too weak to illuminate this issue or its proposed fix, so we plan to consider refinements in future work.

5 Related Work

Security-typed languages for enforcing information flow control are a rich area of research [14]. Security policies are expressed as labels on terms and a principal hierarchy defining delegation relationships; in most systems this hierarchy is fixed at compile-time. Jif [10] and recent formal work [17, 19] support *runtime principals*, which make it possible for the hierarchy to grow at runtime, but do not allow revocations. Our calculus is the first to address generalized, dynamic updates to the principal hierarchy.

Security-type systems are intended to provide a *noninterference* guarantee [13, 4, 5, 9], modulo certain small-bandwidth information channels permitted for performance reasons (timing and termination channels) and an explicit “escape hatch” in the form of a robust downgrading mechanism. The introduction of such downgrading into these languages opened a new chapter in discussions about the meaning of noninterference that is still ongoing [18, 11, 8, 15]. As we have described, our dynamic policy updating is complementary to declassification.

Declassification, as typically used, relabels a data value from one label to another; policy updates as considered in this paper permit the relationship between the labels to change over time. Both features are necessary in practice, and both can potentially be abused—it is possible that work on structured uses of declassification, as provided by robustness [18, 11] or intransitive-noninterference [8] may apply to policy updates as well. We believe our discussion of dynamic policy updating here provides a new avenue for understanding the meaning of noninterference policies for realistic programs.

This work was inspired by a similar system called Proteus that we developed for ensuring type-safety of dynamic software updates [16]. In Proteus, users can define named types T . When given that type $T = \tau$ (for some type τ), treating a value of type T as a τ or vice versa requires an explicit coercion. When a program is dynamically updated to change the definition of T to be τ' , a dynamic analysis can check for these coercions in any functions not being updated (updated functions are assumed compatible with the new definition). If such a coercion is found then the update is only allowed if $\Gamma \vdash \tau' \preceq \tau$, where Γ is the updated type environment. This dynamic analysis is analogous to dynamic tag checking $\Pi \vdash e$, which essentially ensures for the new Π that $\Pi \vdash \ell_1 \sqsubseteq \ell_2$ for all tags $[\ell_1 \sqsubseteq \ell_2]$ in e . In Proteus, this dynamic analysis can be replaced with a simpler run-time given certain static information; we conjecture a similar result for λ_{tag}^{Π} .

Primarily in the context of public key infrastructures (PKI), the specific case of credentials revocation has been the subject of considerable study [6, 2]. This work has focused on the exploration of the fundamental tradeoff between security and cost. To simplify, on-line revocation servers effectively permit only a very small window of vulnerability for illicit use of compromised credentials, but often incur a high computation cost. Off-line systems provide a lower computational cost, but do so at the expense of longer latencies for receiving revocation notification. The issue of introducing policy revocation into a running program in a way that maintains sound execution has not been explored in the literature.

6 Conclusion

We have presented a new security-typed language that allows dynamic updating of information-flow policies, in particular the delegation relations in the principal hierarchy. Assumptions needed for sound execution can be represented within the program as *permission tags*, and a run-time tag checking mechanism can be used to prevent illegal updates to the principal hierarchy. Tags are implemented as run-time coercions that capture dynamic labeling behavior, which can prevent spurious rejection of legal policy updates. Tags are added to programs via an automatic translation from a standard source language. We are the first to formalize an information flow language that is sound yet permits dynamic revocations. Our language also satisfies *noninterference between updates*, which seems to us be a reasonable security property in the presence of updates. We hope that our work stimulates interest in making security-typed languages expressive enough to be used in real systems, where policies regularly change.

Acknowledgments We thank Jeff Foster, Peter Sewell, and the anonymous reviewers for their comments. Hicks is supported in part by NSF grant CCF-0346989 (*CAREER: Programming Languages for Reliable and Secure Low-level Systems*). Zdancewic and Tse are supported in part by NSF grant CCR-0311204 (*Dynamic Security Policies*) and NSF grant CNS-0346939 (*CAREER: Language-based Distributed System Security*).

References

- [1] Alexander Aiken, Edward L. Wimmers, and T. K. Lakshman. Soft typing with conditional types. In *Proceedings of the ACM Symposium on Principles of Programming Languages*, pages 163–173, 1994.
- [2] D. Boneh, X. Ding, G. Tsudik, and M. Wong. A method for fast revocation of public key certificates and security capabilities. In *Proceedings of USENIX Security Symposium*, pages 297–308, Aug 2001.

- [3] Val Breazu-Tannen, Thierry Coquand, Carl Gunter, and Andre Scedrov. Inheritance as implicit coercion. *Information and Computation*, 93:172–221, 1991.
- [4] J. A. Goguen and J. Meseguer. Security policies and security models. In *IEEE Symposium on Security and Privacy*, pages 11–20. IEEE Computer Society Press, April 1982.
- [5] J. A. Goguen and J. Meseguer. Unwinding and inference control. In *IEEE Symposium on Security and Privacy*, pages 75–86. IEEE Computer Society Press, April 1984.
- [6] Carl A. Gunter and Trevor Jim. Generalized Certificate Revocation. In *ACM Symposium on Principles of Programming Languages*, 2000.
- [7] Nevin Heintze and Jon G. Riecke. The SLam Calculus: Programming with Secrecy and Integrity. In *ACM Conference on Principles of Programming Languages (POPL)*, 1998.
- [8] Heiko Mantel and David Sands. Controlled declassification based on intransitive noninterference. In *Proceedings of the 2nd ASIAN Symposium on Programming Languages and Systems, APLAS 2004*, volume 3302 of *LNCS*, pages 129–145. Springer Verlag, 2004.
- [9] John McLean. Security models and information flow. In *IEEE Symposium on Security and Privacy*, pages 180–187. IEEE Computer Society Press, 1990.
- [10] Andrew C. Myers. Mostly-static decentralized information flow control. Technical Report MIT/LCS/TR-783, Massachusetts Institute of Technology, University of Cambridge, January 1999. Ph.D. thesis.
- [11] Andrew C. Myers, Andrei Sabelfeld, and Steve Zdancewic. Enforcing Robust Declassification. In *Proc. of 17th IEEE Computer Security Foundations Workshop*, pages 172–186, Asilomar, CA, June 2004. IEEE Computer Society Press.
- [12] Benjamin Pierce. *Types and Programming Languages*. MIT Press, 2002.
- [13] John C. Reynolds. Syntactic control of interference. In *Proc. 5th ACM Symp. on Principles of Programming Languages (POPL)*, pages 39–46, 1978.
- [14] Andrei Sabelfeld and Andrew C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, January 2003.
- [15] Andrei Sabelfeld and David Sands. Dimensions and principles of declassification. In *Proc. of the 18th IEEE Computer Security Foundations Workshop*, 2005.
- [16] Gareth Stoye, Michael Hicks, Gavin Bierman, Peter Sewell, and Iulian Neamtiu. *Mutatis Mutandis*: Safe and flexible dynamic software updating. In *Proceedings of the ACM Conference on Principles of Programming Languages (POPL)*, January 2005.
- [17] Stephen Tse and Steve Zdancewic. Run-time Principals in Information-flow Type Systems. In *IEEE 2004 Symposium on Security and Privacy*. IEEE Computer Society Press, May 2004.
- [18] Steve Zdancewic and Andrew C. Myers. Robust declassification. In *Proc. of 14th IEEE Computer Security Foundations Workshop*, pages 15–23, Cape Breton, Canada, June 2001.
- [19] Lantian Zheng and Andrew C. Myers. Dynamic security labels and noninterference. In *Proceedings of the 2nd International Workshop on Formal Aspects in Security and Trust (FAST)*, Toulouse, France, August 2004.