

# Encoding Information Flow in AURA

Limin Jia Steve Zdancewic

University of Pennsylvania \*

{liminjia, stevez}@seas.upenn.edu

## Abstract

Two of the main ways to protect security-sensitive resources in computer systems are to enforce access-control policies and information-flow policies. In this paper, we show how to enforce information-flow policies in AURA, which is a programming language for access control. When augmented with this mechanism for enforcing information-flow policies, AURA can further improve the security of reference monitors that implement access control.

We show how to encode security types and lattices of security labels using AURA's existing constructs for authorization logic. We prove a noninterference theorem for this encoding. We also investigate how to use expressive access-control policies specified in authorization logic as the policies for information declassification.

**Categories and Subject Descriptors** D.3.3 [Programming Languages]: Language Constructs and Features; D.4.6 [Operating Systems]: Security and Protection—Information flow controls, Access controls

**General Terms** Languages, Security

**Keywords** Access control, Authorization logic, Information flow control, Declassification, Security type system

## 1. Introduction

Almost all computer systems contain security-sensitive resources that need to be protected from untrusted applications. These include files, network connections, and private data such as a user's password or credit card number. Two of the main mechanisms for protecting these resources are

\* This research was sponsored in part by NSF Grants CNS-0524059, CCF-0716469, CNS-0346939 and DARPA Grant RA06-46. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF or DARPA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PLAS '09 June 15, Dublin, Ireland.

Copyright © 2009 ACM 978-1-60558-645-8/09/06...\$10.00

access control and information-flow analysis. Access control aims to prevent unauthorized principals—human users or other computer systems—from gaining access to the resources. Enforcing information-flow policies focuses on protecting the confidentiality of private data and makes sure that attackers cannot guess secrets by observing the behavior of multiple runs of a program [21, 26]. In this paper, we investigate how to enforce information-flow policies in AURA [29, 16], a language for access control. When augmented with this mechanism for enforcing information-flow policies, AURA can further improve the security of reference monitors that implement access control.

We begin by a brief overview of information-flow analysis and AURA, a language for access control.

**Enforcing information-flow policies** To protect the confidentiality of information, researchers design advanced type systems to enforce that secret input data cannot be leaked by observation of a system's public output. The key idea in information-flow type systems (see the survey by Sabelfeld and Myers [26]) is that program data are given security types that indicate their security levels. For instance, if a secret integer is protected at security level  $H$ , we give this integer the type  $\text{int}_H$ . The type system will guarantee that there is no information flow from high-security data to low-security data in well-typed programs. This property is referred to as noninterference.

However, information-flow policies that disallow any information flow from high security to low security are too draconian for real computer systems. Computer systems need to leak some amount of secret information to be useful. One classic example is the login program that compares user input with the password stored in the system, which is a secret. The boolean result of the comparison is not a secret because the login program has to either allow or deny access to the user. Therefore, attacker can always know if the password he typed in is the correct one or not. Another common example is that the average of all employees' salaries is released, but each individual salary has to be kept secret. Recently, there has been much work on controlled declassification of secret information [31, 22, 25, 18, 27, 10, 20, 6, 7, 8]. In the presence of declassification, the noninterference property does not hold.

**AURA, a language for access control** To ensure that only allowed principals can access protected resources, access-control requirements must be carefully defined and enforced. An *access-control policy* specifies whether a request by a principal to access a resource should be granted.

To clearly specify access-control policies and reason about them formally, researchers have developed authorization logics [5, 11, 13, 1, 2]. In logic-based access-control systems, logical proofs constructed using access-control policies serve as capabilities for accessing resources.

In these authorization logics, the formula  $A \text{ says } P$  expresses principals’ beliefs.  $A \text{ says } P$  states that principal  $A$  believes that  $P$  is true. For instance,  $Alice \text{ says } SkyIsPurple$  means that principal  $Alice$  believes that the sky is purple.  $SkyIsPurple$  is an assertion affirmed by a principal  $Alice$ . However, it is not necessarily the case that  $SkyIsPurple$  is true or that other principals believe it.

One desirable property of authorization logics is that principals should not interfere with each other’s beliefs. Without explicit delegation, what a principal  $A$  believes should not be affected by other principals’ beliefs. Such properties are also referred as noninterference properties [3, 14].

AURA is a language for implementing reference monitors for logic-based access control. AURA provides built-in support for specifying access-control policies. More specifically, the type system of AURA contains a constructive authorization logic based on DCC [2]. Programmers can manipulate authorization logic proofs as they do other language constructs. If implemented in AURA, a safe interface to access resources requires as an additional argument, a proof attesting that the access complies with the access-control policies. For example, a function  $playFor$ , which plays a song  $s$  on behalf of a principal  $p$ , might have the following type, which requires a proof that  $p$  is permitted to play  $s$ :

$$(s : Song) \rightarrow (p : \mathbf{prin}) \rightarrow \mathbf{pf}(\mathbf{self} \text{ says } MayPlay\ p\ s) \rightarrow Unit.$$

**Enforcing information-flow policies in AURA** Our work is inspired by the work on building a library for light-weight information-flow security in Haskell [24]. In that work, information-flow types are encoded as a Haskell data type  $(Sec\ s\ t)$  where  $s$  is the security level.  $Sec$  is implemented as a monad and a module system guarantees that attackers cannot extract secrets hidden in the monad.

We use very similar high-level ideas to encode information-flow types in AURA. Our advantage over the Haskell approach is that we can use constructs for AURA’s authorization logic for the encoding. The main idea of our encoding is that we use principals to represent security labels, and the type for a secret of type  $t$  protected at level  $H$  can be encoded as  $(x : \mathbf{pf}\ H\ \text{says}\ Reveal) \rightarrow t$ . Intuitively, without  $H$ ’s private key, no one can create an assertion of the type  $H \text{ says } Reveal$  and therefore secrets protected at level  $H$  can not flow to public channels.

The noninterference theorem of such encoding depends upon the noninterference properties of the authorization

logic. Furthermore, expressive access-control policies specified in authorization logic can be used to specify the policies for declassification.

**Contributions and roadmap** This paper makes the following contributions.

- We show how to encode information-flow types using authorization logics based on prior work [29, 16].
- We prove the basic noninterference theorem of our encoding. The key components of the proof are mechanized in the proof assistant Coq [12].
- We investigate through examples how declassification can be governed by access-control policies.

The rest of the paper is organized as follows. In Section 2, we review AURA. In Section 3, we explain how to encode information-flow types using AURA’s data types and the **says** monad. Next, in Section 4, we show how to prove the noninterference theorem for our encoding. Then, in Section 5, we extend our encoding and proof of noninterference to accommodate lattices of security labels. In Section 6, we investigate declassification. In the end, we discuss related work in Section 7.

## 2. AURA – A Language for Authorization and Audit

In this section, we give an overview of AURA to set up the background for the encoding of information-flow types in the next section. We will only discuss the high-level ideas. Technical details about the design of AURA can be found in our previous work [29, 16].

AURA is intended to be used to implement reference monitors for access control in security-sensitive settings. A reference monitor mediates access by allowing or denying requests to a resource (based, in this case, on policy specified in an authorization logic). For demonstrating key features of the language, we use an AURA implementation of a jukebox server as a running example.

### 2.1 Language Features

AURA is a call-by-value polymorphic lambda calculus. AURA consists of a “term-level” programming language for carrying out computation and a “proof-level” assertion language for writing proofs of access-control statements. AURA uses **Type** to classify the types of computations, and **Prop** to classify the types of proofs.

**Authorization logic** AURA allows programmers to define propositions like  $MayPlay$  using assertions. The following definition for  $MayPlay$  states that  $MayPlay$  takes a principal and a song as arguments and constructs a proposition.

$$\mathbf{assert}\ MayPlay : \mathbf{prin} \rightarrow Song \rightarrow \mathbf{Prop}$$

While assertions are similar in flavor to datatypes with no constructors, there is a key difference: there is no pattern-matching statement associated with these assertions. Assertions such as *MayPlay* are only used as constants affirmed by principals to specify access-control policies.

In AURA, *a* **says** *P* is a proposition stating that principal *a* believes that proposition *P* is true. There are a few different ways to create a proof for *a* **says** *P* in AURA. We can construct a term of type *a* **says** *P* from a proof *p* of *P* using the operation **return** *a p*. We can also create the proof by chaining other proofs about *a*'s beliefs using the bind operation written as (**bind** *x* : *Q* = *q* **in** *p*). Here *x* stands in for the proof of *Q* encapsulated by *q* and *p* is a proof of *a* **says** *P* using *x*.

For example, consider the principals *a* and *b*, the song *freebird*, and the assertion *MayPlay* introduced earlier. The statements

$$\begin{aligned} \text{ok} & : a \text{ says } (\text{MayPlay } a \text{ freebird}) \\ \text{delegate} & : b \text{ says } ((p : \text{prin}) \rightarrow (s : \text{Song}) \rightarrow \\ & \quad (a \text{ says } (\text{MayPlay } p \ s)) \rightarrow \\ & \quad (\text{MayPlay } p \ s)) \end{aligned}$$

assert that *a* gives herself permission to play *freebird* and delegates to *a* the authority to allow other principals to play the song. These two terms may be used to create a proof of *b* **says** (*MayPlay* *a* *freebird*) as follows:

$$\begin{aligned} \text{bind } d & : ((p : \text{prin}) \rightarrow (s : \text{Song}) \rightarrow \\ & \quad (a \text{ says } (\text{MayPlay } p \ s)) \rightarrow (\text{MayPlay } p \ s)) \\ & = \text{delegate} \\ \text{in return } b & (d \ a \ \text{freebird} \ \text{ok}). \end{aligned}$$

Such a proof could be passed to the *playFor* function if **self** is *b*, or it could be used to form a larger chain of reasoning.

In addition to uses of **return** and **bind**, AURA allows for the introduction of proofs of *a* **says** *P* without corresponding proofs of *P* by providing a pair of constructs, **say** and **sign**, that represent a principal's active affirmation of a proposition. The value **sign**(*a*, *P*) has type *a* **says** *P*; intuitively, we may think of it as a digital signature using *a*'s private key on proposition *P*.

Only the principal *a*—or, equivalently, programs with access to *a*'s private key—should be able to create a term of the form **sign**(*a*, *P*). We thus prohibit such terms from appearing in source programs and introduce the related term **say** *P*, which represents an effectful computation that uses the runtime's current authority—that is, its private key—to sign proposition *P*. When executed, **say** *P* generates a fresh value **sign**(**self**, *P*), where **self** is a built-in principal representing the current run-time authority.

It is worth noting that a principal can assert any proposition, even *False*. Because assertions are confined to the monad—thanks to the noninterference property of DCC—such an assertion can do little harm apart from making that particular principal's own assertions inconsistent.

**Dependent types** AURA incorporates dependent types: proofs in authorization logic can depend upon data, which allows for precise specification of access-control policies. For instance, the type of the proof that the *playFor* function requires is tied to the principal and the file arguments that *playFor* takes.

To simplify the meta-theory, AURA does not employ type-level reduction during type checking; and types only depend on values (i.e., well-formed normal forms). For instance, if *S* is a type constructor of the type  $(x : \text{Nat}) \rightarrow \text{Type}$ , then *S* (1 + 2) cannot be given a type in AURA because 1 + 2 is not a value; but *S* (1) has the type **Type**.

To make use of equalities obtained by run-time comparison of two values, AURA offers a type-refining equality test on *atomic* values—for instance, principals and booleans—as well as an explicit type cast between constructs of equivalent types. For example, when typechecking **if self** = *a* **then** *e*<sub>1</sub> **else** *e*<sub>2</sub>, the fact that **self** = *a* is automatically made available while typechecking *e*<sub>1</sub> (due to the fact that **prin** is an atomic type). Therefore, in *e*<sub>1</sub> proofs of type **self** **says** *P* can be cast to type *a* **says** *P* and vice-versa.

**The proof monad** AURA uses the constant **pf** : **Prop** → **Type** to wrap access-control proofs as program values. Similar to the **says** monad, we can construct terms of the type **pf** *P* by using **return**<sub>p</sub> *p* when *p* is a proof of *P*; or **bind**<sub>p</sub> *x* = *q* **in** *p* to chain proofs together<sup>1</sup>.

Such a separation between proofs and computations is necessary to prevent effectful program expressions from appearing in a proof term. The type of **say** *P* is **pf** (**self** **says** *P*). If **say** *P* was given type **self** **says** *P*, it would be possible to create a bogus “proof”  $\lambda x : \text{Prop. say } x$ ; the meaning of this “proof” would depend on the authority (**self**) of the program that applied the proof object.

**Summary of syntax** To simplify the presentation of AURA, it makes sense to unify as many of the constructs as possible. We thus adopt a lambda-cube style presentation [9] that uses the same syntactic constructs for terms, proofs, types, and propositions. A summary of AURA's core syntax is shown below.

$$\begin{array}{ll} \text{Terms } t & ::= x \mid \text{ctr} \mid \dots \\ & \quad \mid \lambda x : t_1. t_2 \mid t_1 t_2 \mid (x : t_1) \rightarrow t_2 \\ & \quad \mid \text{match } t_1 t_2 \text{ with } \{b\} \mid \langle t_1 : t_2 \rangle \\ \text{Branches } b & ::= \cdot \mid b \mid \text{ctr} \Rightarrow t \end{array}$$

In addition to the above common features ( $\lambda$ -abstraction, application, constructors, pattern matching, type cast, etc.),

<sup>1</sup> In formal definitions, to distinguish the bind and return operation for **says** monad from those for **pf** monad, we annotate the bind and return with a subscript *s* for **says** monad and *p* for **pf** monad. However, the type checker can easily tell them apart; therefore, in AURA programs, bind and return are overloaded for both monads.

the AURA-specific syntax is shown below.

$$\begin{array}{l}
t ::= \dots \mid \mathbf{Type} \mid \mathbf{Prop} \mid \mathbf{Kind} \mid \mathbf{prin} \mid a \mathbf{says} P \\
\mid \mathbf{pf} P \mid \mathbf{self} \mid \mathbf{sign}(a, P) \mid \mathbf{say} P \\
\mid \mathbf{return}_s a p \mid \mathbf{bind}_s x = e_1 \mathbf{in} e_2 \\
\mid \mathbf{return}_p p \mid \mathbf{bind}_p x = e_1 \mathbf{in} e_2 \\
\mid \mathbf{if} v_1 = v_2 \mathbf{then} e_1 \mathbf{else} e_2
\end{array}$$

AURA's value forms are as follows. We use metavariable  $v$  to denote values. We write  $\mathit{val}(e)$  to mean that  $e$  is a value.

$$\begin{array}{l}
v ::= x \mid \lambda x:t.e \mid \mathit{ctr} v_1 \dots v_n \mid \mathbf{self} \mid \mathbf{sign}(v, p) \\
\mid \mathbf{return}_s v p \mid \mathbf{bind}_s x = p \mathbf{in} q \mid \mathbf{return}_p v
\end{array}$$

**Signatures: data declarations and assertions** Programmers can define bundles of mutually recursive datatypes and propositions in AURA just as they can in other programming languages. A signature  $S$  collects these data definitions and, as a consequence, a well-formed signature can be thought of as a map from constructor identifiers to their types.

For instance, we can define the boolean type as follows:

$$\begin{array}{l}
\mathbf{data} \mathit{Bool} : \mathbf{Type} \{ \\
\mid \mathit{tt} : \mathit{Bool} \\
\mid \mathit{ff} : \mathit{Bool} \\
\}
\end{array}$$

Data definitions may be parametrized. For example, the familiar polymorphic list declaration is written as follows:

$$\begin{array}{l}
\mathbf{data} \mathit{List} : \mathbf{Type} \rightarrow \mathbf{Type} \{ \\
\mid \mathit{nil} : (t : \mathbf{Type}) \rightarrow \mathit{List} t \\
\mid \mathit{cons} : (t : \mathbf{Type}) \rightarrow t \rightarrow \mathit{List} t \rightarrow \mathit{List} t \\
\}
\end{array}$$

AURA's type system conservatively constrains **Prop** definitions to be inductive by disallowing negative occurrences of **Prop** constructors. Such a restriction is essential for consistency of the logic, since otherwise it would be possible to write loops that inhabit any proposition, including *False*.

## 2.2 Metatheory

The term typing judgment in AURA is written  $S; E \vdash t : s$ , where  $S$  is the signature containing definitions of data structures and assertions and  $E$  is the environment mapping variables to their types. We write  $S \vdash \diamond$  to denote the well-formed judgments for signatures, and  $S \vdash E$  to denote the well-formed judgments for environments. The small step operational semantics is denoted by  $e \mapsto e'$ .

We proved previously [16], the following properties of AURA. They will be useful in proving the noninterference properties of the information-flow type encoding in Section 4.

**Theorem 1** (Preservation). *If  $S; \cdot \vdash e : t$  and  $e \mapsto e'$ , then  $S; \cdot \vdash e' : t$ .*

**Theorem 2** (Progress). *If  $S; \cdot \vdash e : t$  then either  $\mathit{val}(e)$  or exists  $e'$  such that  $e \mapsto e'$ .*

**Theorem 3** (Typechecking is decidable).

- *If  $S \vdash \diamond$  and  $S \vdash E$ , then  $\forall e, \forall t$ , it is decidable whether there exists a derivation such that  $S; E \vdash e : t$ .*
- *If  $S \vdash \diamond$  then  $\forall E$  it is decidable whether there exists a derivation such that  $S \vdash E$ .*
- *It is decidable whether there exists a derivation such that  $S \vdash \diamond$ .*

We also proved that the **Prop** fragment of AURA is strongly normalizing. This theorem will allow us to conclude that despite the intricate dependencies on data, the authorization logic fragment is still logically consistent.

In AURA's core language, the proofs are computation free, meaning that we do not have reduction rules on proofs. For instance,  $\mathbf{bind}_s x = p \mathbf{in} q$  is a value. This is because the proofs are only meaningful as witnesses to access-control policies; and the reduction of proofs by reference monitors would not contribute significantly to the functionality of the system. We define proof reduction rules for the proofs in AURA, which will further reduce a "value" in the core language to a normal form according to the new reduction rules. We proved the following strong normalization theorem, details can be found in the companion tech report [17].

**Theorem 4** (The proofs in AURA are strongly normalizing). *If  $S; \cdot \vdash e : P$ , and  $S; \cdot \vdash P : \mathbf{Prop}$ , then  $e$  is strongly normalizing under the reduction rules for proofs.*

The noninterference proof also uses the following lemma stating that AURA's operational semantics is deterministic.

**Lemma 5** (AURA's operational semantics is deterministic). *If  $e \mapsto^* v_1$  and  $e \mapsto^* v_2$  and  $\mathit{val}(v_1), \mathit{val}(v_2)$  then  $v_1 = v_2$ .*

## 3. Encoding Information Flow Types

In this section, we explain how to use AURA's authorization logic constructs to encode information-flow types. These types are indexed by the security level, at which data is protected. For lucid explanation of the main ideas, we assume there is only one security level  $H$  and all secrets are protected at level  $H$ . We will extend this encoding in Section 5 to accommodate standard lattices for security labels.

In our encoding, security labels as treated as principals. To support the definitions of security lattices (here the lattice only contains one security label), we extend AURA's signature to allow the definitions of constants of the type **prin** for declaring security labels. We can declare  $H$  as follows:

**const**  $H : \mathbf{prin}$

Next, we define the assertion *Reveal*.

**assert** *Reveal* : **Prop**

In this simple encoding, we use a value of the type **pf** ( $H \mathbf{says} \mathit{Reveal}$ ) as the capability to access secrets protected at level  $H$ . *Reveal* is the same kind of assertion as *MayPlay* shown in the previous section. In AURA, there is no term witnessing the proof of *Reveal*; therefore, a proof of  $H \mathbf{says} \mathit{Reveal}$  can only be created by principal  $H$  actively

affirming it by signing *Reveal* using its private key. Furthermore, we assume that  $H$  is not the run-time authority **self**, whose private key is the only private key that programmers have access to. With the above two conditions, we know that programmers cannot produce a term that is a proof of the proposition  $H$  **says** *Reveal*. We define a data type for secrets protected at level  $H$  below:

```

data SecH : Type → Type {
  | mkSec : (t : Type)
    → (pf(H says Reveal) → t)
    → SecH t
}

```

*SecH* is a polymorphic type constructor. For instance, *SecH Bool* is the type for boolean expressions protected at  $H$ . The data constructor *mkSec* takes two arguments. The first argument is a type  $t$ . The second argument is a function that when applied to a term of type **pf** ( $H$  **says** *Reveal*), yields the secret of type  $t$ . The secret data is in effect guarded by a capability of type **pf** ( $H$  **says** *Reveal*). For example,  $s = \lambda x : \mathbf{pf} (H \text{ says } Reveal).3$  is a secret integer protected at level  $H$ . If there is a value  $v$  of the type **pf** ( $H$  **says** *Reveal*), evaluating  $s v$  will reveal the secret 3.

A term  $e$  of the type **pf**( $H$  **says** *Reveal*) belongs to the **Type** universe, meaning  $e$  is a computation. In AURA, programmers could write a non-termination computation  $\Omega$  of the type **pf**( $H$  **says** *Reveal*). However, this does not compromise our secret hidden in  $s$ , because AURA is call-by-value. Any attempt to execute  $s \Omega$  and extract the term of type  $t$  from  $s$  will result in non-termination.

The only way to get hold of a value of type **pf**( $H$  **says** *Reveal*) is when constructing another secret of type (*SecH*  $s$ ) using *mkSec*  $s$  ( $\lambda k : \mathbf{pf} (H \text{ says } Reveal).e$ ). Here  $k$  is a capability for access secrets protected at  $H$ , and  $k$  is available in  $e$ . This means that terms of type *SecH*  $t$  operate like a monad; a computation that manipulates a secret has to have type *SecH*  $t$ . We can encode the standard return and bind operation for *SecH*  $t$  monad.

To create an expression of type *SecH*  $t$  from an expression of type  $t$ , we can use the following *Return* function.

```

Return : (t : Type) → (d : t) → (SecH t) =
  λt : Type. λd : t. mkSec t (λkey : pf(H says Reveal).d)

```

```

1 Bind : (t : Type) → (s : Type) → (d : SecH t)
2   → (f : t → SecH s) → SecH s
3 = λt : Type. λs : Type. λd : SecH t. λf : t → SecH s.
4   mkSec s
5     (λk : pf(H says Reveal).
6       (match d with {
7         | mkSec →
8           λdt : (pf(H says Reveal) → t).
9             match (f (dt k)) with {
10            | mkSec →
11              λds : (pf(H says Reveal) → s). (ds k) }
12            })))

```

To operate on secrets, we can use the *Bind* function shown below. Given an expression  $d$  of type *SecH*  $t$ , and a function  $f$  that takes an expression of type  $t$  and produces an expression of type *SecH*  $s$ , *Bind* will apply  $f$  to the secrets in  $d$  and produce a term of type *SecH*  $s$ .

In the body of *Bind*, we need to apply function  $f$  (line 3) to the secret hidden in  $d$ . To extract the secret in  $d$ , we need a capability of type **pf** ( $H$  **says** *Reveal*). We can use such a capability  $k$  (line 5) in the body of the function we construct between line 5 and 12. We know  $d = \mathbf{mkSec} H dt$  by pattern matching on  $d$  on line 6. The term  $dt k$  has type  $t$  because  $dt$  has type **pf** ( $H$  **says** *Reveal*)  $\rightarrow t$  (line 8).  $dt k$  is the secret in  $d$ . The function application  $f (dt k)$  on line 9 has type *SecH*  $s$ . We need to construct a term of type  $s$ , because the function between line 5 and 12 has type **pf** ( $H$  **says** *Reveal*)  $\rightarrow s$ . We pattern match on  $(f (dt k))$  (line 9 – 11) and use  $k$  on line 11 to reveal the secret of type  $s$ .

Our encoding hides the expression that is a secret under a lambda abstraction, and because AURA does not evaluate under lambda abstractions, the computation in *SecH*  $t$  is lazy. A secret will not be evaluated until a capability for accessing the secret is provided.

## 4. Proof of Noninterference

To demonstrate that our encoding indeed protects secrets properly, we prove the noninterference theorem for our encoding. The main part of the proof is mechanized in Coq. The only paper proof is the proof of the noninterference property of AURA's authorization logic.

A noninterference proof is a proof of program equivalence. We want to prove that two programs containing different secrets should *behave the same* to the public observer. Here we use the termination-insensitive definition. We only enforce the equivalence between two programs when they both terminate. We use the squared semantics proof technique introduced by Pottier and Simonet [23]. The main idea of this approach is to define an extended language with a pair expression. The execution trace of a pair expression captures a pair of execution traces that could potentially contain different secrets. Proving the noninterference theorem is reduced to showing that two execution traces containing different secrets result in the same value in the extended language. Some of the challenges of using this techniques are 1) deciding where to introduce the pair expression so that the operational semantics can capture a pair of evaluation traces containing different secrets, and 2) introducing the pair expression in AURA in such a way that it works correctly with AURA's other language features such as dependent types.

The rest of this section is organized as follows. First, in Section 4.1, we introduce the design of AURA-PAIR, AURA extended with a pair construct. Next, in Section 4.2, we build connections between AURA and AURA-PAIR through a set of lemmas mapping the typing and evaluation relations be-

tween the two languages. Finally in Section 4.3, we discuss the proof of the noninterference theorem.

## 4.1 AURA-PAIR

We define AURA-PAIR by extending AURA with an expression denoting a pair of AURA expressions.

### 4.1.1 Syntax

We use meta-variables  $\hat{t}$  and  $\hat{e}$  to denote terms in AURA-PAIR, and use  $t$  and  $e$  to denote AURA terms. A summary of the syntax of AURA-PAIR is shown below. In addition to all the constructs in AURA, the definition of  $\hat{t}$  includes a new construct  $\langle t_1 \mid t_2 \rangle$ . Since we syntactically require that  $t_1$  and  $t_2$  are terms from AURA, nested pair expressions are ruled out by this definition.

AURA-PAIR Terms

$$\hat{t} ::= x \mid ctr \mid \lambda x:\hat{t}_1.\hat{t}_2 \mid \hat{t}_1 \hat{t}_2 \mid (x:\hat{t}_1) \rightarrow \hat{t}_2$$

AURA-PAIR Values

$$\hat{v} ::= \lambda x:\hat{t}.\hat{v} \mid \dots \mid \langle v_1 \mid v_2 \rangle$$

We also extend the values to include pair values, where each component in the pair is a value in AURA.

Before we define typing rules for AURA-PAIR, we introduce a few auxiliary definitions. First, we define a floor function that takes a term in AURA-PAIR and returns an AURA term that corresponds to either the left ( $i = 1$ ) or the right ( $i = 2$ ) part of the pair.

$$\boxed{[\hat{t}]_i = t} \quad [x]_i = x \quad [c]_i = c$$

$$[ctr]_i = ctr \quad [\hat{t}_1 \hat{t}_2]_i = [\hat{t}_1]_i [\hat{t}_2]_i$$

$$\dots \quad [\langle t_1 \mid t_2 \rangle]_i = t_i$$

For most constructs, the floor function is pushed into the sub-terms. For the pair expression, we return the sub-components of the pair right away. For simplicity of presentation, we assume there is an implicit injection from an AURA term to an AURA-PAIR term. In our Coq proof, we defined such a function explicitly.

Since the pair expressions are not allowed to be nested in AURA-PAIR, we define a special capture-avoidance substitution for AURA-PAIR as follows.

$$\boxed{\hat{u}[\hat{t}/x] = \hat{u}'}$$

$$(\hat{u}_1 \hat{u}_2)[\hat{t}/x] = (\hat{u}_1[\hat{t}/x]) (\hat{u}_2[\hat{t}/x])$$

$$\dots$$

$$\langle \langle u_1 \mid u_2 \rangle \rangle[\hat{t}/x] = \langle u_1 \{ [\hat{t}]_1/x \} \mid u_2 \{ [\hat{t}]_2/x \} \rangle$$

For most cases, the substitution is standard. For the pair expression (last rule above), we use the term substitution in AURA, and substitute the floor of the term to be substituted ( $\hat{t}$ ) for the variables in the sub-components of the pair ( $u_1$  and  $u_2$ ). Notice that  $u_i \{ [\hat{t}]_i/x \}$  is an AURA term. If we substitute an expression containing a pair into another pair expression, this substitution will make sure that the resulting expression does not contain nested pairs.

### 4.1.2 Operational Semantics

We use  $\hat{e} \mapsto_p \hat{e}'$  to denote the small-step operational semantics of AURA-PAIR. Most evaluation rules are the same as the ones in AURA. The interesting reduction rules for AURA-PAIR are shown in Figure 1. For the APP rule, we use the special substitution defined above. Three additional rules are defined for evaluating the pair expression. The first two evaluate the terms inside a pair using the reduction rules in AURA. The last one lifts the pair when an application occurs. In Pottier and Simonet's original system, there is one lifting rule for each beta redex. We only have one such lifting rule for AURA-PAIR despite the fact that AURA has many beta redexes such as (**match** ( $c v_1 \dots v_n$ )  $t$  **with**  $\{b\}$ ). The reason is that the typing judgments for AURA-PAIR restrict the appearance of the pair expression to function applications. This drastically simplifies the design of AURA-PAIR since we eliminated unnecessary lifting rules.

$$\frac{val(\hat{v})}{(\lambda x:\hat{t}.\hat{e}) \hat{v} \mapsto_p \hat{e}[\hat{v}/x]} \text{ APP}$$

$$\frac{e_1 \mapsto_p e'_1}{\langle e_1 \mid e_2 \rangle \mapsto_p \langle e'_1 \mid e_2 \rangle} \text{ PAIR-1} \quad \frac{e_2 \mapsto_p e'_2}{\langle e_1 \mid e_2 \rangle \mapsto_p \langle e_1 \mid e'_2 \rangle} \text{ PAIR-2}$$

$$\frac{val(v_1) \quad val(v_2) \quad val(\hat{v}_3)}{\langle v_1 \mid v_2 \rangle \hat{v}_3 \mapsto_p \langle v_1 [\hat{v}_3]_1 \mid v_2 [\hat{v}_3]_2 \rangle} \text{ LIFT}$$

Figure 1. Operational Semantics

### 4.1.3 Typing Rules

The typing judgment for AURA-PAIR is written  $S; E \vdash^p \hat{e} : \hat{t}$ . The only new typing rule is the rule for the pair expression, shown below. All other rules are the same as those in AURA.

$$S \vdash^p E \quad S; [E]_i \vdash t_i : [(x:u_1) \rightarrow u_2]_i$$

$$S; \cdot \vdash^p (x:u_1) \rightarrow u_2 : k$$

$$\frac{\nexists v \text{ such that } val(v) \text{ and } S; \cdot \vdash v : [u_1]_i}{S; E \vdash^p \langle t_1 \mid t_2 \rangle : (x:u_1) \rightarrow u_2} \text{ PAIR}$$

We assign an arrow type  $(x:t_k) \rightarrow t$  to the pair expression, because the pair expression represents a pair of secrets, which have type (**pf** ( $H$  **says** *Reveal*)  $\rightarrow t$ ). The first argument of the arrow is the capability that cannot be forged. We enforce this by requiring that there is no value of such type under an empty context. Each sub-component of the pair is type checked under the floor of the result type. The floor operation is crucial for us because AURA is dependently typed, and the types may contain pair expressions as well.

It is strange to have a negation in the typing rules. The PAIR rule is still inductively defined because we are using the already-defined AURA's typing relation, and we have proven the decidability of the typing relation in AURA. Furthermore, this type system is never meant to be used to check programs. It is used to illustrate the noninterference properties of AURA. We do not have to consider the efficiency of using such a typing rule.

We proved progress and preservation theorems for AURA-PAIR. Since we already have Coq proofs for AURA, it was not too hard to change the proofs to prove the soundness of AURA-PAIR. In Pottier and Simonet’s original paper, only preservation of the extended language is proven. The progress property simplifies the noninterference proof since we do not need to consider situations where AURA-PAIR might get stuck.

**Theorem 6 (Preservation).** *If  $S; \cdot \vdash^p \hat{e} : \hat{t}$  and  $\hat{e} \mapsto_p \hat{e}'$ , then  $S; \cdot \vdash^p \hat{e}' : \hat{t}$ .*

**Theorem 7 (Progress).** *If  $S; \cdot \vdash^p \hat{e} : \hat{t}$  then either  $\text{val}(\hat{e})$  or exists  $\hat{e}'$  such that  $\hat{e} \mapsto_p \hat{e}'$ .*

## 4.2 Connections Between AURA and AURA-PAIR

The point of defining AURA-PAIR is to compare two AURA programs. Here we establish the connection between programs in AURA and AURA-PAIR at both the typing and operational levels.

First, we establish the mapping between the special substitution in AURA-PAIR and the substitution in AURA.

**Lemma 8 (Floor of Substitution).**

$$[\hat{e}_2[\hat{e}_1/x]]_i = [\hat{e}_2]_i\{[\hat{e}_1]_i/x\}$$

Lemmas 9 and 10 concern the mapping of typing relations between AURA and AURA-PAIR. Lemma 9 states that if an expression  $\hat{e}$  is well-typed in AURA-PAIR, then both its left and right projection are well-typed in AURA. Lemma 10 states that a well-typed term in AURA is also well-typed in AURA-PAIR. We define  $\lfloor E \rfloor_i$  to be the point-wise lifting of the floor function on the environment  $E$ .

**Lemma 9 (Typing Soundness of AURA-PAIR).**

*If  $S; E \vdash^p \hat{e} : \hat{t}$  then  $S; \lfloor E \rfloor_i \vdash \lfloor \hat{e} \rfloor_i : \lfloor \hat{t} \rfloor_i$ .*

**Lemma 10 (Typing Completeness of AURA-PAIR).**

*If  $S; E \vdash e : t$  then  $S; E \vdash^p e : t$ .*

The next two lemmas concern the evaluation behavior. The first lemma, Lemma 11, states that if a term  $\hat{e}$  in AURA-PAIR evaluates to a value  $\hat{v}$ , then both the left and right projection in  $\hat{e}$  should evaluate to values in AURA. This lemma tells us that AURA-PAIR adequately represents two traces of evaluation in AURA. The next lemma, Lemma 12, states that if both of the right and left projection of  $\hat{e}$  evaluate to values in AURA, then  $\hat{e}$  should evaluate to a value in AURA-PAIR. This lemma tells us that AURA-PAIR faithfully models the termination behavior of AURA.

**Lemma 11 (Soundness of the Evaluation of AURA-PAIR).**

*If  $S; \cdot \vdash^p \hat{e} : \hat{t}$  and  $\hat{e} \mapsto_p^* \hat{v}$  then  $\lfloor \hat{e} \rfloor_i \mapsto^* \lfloor \hat{v} \rfloor_i$*

**Lemma 12 (Completeness of the Evaluation of AURA-PAIR).** *If  $S; \cdot \vdash^p \hat{e} : \hat{t}$  and  $\lfloor \hat{e} \rfloor_i \mapsto^* v_i$  where  $v_i$  is a value and  $i \in \{1, 2\}$  then  $\exists \hat{u}$  such that  $\hat{e} \mapsto_p^* \hat{u}$  and  $\hat{u}$  is a value.*

## 4.3 Noninterference

We use the following macros throughout this section.

$HKey = \mathbf{pf} (H \text{ says } \text{Reveal}) \quad SecHB = SecH \text{ Bool}$

We define a function  $CTROF(S, T)$  that takes a signature  $S$  and a type constructor  $T$  as arguments and returns the list of data constructors associated with  $T$ . For instance,  $CTROF(S, Bool) = \{tt, ff\}$ , if  $S$  contains the definition of  $Bool$ .

As we have mentioned in previous sections, the key idea of the encoding is to use  $HKey$  to guard secret data. We state this in the following lemma.

**Lemma 13 (Secret).**  $\nexists v, \text{val}(v)$  and  $S; \cdot \vdash v : HKey$ .

*Proof.* By contradiction. We use the strong normalization result of AURA, and the fact that programmers cannot generate the value  $\text{sign}(H, \text{Reveal})$ .

Assume  $S; \cdot \vdash v : HKey$

By Canonical Form,  $HKey = \mathbf{pf} (H \text{ says } \text{Reveal})$ ,

$$v = \mathbf{return}_p q \tag{1}$$

By Inversion of  $S; \cdot \vdash v : HKey$ ,

$$S; \cdot \vdash q : H \text{ says } \text{Reveal} \tag{2}$$

By Strong Normalization results of AURA,

$$q \mapsto^* q' \text{ and } q' \text{ is in normal form} \tag{3}$$

By Canonical Form, and  $q' \neq \text{sign}(H, \text{Reveal})$ ,

$$q' = \mathbf{return}_s H c \text{ and } S; \cdot \vdash c : \text{Reveal} \tag{4}$$

By Canonical Form,

$$c \in CTROF(S, \text{Reveal}) = \{ \} \tag{5}$$

Contradiction

□

The lemma assures us that no one can fabricate a value that has type  $HKey$ .

We prove the following noninterference theorem.

**Theorem 14 (Noninterference).** *If  $S; x : SecHB \vdash e : Bool$  and given any two values  $v_1, v_2$  such that  $S; \vdash v_1 : SecHB$   $S; \vdash v_2 : SecHB$  and  $e\{v_1/x\} \mapsto^* w_1$  and  $e\{v_2/x\} \mapsto^* w_2$  where  $w_1, w_2$  are values, then  $w_1 = w_2$ .*

The proof is shown in Figure 2. To clearly present the structure of the proof, we write the proof in two columns. The left column contains statements in AURA and the right one contains statements in AURA-PAIR. The arrows between the two columns are labeled with lemmas from Section 4.2 that connect the properties of AURA and AURA-PAIR. The statements in gray boxes are assumptions of the noninterference theorem. The statement in the framed box is the conclusion.

The proof starts from the left column. First, we examine the values  $v_1$  and  $v_2$  and extract the sub-terms  $f_i$ , which contain secrets guarded by  $HKey$ . Next, using Lemma Secret (Lemma 13), we conclude that there is no value of type  $HKey$ , which allows us to go to the AURA-PAIR side and construct a value pair  $\langle f_1 \mid f_2 \rangle$ . Now the evaluation of expression  $e[\hat{v}/x]$  captures the two evaluation traces containing different secrets. We stay on the AURA-PAIR side until we know that  $e[\hat{v}/x]$  evaluates to a value  $\hat{u}$  using the Evaluation Completeness Lemma (Lemma 12). Using the Evalua-

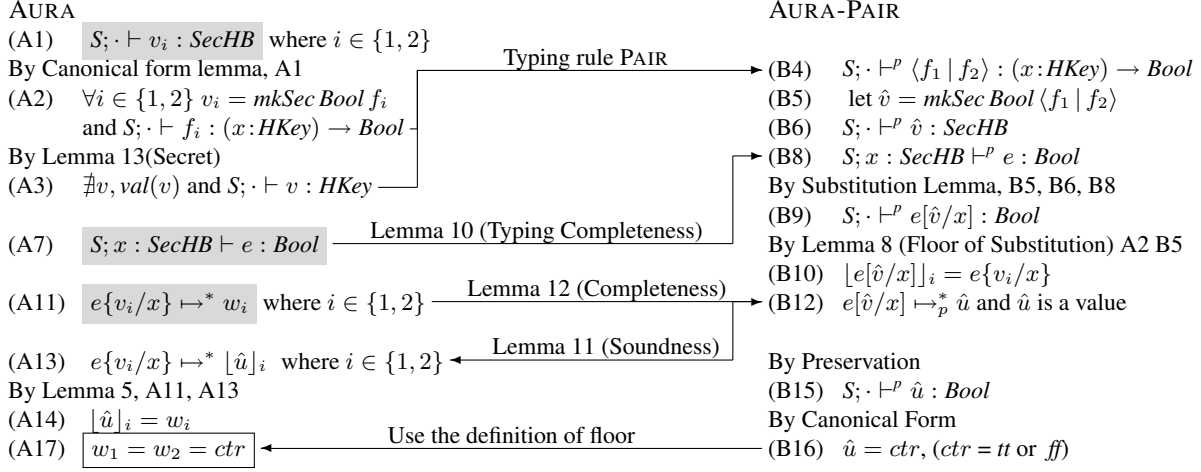


Figure 2. Proof of non-interference theorem

tion Soundness Lemma, we go back to AURA and conclude that  $e\{v_i/x\}$  evaluates to the floor of  $\hat{u}$ . Because AURA's reduction rules are deterministic (Lemma 5), we know that  $w_i$  is the same as the floor of  $\hat{u}$ . Now, we go to the AURA-PAIR side and gather more facts about  $\hat{u}$ . Because value  $\hat{u}$  is of type *Bool*, we know that  $\hat{u}$  has to be either the data constructor *tt* or *ff*. Because the floor of a constructor is itself, we know that both  $w_1$  and  $w_2$  have to be the same constructor.

## 5. Extension to Lattices

So far, we only considered single-level security where all secrets are protected at *H*. It is useful to have multi-level security where information is protected at several different security levels. For instance, a document could be classified as top secret, secret or public. We use a security lattice  $\langle \mathcal{L}, \sqsubseteq \rangle$  to model multi-level security.  $\mathcal{L}$  is a set of labels and  $\sqsubseteq$  is a partial order on labels in  $\mathcal{L}$ . The information-flow policy captured by the security lattice is that if  $\ell_1 \sqsubseteq \ell_2$ , then information protected at  $\ell_2$  is more secret than information protected at  $\ell_1$ , and information can only flow from  $\ell_1$  to  $\ell_2$ .

We extend our encoding to enforce information-flow policies specified by security lattices. Throughout this section, we consider a two-point security lattice with labels: *H* and *L*, and the partial order between them:  $L \sqsubseteq H$ . The techniques for encoding the security lattice and proving noninterference can be carried over to handle more general security lattices.

### 5.1 Extended Encoding

Both *H* and *L* are constants of type **prin**. The partial order  $L \sqsubseteq H$  is encoded using delegation in authorization logic as  $L2H : L \text{ says } (H \text{ says } \text{Reveal} \rightarrow \text{Reveal})$

In an implementation of the lattice in AURA,  $L2H$  can be the expression  $\text{sign}(L, H \text{ says } \text{Reveal} \rightarrow \text{Reveal})$ . It is an active affirmation by principal *L* by signing the proposition  $H \text{ says } \text{Reveal} \rightarrow \text{Reveal}$  using its private key.

Using  $L2H$  and  $hk : \text{pf}(H \text{ says } \text{Reveal})$ , we can construct a term of the type  $\text{pf}(L \text{ says } \text{Reveal})$  as follows:

$$lk : \text{pf}(L \text{ says } \text{Reveal}) =$$

$$\text{bind } h : H \text{ says } \text{Reveal} = hk \text{ in}$$

$$\text{bind } del : (H \text{ says } \text{Reveal} \rightarrow \text{Reveal}) = L2H \text{ in}$$

$$(\text{return } (\text{return } L (del\ h)))$$

Whenever we have a capability to reveal secrets protected at level *H*, we can obtain a capability to reveal secrets protected at level *L*.

We define the type constructor of security types below. The type constructor *Sec* takes as the first argument, the security level at which data is protected.

```

data Sec : prin → Type → Type {
  | mkSec : (l : prin) → (t : Type)
    → (pf(l says Reveal) → t)
    → Sec l t
}

```

Using the above definition, we can define the type for booleans protected at security level *L* as  $Sec\ L\ Bool$ , and the type for booleans protected at security level *H* as  $Sec\ H\ Bool$ .

The encodings of *return* and *bind* are similar to the ones in Section 3. The only difference is that we need to propagate the security label in the encoding.

$$\text{Return} : (l : \text{prin}) \rightarrow (t : \text{Type}) \rightarrow (d : t) \rightarrow (Sec\ l\ t) =$$

$$\lambda l : \text{prin}. \lambda t : \text{Type}. \lambda d : t.$$

$$mkSec\ t (\lambda key : \text{pf}(l \text{ says } \text{Reveal}). d)$$

$$\text{Bind} : (l : \text{prin}) \rightarrow (t : \text{Type}) \rightarrow (s : \text{Type}) \rightarrow (d : Sec\ l\ t)$$

$$\rightarrow (t \rightarrow Sec\ l\ s) \rightarrow Sec\ l\ s$$

$$= \lambda l : \text{prin}. \lambda t : \text{Type}. \lambda s : \text{Type}. \lambda d : Sec\ l\ t. \lambda f : t \rightarrow Sec\ l\ s.$$

$$\text{match } d \text{ with } \{$$

$$| mkSec \rightarrow$$

$$\lambda dt : (\text{pf}(l \text{ says } \text{Reveal}) \rightarrow t).$$

$$mkSec\ l\ s (\lambda key : \text{pf}(l \text{ says } \text{Reveal}).$$

$$\text{match } (f(dt\ key)) \text{ with} \{$$

$$| mkSec \rightarrow$$

$$\lambda ds : (\text{pf}(l \text{ says } \text{Reveal}) \rightarrow s). ds\ key$$

$$\} \}$$



We can treat secrets protected at level  $L$  as if they are protected at level  $H$ , since there are more information-flow restrictions on data protected at level  $H$  than at level  $L$ . We define a function  $LtoH$  that takes an expression of type  $Sec\ L\ t$  and return an expression of the type  $Sec\ H\ t$ .

```

1  $LtoH : (t : \mathbf{Type}) \rightarrow Sec\ L\ t \rightarrow Sec\ H\ t =$ 
2  $\lambda t : \mathbf{Type}. \lambda d : Sec\ L\ t.$ 
3  $mkSec\ H\ t$ 
4  $(\lambda phk : \mathbf{pf}(H\ \mathbf{says}\ Reveal).$ 
5  $\mathbf{match}\ d\ \mathbf{with}\ \{$ 
6  $\quad | mkSec\ \rightarrow$ 
7  $\quad \lambda dl : (\mathbf{pf}(L\ \mathbf{says}\ Reveal) \rightarrow t).$ 
8  $\quad dl(\mathbf{bind}\ hk : (H\ \mathbf{says}\ Reveal) = phk\ \mathbf{in}$ 
9  $\quad \mathbf{bind}\ del : (H\ \mathbf{says}\ Reveal \rightarrow Reveal) = L2H$ 
10  $\quad \mathbf{in}\ \mathbf{return}\ (\mathbf{return}\ L\ (del\ hk)))$ 
11  $\quad \})$ 

```

The body of  $LtoH$  changes the capability guarding the secret in  $d$  from  $\mathbf{pf}(L\ \mathbf{says}\ Reveal)$  to  $\mathbf{pf}(H\ \mathbf{says}\ Reveal)$ . We start by using  $mkSec$  to construct a term protected at  $H$ . On line 4, the variable  $phk$  is the new capability associated with the secret data in  $d$ . The pattern-matching expression between line 5 and 11 constructs a term of type  $t$  by revealing the secret in  $d$ . To do so, we need a capability of the type  $\mathbf{pf}(L\ \mathbf{says}\ Reveal)$ . We obtain such capability between line 8 and 10 by using  $L2H$ , which is the delegation from  $L$  to  $H$  and variable  $phk$ , which is the capability to access secrets protected at  $H$  (defined on line 4). The bind expression between line 8 and 10 is the same as  $lk$  we defined earlier in this section. The secret of type  $t$  hidden in  $d$  is revealed by applying  $dl$  (line 7) to the capability of type  $\mathbf{pf}(L\ \mathbf{says}\ Reveal)$ , constructed at line 8–10.

## 5.2 Noninterference

The encoding in Section 5.1 also has the noninterference property. Intuitively, by the noninterference properties of the authorization logic, we cannot prove  $H\ \mathbf{says}\ Reveal$  from  $L\ \mathbf{says}\ (H\ \mathbf{says}\ Reveal \rightarrow Reveal)$  and  $L\ \mathbf{says}\ Reveal$ . Therefore, when constructing computations protected at the security level  $L$ , we cannot use any data protected at level  $H$ , which are guarded by a capability of type  $\mathbf{pf}(H\ \mathbf{says}\ Reveal)$ .

Our proofs rely on the strong normalization results on AURA's authorization logic (Theorem 4). The idea is that any AURA term of type  $H\ \mathbf{says}\ Reveal$  can be normalized using proof reduction rules to a normal form, and we prove that no normal form has type  $H\ \mathbf{says}\ Reveal$ .

We define the normal forms for proofs below.

Normal Forms

$$\begin{aligned}
nf \quad ::= & \lambda x : t_1. nf \mid c\ nf_1 \cdots nf_n \mid \mathbf{return}_p\ nf \\
& \mid \mathbf{return}_s\ nf_1\ nf_2 \mid \mathbf{sign}(\mathbf{self}, nf) \\
& \mid \mathbf{bind}_s\ x = nf_e\ \mathbf{in}\ nf \mid \mathbf{self} \\
& \mid \mathbf{bind}_s\ x = \mathbf{sign}(\mathbf{self}, nf)\ \mathbf{in}\ nf \mid nf_e \\
& \mid \mathbf{Kind} \mid \mathbf{Type} \mid \mathbf{Prop} \mid \mathbf{prin} \mid \mathbf{pf}\ t \\
& \mid (x : t_1) \rightarrow t_2
\end{aligned}$$

Elimination Normal Forms

$$nf_e \quad ::= x \mid nf_e\ nf \mid con$$

The last two lines of the definition of  $nf$  are types. AURA has no reduction rules at the type level, so all the types are in normal form. The two  $\mathbf{bind}_s$  expressions are stuck computations. Other stuck computations, such as  $x\ y$ , that are not  $\mathbf{bind}_s$  expressions, are denoted by  $nf_e$ . We make a distinction between stuck computations that are  $\mathbf{bind}_s$  expressions and those are not because we have a special commuting reduction rule on terms of the form  $\mathbf{bind}_s\ x = (\mathbf{bind}_s\ y = t_1\ \mathbf{in}\ t_2)\ \mathbf{in}\ t_3$  see [17] for details .

The constants denoted by  $con$  include principals such as  $H$  and  $L$  defined for the lattice. We treat  $L2H$  that defines the partial order between  $L$  and  $H$  as a constant as well. This is because ordinary programmers cannot get hold of either  $H$  or  $L$ 's private key, so the normal form of the programmer's code cannot include expressions of the form of  $\mathbf{sign}(H, P)$  or  $\mathbf{sign}(L, P)$ . When treating  $L2H$  as an opaque constant, the definitions of  $nf$  and  $nf_e$  above generate the same normal form for programmers' code that makes use of  $L2H$  as a constant and that treats  $L2H$  as  $\mathbf{sign}(L, H\ \mathbf{says}\ Reveal \rightarrow Reveal)$ .

We prove the following lemma, which is analogous to Lemma Secret (Lemma 13). This lemma assures us that we cannot construct a term witnessing  $H\ \mathbf{says}\ Reveal$ , even if we assume  $L$  can make arbitrary assertions.

**Lemma 15.** *Secret H key*

$\mathit{if} \forall con \in dom(S), S(x) = \mathbf{prin}$ , or  $S(x) = L\ \mathbf{says}\ t$ ,

- $\mathcal{E} :: S; \cdot \vdash nf : t$  then  $t \neq H\ \mathbf{says}\ Reveal$ , and  $t \neq Reveal$
- $\mathcal{E} :: S; \cdot \vdash nf_e : t$  then  $t = L\ \mathbf{says}\ P$  or  $t = \mathbf{prin}$

**Proof (sketch):** By mutual induction on derivation  $\mathcal{E}$ .  $\square$

The signature we care about is  $SS$ , which contains the definition of data type  $Bool$ ,  $Sec$ , assertion  $Reveal$ , and constants representing the two-point security lattice.

$$\begin{aligned}
SS = & \mathbf{data}\ Bool : \mathbf{Type} \dots, \\
& \mathbf{assert}\ Reveal : \mathbf{Prop}, \\
& \mathbf{data}\ Sec : \mathbf{prin} \rightarrow \mathbf{Type} \rightarrow \mathbf{Type} \dots, \\
& \mathbf{const}\ H : \mathbf{prin}, \mathbf{const}\ L : \mathbf{prin}, \\
& \mathbf{const}\ L2H : L\ \mathbf{says}\ (H\ \mathbf{says}\ Reveal \rightarrow Reveal)
\end{aligned}$$

As a corollary of Lemma 15, we can prove that we cannot construct a term of type  $\mathbf{pf}(H\ \mathbf{says}\ Reveal)$  from the lattice definitions and  $L\ \mathbf{says}\ Reveal$ .

**Lemma 16** (H Secret).

$\nexists v, val(v)$  and  $SS, \mathbf{const}\ LK : L\ \mathbf{says}\ Reveal; \cdot \vdash v : HKey$ .

**Proof (sketch):** Using the strong normalization result and Lemma 15.  $\square$

We use the following macros for the rest of this section.

$$\begin{aligned} LKey &= \mathbf{pf} (L \mathbf{says} \mathit{Reveal}) & SecLB &= Sec L Bool \\ HKey &= \mathbf{pf} (H \mathbf{says} \mathit{Reveal}) & SecHB &= Sec H Bool \end{aligned}$$

This noninterference theorem below states that with two different secret inputs protected at security label  $H$ , the output values at level  $L$  are the same. The statement of the noninterference theorem with security lattices becomes more complicated because now we have to state that if two input values are the same for observers at security level  $L$ , then the output values of type  $SecLB$  are the same for observers at security level  $L$ . We indicate the presence of  $L$  observers by including a constant  $LK$  of the type  $(L \mathbf{says} \mathit{Reveal})$  in the signature for type checking the input values  $v_i$ . This is equivalent to saying that the observers at level  $L$  can see any secrets protect at level  $L$ , because  $\mathbf{return}_p LK$  has type  $\mathbf{pf} (L \mathbf{says} \mathit{Reveal})$ . We cannot simply use the syntactic equality to state the equality of two values of type  $SecLB$ , because those values contain sub-terms that are functions. We need to specify that those functions evaluate to the same values when applied to the same arguments.

**Theorem 17** (Noninterference).

If  $SS; x : SecHB \vdash e : SecLB$  and given any two values  $v_1, v_2$  such that  $SS, \mathbf{const} LK : L \mathbf{says} \mathit{Reveal}; \cdot \vdash v_i : SecHB$  and  $e\{v_1/x\} \mapsto^* w_i$  where  $w_1, w_2$  are values, then  $w_i = mkSec L Bool f_i$  and if  $(f_i (\mathbf{return}_p LK)) \mapsto^* u_i$  where  $u_i$  are values, then  $u_1 = u_2$ .

The structure of the proof is very similar to the one shown in Figure 2. Due to space constraints, we omit the details. We explain two points in the proof: where Lemma Secret (Lemma 16) is used, and why the outputs are compared in the presence of  $LK$ .

In the proof, we know by the Canonical Forms Lemma that  $v_i = mkSec H Bool g_i$ . Lemma Secret (Lemma 16) allows us to construct a well-typed pair  $\langle g_1 \mid g_2 \rangle$  in AURA-PAIR. This means that  $g_1$  and  $g_2$  are secrets given the current context and, therefore, could be put into a pair.

In the end, we know that  $e[\hat{v}/x] \mapsto_p^* s$  and  $w_i = \lfloor s \rfloor_i$ . By canonical forms, we know that  $s = mkSec L Bool q$ , and  $SS, \mathbf{const} LK : L \mathbf{says} \mathit{Reveal}; \cdot \vdash^p q : LKey \rightarrow Bool$ . With  $LK$  in the signature, the canonical form will tell us that  $q$  has to be a lambda abstraction. Without  $LK$ ,  $q$  itself could be a pair of functions. For observers at level  $L$ ,  $q$  could not have been a pair because  $q$  does not contain information of higher secrecy than  $L$ .

## 6. Declassification

Information-flow polices that do not allow any information flow from high security to low security are typically too restrictive for practical use. To build useful systems, we often find it necessary to leak some amount of secret information.

In this section, we explore through several examples the design space for using access-control policies to specify declassification policies in AURA.

### 6.1 Simple Declassification Policies

**Escape hatches** We can define a declassify operation similar to escape hatches [25]. The *declassify* function will reveal a secret protected at level  $H$ . If we assume that *declassify* is running under the authority  $H$ , the term  $\mathbf{say} \mathit{Reveal}$  is a capability for revealing the secret, and we can implement *declassify* in AURA as follows.

$$\begin{aligned} declassify &: Sec H t \rightarrow Maybe t \\ &= \lambda d : Sec H t. \\ &\quad \mathbf{match} d \mathbf{with} \\ &\quad \mid mkSec \rightarrow \\ &\quad \quad \lambda dt : \mathbf{pf}(H \mathbf{says} \mathit{Reveal}) \rightarrow t. \\ &\quad \quad \mathbf{if} H = \mathbf{self} \\ &\quad \quad \mathbf{then} Just (dt \langle (\mathbf{say} \mathit{Reveal}) : \mathbf{pf} (H \mathbf{says} \mathit{Reveal}) \rangle) \\ &\quad \quad \mathbf{else} Nothing \end{aligned}$$

Since  $H$  is the same as  $\mathbf{self}$ , in the true branch of the if expression we can use the explicit type cast to give the expression  $(\mathbf{say} \mathit{Reveal})$  the type  $\mathbf{pf} (H \mathbf{says} \mathit{Reveal})$ , which is used to reveal the secret hidden in  $d$ .

**When** More interestingly, we can use access-control policies to specify *when* information leaks are allowed. We can provide the following generic declassification interface:

$$declassify : Sec H t \rightarrow \mathbf{pf} (H \mathbf{says} \mathit{Reveal}) \rightarrow t$$

*declassify* takes two arguments: a secret protected at level  $H$  and the capability to reveal secrets protected at level  $H$ . *declassify* returns the secret hidden in the first argument.

We can define access-control policies that can be used to construct a proof of  $\mathbf{pf} (H \mathbf{says} \mathit{Reveal})$ . For instance,  $pol_1$ , below, specifies that if payment has been made, then the secret can be released. We use *Cashier* to represent the principal that controls the payment process. *Paid* is an assertion defined in the same way as *Reveal*.

$$pol_1 : H \mathbf{says} (\mathit{Cashier} \mathbf{says} \mathit{Paid} \rightarrow \mathit{Reveal})$$

We can further define policies to specify when *Cashier* will affirm that payment has been made. For example, the following policy states that if *PNCBank* affirms that deposit has been made to account (Num), then *Cashier* will agree that payment has arrived.

$$pol_c : \mathit{Cashier} \mathbf{says} (\mathit{PNCBank} \mathbf{says} (\mathit{Deposited} \mathit{Num}) \rightarrow \mathit{Paid})$$

Alternatively, we could give the declassification interface the following more informative type.

$$declassify : Sec H t \rightarrow \mathbf{pf} (\mathit{Cashier} \mathbf{says} \mathit{Paid}) \rightarrow t$$

**Who** We can also specify to whom information is released. In the following example, we allow privileged users to access secret information. We use *Sys* to denote the principal System, who is in charge of deciding who are privileged principals. The predicate *Privileged p* means that principal *p* is a privileged principal and is defined below.

**assert** *Privileged* : **prin** → **Prop**

Policy  $pol_2$  allows the capability to access secrets protected at level *H* to be obtained by constructing a proof of *Sys says (Privileged p)*.

$pol_2 : H \text{ says } (Sys \text{ says } (Privileged p)) \rightarrow Reveal$

The declassify interface that allows privileged principals to access secrets protected at *H* is shown below.

$declassify : (p : \text{prin}) \rightarrow Sec H t$   
 $\rightarrow Sys \text{ says } (Privileged p) \rightarrow t$

The first argument of declassify is the principal to whom the information is released. The second argument is a secret protected at *H*. The third argument is a proof that *Sys* believes that *p* is a privileged user. The body of declassify uses  $pol_2$  and returns the secret.

## 6.2 More Elaborate Policies

**Refinement of secrets** We can refine our encoding so that instead of using a single capability for all secrets, we can define different classes of secrets guarded by different capabilities. For example, the salaries of employees in the Engineer College are secrets. However, there are several different kinds of employees. We can use different capabilities to guard graduate students' salaries (**pf** (*H says GradSalary*)), postdocs' salaries (**pf** (*H says PostDocSalary*)) and professors' salaries (**pf** (*H says ProfSalary*)).

Here, *GradSalary*, *PostDocSalary* and *ProfSalary* are all assertions defined in the same way as *Reveal*. Now the security types need to indicate the class of secrets as well. For instance, instead of *Sec H t*, we use *Sec H Reveal t*.

We can write policies to declassify certain kinds of secrets. For example, the following statements declare that the Dean believes that postdocs and grad students are temporary employees. Predicate (*tmpE P*) means that *P* is a proposition used for guarding the salary info of temporary employees.

$s1 : Dean \text{ says } (tmpE PostDocSalary)$   
 $s2 : Dean \text{ says } (tmpE GradSalary)$

The following declassification interface downgrades all information about temporary employees.

$declassify : (R : \text{Prop}) \rightarrow Sec H R t$   
 $\rightarrow Dean \text{ says } (tmpE R) \rightarrow t$

**Nonces** One problem with the declassification interfaces shown so far is that a replay attack could cause unwanted information leaks. For instance, in the example where a

proof of **pf** (*Cashier says Paid*) is required for the release of secrets, an attacker can use an old proof to learn all the secrets protected by **pf** (*H says Reveal*).

A standard way to prevent such replay attacks is to include a fresh nonce in the proofs, thereby preventing old proofs from being re-used. We can refine our encoding and include a nonce in the declassification interface.

**data** *Nonce* : **Type**{  
 | *n1* : *Nonce*  
 ...  
 }  
**assert** *Paid* : *Nat* → **Prop**  
 $declassify : (n : Nonce) \rightarrow Sec H t$   
 $\rightarrow \text{pf } (Cashier \text{ says } (Paid n))$   
 $\rightarrow Maybe t$

A trusted nonce-generation function becomes part of the declassification interface. It produces a fresh nonce *m* for each declassification. In the body of the *declassify* function, the nonce a user passes in is checked against the stored current nonce for equality. Only when they are equal, the proof passed in by the user can be casted to a proof of **pf** (*Cashier says (Paid m)*), where *m* is the value of the stored current nonce. Therefore, an old proof with an expired nonce will not reveal the secret. We are in effect implementing a release-once policy.

However, the stored current nonce also becomes part of the trusted declassification interface. This means that implementing this version of the *declassify* function requires mutable state, which is not supported by AURA at this time.

## 6.3 Discussion

We have not studied the formal properties of these declassification policies. We suspect that the noninterference property of the authorization logic will allow us to express properties such as: "if a leak happens then certain principals must have made certain assertions". This kind of property would be useful for auditing purposes.

AURA does not support a module system or key management for run-time keys. Each process is associated with one run-time authority. This made it difficult to specify that *declassify* has to be run in trusted space and cannot be exploited by attackers. Furthermore, AURA's lack of support for mutable state prevents us from implementing declassification policies involving nonces, as shown in the example. If we were to add state to AURA, our encoding would need to be refined to consider possible information leaks from state changes. We speculate that techniques from prior work, such as those used for building a library for light-weight information-flow security in Haskell [24] would apply. We leave these investigations for future work.

## 7. Related Work

**Information flow type systems** There has been much work on using language-based approaches to protect the confiden-

tiality of information (Cf. [30, 15, 21, 32, 28]). Most of these works use security-label indexed types to indicate the security level of the data, and type systems enforce information-flow policies. Abadi *et al.* pointed out that information-flow analysis is a dependency analysis, and the noninterference property holds in the dependency core calculus (DCC) [4]. In DCC, security types are treated as security-label indexed monads. Abadi later showed that DCC can be used as a calculus for access control [2]. When DCC's monads are interpreted as principal-indexed types expressing principal's beliefs, DCC is isomorphic to an authorization logic. AURA contains a constructive authorization logic based on DCC in its type system [29, 16]. However, the principal-indexed monads in AURA cannot be used directly as security types. For example, we cannot use  $H \text{ says } int$  as the type for an integer protected at level  $H$ . This is because AURA has two separate universes: one for logical proofs and propositions, which is pure; and the other for computations which may have effects such as non-termination. The separation is necessary for maintaining the soundness of AURA's authorization logic. The **says** monads are logical assertions, whereas the security types are the types for data and computations. This paper provides an encoding of security types for data using the principal-indexed types.

Another approach to enforcing information-flow policies is to encode security types as libraries for existing functional languages, notably Haskell. Li *et al.* showed how to enforce information-flow policies in Haskell by encoding information-flow types using the arrow combinator [19]. A light-weight encoding of information-flow types using Haskell's type classes and abstract data types is later presented by Russo *et al.* [24]. Both of these encodings rely on Haskell's type classes and abstract data types to ensure that the information-flow policies are enforced. Our encoding relies on the noninterference properties of the authorization logic to enforce information-flow policies. One significant technical contribution of our work is that we proved a noninterference theorem for our encoding using the squared semantics approach [23], and that all aspects of our proofs related to the squared semantics are mechanized in Coq. To our knowledge, the noninterference proof in Russo's work [24] is done for an abstraction of the implementation. There is no formal proof about whether the abstraction faithfully models the implementation. Our proof of noninterference is done for the implementation itself, which had not been done. We acknowledge that Haskell is significantly more complicated than AURA, and our encoding does not consider side effects such as mutable references or IO, because AURA does not have these features. Another important contribution of our work is that we studied aspects of using access-control policies to declassify information.

**Noninterference proofs of authorization logics** The noninterference theorems of our encoding rely on the noninterference properties of AURA's authorization logic. We

need to demonstrate that there is no value of the type  $\text{pf } (H \text{ says } \textit{Reveal})$  (stated in Lemmas 13 and 16), which is a form of the noninterference property of the authorization logic in AURA. The first noninterference proof of a constructive authorization logic was done by Garg [14]. In Garg's proofs, the sub-formula properties of a cut-free sequent calculus are used to identify the assumptions that contribute to the conclusion. We could not use Garg's noninterference results directly because Garg's logic has different rules than ours, and it is first-order, while AURA's logic is second-order. In our proof, we use the strong-normalization results of the authorization logic in AURA. We examine all the possible normal forms of proofs that could be constructed using existing assumptions for encoding the lattice, and conclude that certain proofs are not possible. The statements of our noninterference theorem of the authorization logic (Lemmas 13 and 16) are not as general as Garg's. Encoding different lattices need different formulas, and we need to prove a lemma similar to Lemma 16 for each lattice. However, the techniques of our proofs are general enough for constructing proofs for other lattices. What's interesting in our work is that we demonstrate how to apply the noninterference property of an authorization logic to encoding information-flow types that have noninterference property.

Abadi also proved noninterference for CDD [3], a cut down version of DCC. However, in CDD, the lattice of principals does not correspond to delegation between them. We use explicit delegation between principals to encode lattices. Consequently, the noninterference proofs of CDD are not really applicable in our setting.

## References

- [1] M. Abadi. Logic in access control. In *Proceedings of the 18th Symposium on Logic in Computer Science (LICS)*, June 2003.
- [2] M. Abadi. Access control in a core calculus of dependency. In *Proceedings of the 11th ACM SIGPLAN International Conference on Functional Programming, ICFP*, 2006.
- [3] M. Abadi. Access control in a core calculus of dependency. *Computation, Meaning, and Logic: Articles dedicated to Gordon Plotkin ENTCS*, 172:5–31, April 2007.
- [4] M. Abadi, A. Banerjee, N. Heintze, and J. Riecke. A core calculus of dependency. In *Proc. 26th ACM Symp. on Principles of Programming Languages (POPL)*, Jan. 1999.
- [5] M. Abadi, M. Burrows, B. W. Lampson, and G. D. Plotkin. A calculus for access control in distributed systems. *Transactions on Programming Languages and Systems*, 15(4):706–734, Sept. 1993.
- [6] A. Askarov and A. Sabelfeld. Gradual release: Unifying declassification, encryption and key release policies. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy*, 2007.
- [7] A. Askarov and A. Sabelfeld. Localized delimited release: combining the what and where dimensions of information release. In *PLAS '07: Proceedings of the 2007 workshop on Programming languages and analysis for security*, 2007.

- [8] A. Banerjee, D. A. Naumann, and S. Rosenberg. Expressive declassification policies and modular static enforcement. *Security and Privacy, IEEE Symposium on*, 2008.
- [9] H. P. Barendregt. Lambda calculi with types. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 2, pages 117–309. Clarendon Press, Oxford, 1992.
- [10] N. Broberg and D. Sands. Flow locks: Towards a core calculus for dynamic flow policies. In *In ESOP 2006: the 15th European Symposium on Programming*, 2006.
- [11] J. Cederquist, R. Corin, M. Dekker, S. Etalle, and J. den Hartog. An audit logic for accountability. In *The Proceedings of the 6th IEEE International Workshop on Policies for Distributed Systems and Networks*, 2005.
- [12] The Coq Development Team. *The Coq Proof Assistant Reference Manual version 8.1*, 2007. Available from <http://coq.inria.fr/>.
- [13] H. DeYoung, D. Garg, and F. Pfenning. An authorization logic with explicit time. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF-21)*, June 2008.
- [14] D. Garg and F. Pfenning. Non-interference in constructive authorization logic. In *CSFW '06: Proceedings of the 19th IEEE workshop on Computer Security Foundations*, 2006.
- [15] N. Heintze and J. G. Riecke. The SLam calculus: Programming with secrecy and integrity. In *Proc. 25th ACM Symp. on Principles of Programming Languages (POPL)*, Jan. 1998.
- [16] L. Jia, J. A. Vaughan, K. Mazurak, J. Zhao, L. Zarko, J. Schorr, and S. Zdancewic. Aura: A programming language for authorization and audit,. In *13th ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2008.
- [17] L. Jia and S. Zdancewic. Encoding information flow in AURA, technical appendix. Technical Report MS-CIS-09-08, University of Pennsylvania, 2009. Available at <http://www.seas.upenn.edu/~liminjia/research/papers/aura-flow-tr.pdf>.
- [18] P. Li and S. Zdancewic. Downgrading Policies and Relaxed Noninterference. In *Proc. ACM Symp. on Principles of Programming Languages (POPL)*, pages 158 – 170, 2005.
- [19] P. Li and S. Zdancewic. Encoding information flow in Haskell. In *CSFW '06: Proceedings of the 19th IEEE workshop on Computer Security Foundations*, 2006.
- [20] H. Mantel and A. Reinhard. Controlling the what and where of declassification in language-based security. In *In ESOP 2007: the 16th European Symposium on Programming*, 2007.
- [21] A. C. Myers. JFlow: Practical mostly-static information flow control. In *Proc. 26th ACM Symp. on Principles of Programming Languages (POPL)*, Jan. 1999.
- [22] A. C. Myers, A. Sabelfeld, and S. Zdancewic. Enforcing robust declassification and qualified robustness. *Journal of Computer Security*, 14(2):157–196, 2006.
- [23] F. Pottier and V. Simonet. Information flow inference for ML. *ACM Trans. Program. Lang. Syst.*, 25(1):117–158, 2003.
- [24] A. Russo, K. Claessen, and J. Hughes. A library for lightweight information-flow security in Haskell. In *Proceedings of the first ACM SIGPLAN symposium on Haskell*, 2008.
- [25] A. Sabelfeld and A. C. Myers. A Model for Delimited Release. In *International Symposium on Software Security*, 2003.
- [26] A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, Jan. 2003.
- [27] A. Sabelfeld and D. Sands. Dimensions and principles of declassification. In *CSFW '05: Proceedings of the 18th IEEE workshop on Computer Security Foundations*, 2005.
- [28] V. Simonet. Flow Caml in a nutshell. In *Proceedings of the first APPSEM-II workshop*, Mar. 2003.
- [29] J. A. Vaughan, L. Jia, K. Mazurak, and S. Zdancewic. Evidence-based audit. In *Proc. of the IEEE Computer Security Foundations Symposium*, 2008.
- [30] D. Volpano, G. Smith, and C. Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3):167–187, 1996.
- [31] S. Zdancewic and A. C. Myers. Robust declassification. In *Proc. of 14th IEEE Computer Security Foundations Workshop*, pages 15–23, Cape Breton, Canada, June 2001.
- [32] S. Zdancewic and A. C. Myers. Secure information flow and CPS. In *Proc. of the 10th European Symposium on Programming*, Apr. 2001.