

# An Algebraic Approach to Asymmetric Delegation and Polymorphic Label Inference (Technical Report)

Silei Ren<sup>1</sup>[0000-0002-4182-0211], Coşku Acay<sup>2</sup>[0000-0002-0487-1167], and  
Andrew C. Myers<sup>1</sup>[0000-0001-5819-7588]

<sup>1</sup> Cornell University, Ithaca NY 14850, USA  
{sr2262, andru.cs}@cornell.edu

<sup>2</sup> Observe, Inc., San Mateo CA 94402, USA  
coskuacay@gmail.com

**Abstract.** Language-based information flow control (IFC) enables reasoning about and enforcing security policies in decentralized applications. While information flow properties are relatively extensional and compositional, designing expressive systems that enforce such properties remains challenging. In particular, it can be difficult to use IFC labels to model certain security assumptions, such as semi-honest agents.

Motivated by these modeling limitations, we study the algebraic semantics of lattice-based IFC label models, and propose a semantic framework that allows formalizing asymmetric delegation, which is partial delegation of confidentiality or integrity. Our framework supports downgrading of information and ensures their safety through nonmalleable information flow (NMIF).

To demonstrate the practicality of our framework, we design and implement a novel algorithm that statically checks NMIF and a label inference procedure that efficiently supports bounded label polymorphism, allowing users to write code generic with respect to labels.

## 1 Introduction

*Information Flow Control (IFC)* [23, 37] is a well-established approach for enforcing security in decentralized applications. Using *labels*, IFC systems specify fine-grained policies on information flow that can be fully or partly enforced through compile-time analysis. These policies articulate the confidentiality and integrity goals of IFC systems, which are *security conditions* [23]: hyperproperties [12] that constrain the set of system behaviors.

The most prominent security condition in IFC systems is *noninterference* [18], but it is too restrictive in practice. A major challenge for adopting language-based IFC is providing developers with expressive yet intuitive ways to specify their intended security policies. To capture more nuanced security policies, the expressiveness of IFC systems is enhanced by *downgrading mechanisms* such as

*declassification* of confidential information and *endorsement* of untrusted information. Misuse of these mechanisms is further mitigated by enforcing *nonmalleable information flow* [9] (NMIF), a security condition controlling downgrading.

*Delegation* is another common mechanism for specifying security. Delegation allows one principal to grant (delegate) power to another, expressing that the first principal trusts the second. Delegation can compactly represent important aspects of the system’s security policy: when there is delegation between two principals, ensuring the delegator’s security also necessitates enforcing the delegatee’s security. Delegation is commonly supported not only in information flow control systems [4, 5, 30, 34], but also in a wide range of enforcement mechanisms, including access control [7, 13, 38] (where delegation is often referred to as a *principal* or *role hierarchy*), authorization logics [1, 2, 4, 20, 21], and capability systems [25, 27].

The expressive power of delegation can be increased through what we call *asymmetric delegation*: fine-grained delegation of either *confidentiality* or *integrity*. Intuitively, when a principal Alice delegates her confidentiality to another principal Bob, she allows Bob to observe all information visible to her. When Alice delegates her integrity to Bob, she trusts that all information accepted by Bob has not been maliciously modified. With asymmetric delegation, we can model security settings like the semi-honest trust assumption in cryptographic applications and the security setting of blockchains. In the semi-honest setting, principals trust each other to follow the protocol (trust each other with integrity), but do not trust each other with their secrets (but not confidentiality). In the blockchain setting, principals do not trust each other to follow protocols, but all information is public: they effectively trust each other with respect to confidentiality.

While asymmetric delegation increases the expressive power of IFC systems, its precise role—particularly in the presence of downgrading—remains poorly understood. We address this gap by presenting a general and expressive semantic framework for IFC labels that formalizes both asymmetric delegation and its interaction with downgrading. Although prior work [5, 43] develops IFC systems that support certain forms of asymmetric delegation, these systems lack sound and complete NMIF enforcement. Building on our framework, we develop algorithms for verifying the associated semantic security conditions.

Experience with language-based security highlights the importance of IFC label inference to reduce the burden on programmers [3, 34]. In addition, allowing programmers to write code that is generic with respect to labels enhances modularity and code reuse. We support such generic programming through an efficient label inference procedure that supports bounded label polymorphism.

To evaluate our approach, we update the label model of the Viaduct compiler [3] and extend its static information flow analysis. Our implementation features a more concise and modular syntax for specifying trust assumptions, as well as a more efficient label inference procedure.

The rest of the paper is structured as follows:

- Section 2 motivates asymmetric delegation using a semi-honest secure multi-party computation (MPC) program.

- Section 3 studies the effects of asymmetric delegation on security conditions using a novel semantic framework.
- Section 4 presents algorithms that statically enforce the security conditions.
- Section 5 introduces an inference procedure supporting bounded label polymorphism.

```

1 host Alice : {A  $\wedge$  B←}
2 host Bob   : {B  $\wedge$  A←}
3
4 val a: {A  $\wedge$  B←} = Alice.input
5 val b: {B  $\wedge$  A←} = Bob.input
6 val w: {A  $\wedge$  B} = a > b
7
8 Alice.output(
9   declassify w to {A  $\wedge$  B←})
10 Bob.output(
11   declassify w to {B  $\wedge$  A←})

```

**Fig. 1.** Yao’s Millionaires’ problem in Viaduct [3]. The programmer must manually assign labels to hosts.

```

1 host Alice, Bob
2 assume Alice = Bob for integrity
3
4 val a: {Alice} = Alice.input
5 val b: {Bob}   = Bob.input
6 val w: {Alice  $\sqcup$  Bob} = a > b
7
8 Alice.output(
9   declassify w to {Alice})
10 Bob.output(
11   declassify w to {Bob})

```

**Fig. 2.** Yao’s Millionaires’ problem implemented with delegation. Alice  $\sqcup$  Bob is shorthand for  $\langle$ Alice  $\wedge$  Bob, Alice  $\vee$  Bob $\rangle$ .

## 2 A Case for Delegation

### 2.1 Semi-Honest Attackers in Cryptography

Asymmetric delegation can capture a wide variety of security settings. Already mentioned is the semi-honest threat model, widely studied in the cryptography literature [44]. In this model, principals correctly follow the protocol, but attempt to improperly learn other principals’ secrets. Modeling the semi-honest setting in IFC systems remains a challenge. We first give an example of modeling semi-honest security in Viaduct [3], a state-of-the-art compiler that translates information flow policies to cryptographic protocols. We then illustrates how delegation improves usability and modularity.

Consider Yao’s well-known Millionaires’ Problem [44], where Alice and Bob wish to compare their wealth without revealing actual numbers. Figure 1 shows a Viaduct implementation. Lines 1 and 2 declare the hosts Alice and Bob and assign them information flow labels that capture the security assumptions. The hosts are assigned different and incomparable confidentiality labels (A for Alice and B for Bob) but the same integrity label (A  $\wedge$  B) to reflect the trust relation in the semi-honest model. Lines 4 and 5 gather input from the hosts; input from a host has the same label as that host. Line 6 stores the result of the comparison in w, which has a label following standard IFC rules: the result of a computation is more

secret and less trusted than all of its inputs. Specifically,  $w$  has a confidentiality of  $A \wedge B$  since it is derived using secret data from both hosts, and has an integrity of  $A \wedge B$  since that is the integrity of both inputs. Finally, lines 9 and 11 output  $w$  to Alice and Bob, respectively. Note that sending  $w$  to Alice leaks information about Bob’s secret data ( $b$ ), which violates noninterference. Viaduct requires an explicit **declassify** statement to indicate that information leakage is intentional.

## 2.2 Modeling Security with Delegation

Viaduct models security by encoding trust into labels, but this approach has problems. First, programmers must encode security assumptions by carefully crafting host labels, which becomes tricky in large systems with many assumptions. Second, this encoding pollutes the entire program. In fig. 1, every label annotation must acknowledge the semi-honest assumption by carrying around additional integrity (i.e.,  $A^\leftarrow$  or  $B^\leftarrow$ ). And third, the encoding breaks modularity. For example, to add a new host Chuck to the program, we would need to edit every label annotation to carry an extra integrity component of  $C^\leftarrow$ , requiring changes throughout the program even though Chuck is not involved in this portion of the computation. Delegation addresses all of these issues.

Figure 2 implements Yao’s Millionaires’ Problem using delegation. Hosts are no longer assigned cryptic information flow labels; instead, line 2 directly states the security assumption: Alice and Bob trust each other for integrity. Variables have intuitive labels that do not need to repeat the semi-honest security assumption: input from Alice has label Alice. Finally, adding a new host Chuck requires no edits to existing code; we only need to add the following lines:<sup>3</sup>

```

1 host Chuck
2 assume Alice = Chuck for integrity

```

## 2.3 Nonmalleable Information Flow

Downgrading statements (**declassify** and **endorse**) deliberately violate noninterference, so their unrestricted use poses a threat to security. Prior work [33, 46] identifies cases where the attacker can exploit downgrading to gain undue influence over the execution, and proposes *robust declassification* and *transparent endorsement* to limit such cases.

Robust declassification requires that untrusted data is not declassified, and transparent endorsement requires that secret data is not endorsed. NMIF combines these two restrictions, which are key to enabling the Viaduct compiler to securely instantiate programs with cryptography [3].

Here, “secret” and “trusted” are relative to a given attacker, and NMIF must hold for all attackers. In practice, the program cannot be type-checked

<sup>3</sup> In fact, we could even support separate compilation as the program need not be type-checked again: a program considered secure with fewer assumptions is secure with more assumptions.

separately for every possible attacker, so a conservative condition is enforced: downgraded data must be at least as trusted as it is secret. For our example program, this condition means  $w$  must have integrity stronger than or equal to its confidentiality. This condition is immediate in Viaduct since  $w$  has label  $\langle \text{Alice} \wedge \text{Bob}, \text{Alice} \wedge \text{Bob} \rangle$  in fig. 1. On the other hand, the same variable  $w$  in fig. 2 has label  $\langle \text{Alice} \wedge \text{Bob}, \text{Alice} \vee \text{Bob} \rangle$ , which seemingly has weaker integrity than confidentiality (logically,  $\text{Alice} \vee \text{Bob}$  does *not* imply  $\text{Alice} \wedge \text{Bob}$ ). However,  $\text{Alice} = \text{Bob}$  for integrity, so this label is equivalent to  $\langle \text{Alice} \wedge \text{Bob}, \text{Alice} \wedge \text{Bob} \rangle$  using the following derivation:

$$\text{Alice} \vee \text{Bob} = \text{Alice} \vee \text{Alice} = \text{Alice} = \text{Alice} \wedge \text{Alice} = \text{Alice} \wedge \text{Bob}$$

Delegation necessitates equational reasoning under assumptions, and NMIF creates an interaction between confidentiality and integrity. The combination of these two features is what makes asymmetric delegation tricky: the cleaner syntax comes at the cost of additional technical complexity. The following sections tame this complexity by developing a semantic framework for labels, and algorithms that follow the semantics.

### 3 Semantic Framework

#### 3.1 The Lattice of Principals

We build our semantic framework upon the *lattice of principals*, used in prior work in authorization logics and information flow systems [3–5, 31, 34, 41, 43].

A principal  $p \in \mathbb{P}$  refers to an entity in decentralized systems that can be either concrete, such as users or server machines, or abstract, such as RBAC roles [38] or quorums [50]. In IFC systems, they are often used as labels to annotate policies on use of information [4]. For example,  $a: \text{Alice}$  in fig. 2 requires the variable  $a$  to only be written by principals that can influence Alice’s data, and to remain secret to principals who cannot observe Alice’s information.

Principals are ordered by authority. When  $q$  delegates trust to  $p$ , we say  $p$  *acts for*  $q$ , written as  $p \Rightarrow q$ . Conjunction (the “and” logic connective) between principals  $p \wedge q$  represents the least combined authority of  $p$  and  $q$ , and disjunction (“or”)  $p \vee q$  represents greatest common authority. The maximal authority  $\perp$  acts for all other authorities, and the minimal authority  $\top$  trusts all other authorities. More authority is associated with elements lower in the lattice:  $\perp \Rightarrow p \wedge q \Rightarrow p \Rightarrow p \vee q \Rightarrow \top$ .<sup>4</sup>

Additionally, we use a *delegation context*, with the form  $\theta = p_1 \Rightarrow q_1, \dots, p_n \Rightarrow q_n$ , to specify delegations that are not implied by the logical structure of the principal lattice. The declaration **assume**  $\text{Alice} = \text{Bob}$  **for integrity** from fig. 2 is an example of a delegation context, specifying both  $\text{Alice} \Rightarrow \text{Bob}$  and  $\text{Bob} \Rightarrow \text{Alice}$ .

<sup>4</sup> It might seem odd to represent maximal authority as  $\perp$  and minimal authority with  $\top$ , since some prior work (e.g., [4]) makes the opposite choice. An intuitive justification: the “false” proposition entails everything, so no real principal can have authority  $\perp$ , but all are trusted with  $\top$ .

Using the delegation context is compatible with much prior work in IFC. For example, the *trust configuration* from FLAM [4], *meta-policies* from the Rx model [43], *interpretation function* from DLM [29], and *authority lattice* from label algebra [28] are all delegation contexts written differently.

The lattice of principals can be interpreted as an authorization logic [1, 17] where each principal is a proposition about authorization policy. As authorization logics are often built upon propositional intuitionistic logics, whose algebraic models are *Heyting algebra* [36], we assume  $\mathbb{P}$  is distributive.<sup>5</sup>

### 3.2 Delegation and Attackers

In the MPC example, a delegation context makes some principals equivalent. In this subsection, we give delegation a precise semantics.

In systems with decentralized trust, all principals see other principals as potential attackers. In the extreme case where no principal trusts another, the attacker with respect to each principal controls all other principals. Formally, we characterize an attacker  $A \in 2^{\mathbb{P}}$  by the set of principals it controls.

To trust a principal is to disregard the case where it is the attacker. Conversely, when a principal is attacker-controlled, so are the principals it acts for. Formally:

**Definition 1 (Consistent Attackers).** *A is consistent with  $\theta$  ( $A \models \theta$ ) when:*

$$\forall p, q \in \mathbb{P} . ((p \Rightarrow q) \vee (q \Rightarrow p) \in \theta) \implies (p \in A \implies q \in A)$$

The set of consistent attackers is an *attacker model*:  $\mathbb{A}_{|\theta} = \{A \in \mathbb{A} \mid A \models \theta\}$ .

The semantic trust levels of principals can be compared based on the set of consistent attackers that control the principals. Formally:

**Definition 2 (Acts-for Semantics).** *p acts for q (written  $\theta \models p \leq q$ ) when:*

$$\forall A \in \mathbb{A}_{|\theta} . p \in A \implies q \in A$$

We use “ $\leq$ ” to denote the semantics of the acts-for relation “ $\Rightarrow$ ”. When viewing principals as authorization propositions, “acts-for” stands for “implies”, and a delegation context is a theory (list of propositions). Each consistent attacker is a consistent interpretation (truth assignment) of the principals and the delegation context, where 1 is assigned to the principals the attacker controls.

In fact, a delegation context determines a *congruence relation*<sup>6</sup> over the lattice of principals, where  $p \equiv_{\theta} q$  is defined as  $(\theta \models p \leq q) \wedge (\theta \models q \leq p)$ . As a result,  $\equiv_{\theta}$  induces a quotient lattice  $\mathbb{P}/\equiv_{\theta}$  where mutually delegating principals are in the same equivalence class. This quotient lattice is precisely the *Lindenbaum algebra* [8] of the theory  $\theta$ : the “smallest” algebraic model  $\mathbb{P}$  where  $\theta$  hold.

**Theorem 1 (Algebraic Model).** *The algebraic model of the principal lattice  $\mathbb{P}$  under delegation context  $\theta$  is the quotient lattice  $\mathbb{P}/\equiv_{\theta}$ .*

<sup>5</sup> A Heyting algebra is a distributive lattice that supports the *relative pseudocomplement* ( $\rightarrow$ ) operation. We do not need the  $\rightarrow$  operator until label inference.

<sup>6</sup> A congruence is an equivalence relation that preserves the lattice structure.

The semantics of consistent attackers make them the *prime filters* of the lattice of principals. This is a result of the Stone’s Representation Theorem of Distributive Lattices [42], which says elements from distributive lattices can be fully characterized by their prime filters.

**Theorem 2 (Attacker Model).**  $\mathbb{A}_{|\theta}$  is the set of prime filters of  $\mathbb{P}/\equiv_{\theta}$ .

Prime filters have intuitive interpretations, which abstracts and generalizes attacker models from prior work [9, 22].  $A \subseteq \mathbb{P}$  is a prime filter when:

- $\top \in A$ : All attackers control the weakest authority;
- $\perp \notin A$ : No attacker controls the strongest authority;
- If  $p \Rightarrow q$  and  $p \in A$ , then  $q \in A$ : Attackers are consistent;
- If  $p \in A$  and  $q \in A$ , then  $p \wedge q \in A$ : An attacker controls the least combined authority of principals it controls;
- If  $p \vee q \in A$ , then either  $p \in A$  or  $q \in A$ : If two principals are not controlled by an attacker, neither is their greatest common authority.

### 3.3 Labels

Information flow systems mainly consider two aspects of security: confidentiality (read authority) and integrity (write authority). To express differing confidentiality and integrity policies, we use *the lattice of labels*, which are pairs of principals  $\mathbb{L} = \mathbb{P} \times \mathbb{P}$ . For a label  $\ell = \langle p, q \rangle$ ,  $p$  represents confidentiality and  $q$  represents integrity. Like principals, each label reflects the authority required for a principal to access information. For example, information labeled  $\langle \perp, \top \rangle$  can be read by no principal but the strongest  $\perp$ , but it can be influenced by any principal.

A label  $\ell$  act for  $\ell'$  when both  $\ell$  acts for  $\ell'$  for both confidentiality and integrity. Similarly, conjunction/disjunction of labels is defined by the conjunction/disjunction of their confidential and integrity. An asymmetric delegation context is a pair of different delegation contexts  $\Theta = \langle \theta_c, \theta_i \rangle$ .

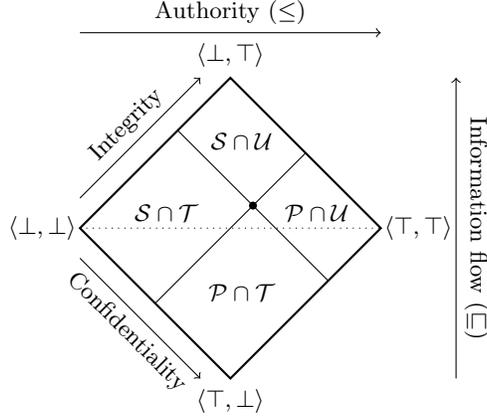
In general, *asymmetric attackers*  $\mathcal{A} \in \mathbf{A} = \mathbb{A} \times \mathbb{A}$  may control different principals for confidentiality and integrity. We write  $\mathcal{A} = \langle C \in \mathbb{A}, I \in \mathbb{A} \rangle$ , where  $C$  represents the principals  $\mathcal{A}$  controls for confidentiality, and  $I$  those for integrity. Consequently, the attacker model  $\mathbf{A}$  is a pair of prime filters.

As visualized in fig. 3, each attacker defines secret ( $\mathcal{S}$ ), public ( $\mathcal{P}$ ), trusted ( $\mathcal{T}$ ) and untrusted ( $\mathcal{U}$ ) sets over the lattice of labels.

$$\begin{aligned} \mathcal{P}_{\langle C, I \rangle} &= \{ \langle p, q \rangle \mid p \in C \} , & \mathcal{U}_{\langle C, I \rangle} &= \{ \langle p, q \rangle \mid q \in I \} , \\ \mathcal{S}_{\langle C, I \rangle} &= \{ \langle p, q \rangle \mid p \notin C \} , & \mathcal{T}_{\langle C, I \rangle} &= \{ \langle p, q \rangle \mid q \notin I \} \end{aligned}$$

### 3.4 Security Hyperproperties and Static Delegation

Our semantic framework formalizes a key insight relating the attacker model and delegation: more delegation means fewer attackers. In turn, fewer attackers should make it easier for programs to be considered secure. To formalize this intuition, we



**Fig. 3.** The lattice of labels  $(\mathbb{L}, \leq)$  and the lattice of information flow  $(\mathbb{L}, \sqsubseteq)$  share the same underlying set, but use a different ordering. The dotted line depicts labels with equal confidentiality and integrity: strictly above are compromised labels, and on or below are uncompromised labels.

study security conditions in the framework of *hyperproperties* [12]. For generality, we leave both the computation model and the system state abstract.

The simple system has states  $\sigma \in \Sigma$ , intentionally left unspecified. An execution of the system emits a trace  $t$ , which is a sequence of states. Define the *behavior*  $B \in \mathbb{B}$  of a program to be the *set* of possible execution traces it can emit. A hyperproperty  $\mathbb{HP} \subseteq \mathbb{B}$  is a set of behaviors. A program satisfies a hyperproperty when its behavior is a member of the hyperproperty.

In decentralized IFC systems, hyperproperties  $\mathbb{HP}_{\mathcal{A}}$  are parameterized by the choice of attacker: a program secure against one attacker may be insecure against another. Let the hyperproperty  $\mathbb{HP}_{\mathbf{A}}$  be the hyperproperty that characterizes programs that are secure against all possible attackers from  $\mathbf{A}$ . A behavior  $B$  of a program falls into  $\mathbb{HP}_{\mathbf{A}}$  precisely when  $B \in \mathbb{HP}_{\mathcal{A}}$  for all  $\mathcal{A} \in \mathbf{A}$ :

$$\mathbb{HP}_{\mathbf{A}} = \bigcap_{\mathcal{A} \in \mathbf{A}} \mathbb{HP}_{\mathcal{A}} \quad \Theta(\mathbb{HP}_{\mathbf{A}}) = \bigcap_{\mathcal{A} \in \mathbf{A}_{|\Theta}} \mathbb{HP}_{\mathcal{A}}$$

When a program assumes a static delegation context  $\Theta$ , the attacker model is further restricted to the ones consistent with  $\Theta$ . Therefore, a delegation context can be understood as a hyperproperty transformer. It follows that hyperproperties accept at least as many programs after a delegation context is added.

**Theorem 3.**  $\mathbb{HP}_{\mathbf{A}} \subseteq \Theta(\mathbb{HP}_{\mathbf{A}})$ .

This matches our intuition about static delegation: the more trust assumptions, the fewer reasonable attackers, the more programs considered secure.

### 3.5 Noninterference and the Lattice of Information Flow

For confidentiality, noninterference [13, 37] demands that information should not flow from high to low (secret to public). Noninterference of integrity requires that information should not flow from low to high (untrusted to trusted). To illustrate how delegation affects noninterference, we adapt Clarkson and Schneider’s [12] definition of *observational determinism* [18].

**Definition 3 (Observational Determinism).** *Let  $\mathcal{L}$  be any set of low labels. Observational determinism  $\mathbb{O}\mathbb{D}$  is the hyperproperty:*

$$\mathbb{O}\mathbb{D}_{\mathcal{L}} = \{B \mid \forall t_1, t_2 \in B. t_1^0 =_{\mathcal{L}} t_2^0 \implies t_1 \approx_{\mathcal{L}} t_2\}$$

State  $t^0$  denotes the initial state of the trace  $t$ . We leave the definition of low equivalence between states unspecified, as in prior work on knowledge-based security [6, 24, 26, 40].

Noninterference for confidentiality  $\text{NI}_{\mathcal{A}}^c = \mathbb{O}\mathbb{D}_{\mathcal{P}_{\mathcal{A}}}$  and integrity  $\text{NI}_{\mathcal{A}}^i = \mathbb{O}\mathbb{D}_{\mathcal{T}_{\mathcal{A}}}$  are mere instantiations of observational determinism over public and trusted labels for some attacker  $\mathcal{A}$ . For an attacker model  $\mathbf{A}_{|\Theta}$ :

$$\Theta(\text{NI}_{\mathbf{A}}) = \bigcap_{\mathcal{A} \in \mathbf{A}_{|\Theta}} (\mathbb{O}\mathbb{D}_{\mathcal{P}_{\mathcal{A}}} \cap \mathbb{O}\mathbb{D}_{\mathcal{T}_{\mathcal{A}}})$$

Prior work [19, 23, 30] enforces low equivalence by ensuring that low-labeled information is not influenced by high-labeled information. Concretely, a dynamic *relabel* is safe when it does not relabel high-labeled information to a low label. We formalize safe relabeling as the *flows-to* relation.

**Definition 4 (Flows-to).** *Label  $\ell$  securely flows to  $\ell'$  ( $\Theta \models \ell \sqsubseteq \ell'$ ) when:*

$$\forall \mathcal{A} \in \mathbf{A}_{|\Theta}. (\ell' \in \mathcal{P}_{\mathcal{A}} \implies \ell \in \mathcal{P}_{\mathcal{A}}) \wedge (\ell' \in \mathcal{T}_{\mathcal{A}} \implies \ell \in \mathcal{T}_{\mathcal{A}})$$

*Equivalently,  $\langle \theta_c, \theta_i \rangle \models \langle p, q \rangle \sqsubseteq \langle p', q' \rangle$  when  $\theta_c \models p' \leq p$  and  $\theta_i \models q \leq q'$ .*

As visualized in fig. 3, flows-to and the acts-for define two lattices on the same underlying set of labels. Information Lattice operators are given by:

$$\langle p_1, q_1 \rangle \sqcup \langle p_2, q_2 \rangle = \langle p_1 \wedge p_2, q_1 \vee q_2 \rangle \quad \langle p_1, q_1 \rangle \sqcap \langle p_2, q_2 \rangle = \langle p_1 \vee p_2, q_1 \wedge q_2 \rangle$$

### 3.6 Downgrading and Nonmalleable Information Flow

Some programs, such as the MPC example from fig. 2, intentionally break noninterference. Prior work on dynamic security policies either achieve downgrading by *relabeling* or by *dynamic delegation*. Visually, relabeling move information downward in fig. 3 and dynamic delegation moves the attacker partition leftward.

**Nonmalleable Information Flow (NMIF)** To prevent misuse of downgrades by relabeling, Cecchetti et al. [9] propose NMIF, a security hyperproperty that combines robust declassification [46] with transparent endorsement.

Robust declassification requires that secret information flow to public only when the information is trusted. This restriction ensures that attackers do not influence disclosure of information to them. A declassification is only robust when it declassifies secret–trusted information ( $\mathcal{S}_A \cup \mathcal{T}_A$ ).

Transparent endorsement is a dual condition that allows untrusted information to influence trusted information only when the information is public to the attacker. It only allows endorsement of public–untrusted information ( $\mathcal{P}_A \cap \mathcal{U}_A$ ).

Therefore, secret–untrusted information should not be downgraded. Indeed, the key recipe to enforcing NMIF is to enforce noninterference for public or trusted information [9]:

$$\text{NI}_{\mathbf{A}}^{\nabla} = \text{OD}_{\mathcal{T}_A \cup \mathcal{P}_A}$$

A label is *compromised* when it is secret–untrusted for some attacker [9, 45]. To enforce NMIF, it suffices to reject downgrading information with compromised labels, so that there is no flow out from compromised labels.

Unfortunately, NMIF against  $\mathbf{A}_{|\theta}$  rejects all downgrades. Namely, all labels (except for  $\langle \top, \perp \rangle$ ) are compromised against the attacker  $\mathcal{A}_{\top} = \langle \{\top\}, \mathbb{P} \rangle$  (it controls every principal for integrity and controls no principal for confidentiality). Therefore, further restrictions over the attacker model is needed.

**NMIF Attackers** Prior work [45, 46] assumes attackers control more confidentiality than integrity. We call them *valid attackers*:

**Definition 5 (Valid Attackers).**  $V = \{ \langle C, I \rangle \in \mathbf{A} \mid I \subseteq C \}$

Restricting attackers to valid ones makes our framework mirror attacker models studied in the cryptography literature [44], where a principal is honest (not controlled by attacker), semi-honest (controlled by attacker for confidentiality), or malicious (controlled by attacker for both confidentiality and integrity). The valid attacker restriction excludes the unrealistic “malicious but not curious” attackers.

In fact, much existing work satisfy the valid attacker assumption by construction. For example, robust declassification [46] originally defines integrity and confidentiality by equivalence relations over system states. The state transitions an active attacker may perform are, by construction, observable by the attacker.

**Definition 6 (Uncompromised Labels).** *Label  $\ell$  is uncompromised under  $\Theta$ , written  $\Theta \models \nabla \ell$ , when  $\ell$  is either public or trusted for all valid attackers:*

$$\forall \mathcal{A} \in V_{|\Theta} \cdot \ell \in \mathcal{P}_A \cup \mathcal{T}_A$$

It follows that labels of principals are uncompromised:  $\Theta \models \nabla \langle p, p \rangle$ .

Uncompromised labels have an alternative characterization: they are the labels with at least as much integrity as confidentiality. Of course, in the presence of asymmetric delegation, we cannot directly compare a label’s confidentiality

and integrity components since each component has a different set of delegations. We circumvent this problem by introducing a witnessing principal  $r$  who has no more integrity than  $q$  and no less confidentiality than  $p$ .

**Theorem 4.**  $\langle \theta_c, \theta_i \rangle \models \blacktriangledown \langle p, q \rangle \iff \exists r \in \mathbb{P}. (\theta_i \models q \leq r) \wedge (\theta_c \models r \leq p)$ .

In the absence of asymmetric delegation, theorem 4 reduces to a simple acts-for check ( $\theta \models q \leq p$ ), which prior systems rely on [9, 46].

## 4 Algorithms

In IFC systems that incorporate delegation, the acts-for relation  $\theta \models p_1 \leq p_2$  is frequently checked by label inference

procedures [34] and even at run time, where it can impose significant overhead [11]. Unfortunately, its semantic definition quantifies over the potentially infinite set of all attackers, which makes direct use of definition infeasible in practice. Similarly, the definition of uncompromised labels  $\Theta \models \blacktriangledown \ell$  does not yield an algorithm. In this section, we assume oracle access to syntactic lattice operations ( $\Rightarrow, \perp, \top, \wedge, \vee$ ) of  $\mathbb{P}$ , and derive sound and complete algorithms that check for the aforementioned relations.

### 4.1 Acts-for Algorithm

Algorithm 1 gives an algorithm for deciding  $\theta \models p \leq q$ . We write  $\theta \vdash p \leq q$  to denote this algorithmic system.

**Algorithm 1 (Acts-for  $\theta \vdash p \leq q$ ).**

$$\begin{array}{c} \mathbb{P}\text{-AXIOM} \\ \frac{p \Rightarrow q}{\cdot \vdash p \leq q} \end{array} \qquad \begin{array}{c} \mathbb{P}\text{-DELEGATION} \\ \frac{\theta \vdash p \wedge q' \leq q \quad \theta \vdash p \leq q \vee p'}{\theta, p' \Rightarrow q' \vdash p \leq q} \end{array}$$

The algorithm recursively applies the lattice axioms and rule  $\mathbb{P}\text{-DELEGATION}$  until the delegation context is empty for all sub-cases. The acts-for relation holds when rule  $\mathbb{P}\text{-AXIOM}$  apply to all sub-cases.

**Theorem 5 (Correctness of Acts-for).** *Algorithm 1 is sound and complete with respect to definition 2:  $\theta \vdash p \leq q \iff \theta \models p \leq q$ .*

### 4.2 NMIF Algorithm

Neither the semantic definition of uncompromised labels (definition 6) nor their alternative characterization (theorem 4) lends itself to an algorithmic implementation. The semantic definition quantifies over all valid attackers (a potentially infinite set), and the alternative characterization conjures up an intermediate principal. Our solution is to use the alternative characterization but with a

“best principal.” For that, we use  $\min_\theta(p)$ , the highest-authority principal in the equivalence class of  $p$ .<sup>7</sup>

**Definition 7 (Minimal Principal).**  $\min_\theta(p) \in \mathbb{P}$  is the (necessarily) unique principal such that  $\theta \models p \leq q$  if and only if  $\min_\theta(p) \Rightarrow q$  for any  $q \in \mathbb{P}$ .

We must demand more structure on  $\mathbb{P}$  to compute  $\min_\theta(p)$ . Specifically, we rely on an oracle that returns *join-prime* factorizations of arbitrary principals. Intuitively, a join-prime principal cannot be written as the join of other principals. In finite lattices, these are the principals of the form  $p = q_1 \wedge \dots \wedge q_n$ . Factorization oracles arise trivially in existing implementations of IFC models, as these are based on lattices with a finite set of principal names [3, 4, 34, 41].

**Algorithm 2 (Min  $\min_\theta(p) = q$ ).**

$$\begin{array}{c}
 \text{MIN-BASE} \\
 \frac{\text{join-prime}(p) \quad \forall(p' \Rightarrow q') \in \theta \bullet p \not\Rightarrow p'}{\min_\theta(p) = p}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{MIN-PICK} \\
 \frac{\text{join-prime}(p) \quad p \Rightarrow p'}{\min_{\theta, p' \Rightarrow q'}(p) = \min_\theta(p \wedge q')}
 \end{array}$$

$$\begin{array}{c}
 \text{MIN-FACTOR} \\
 \frac{\neg \text{join-prime}(p) \quad p = p_1 \vee \dots \vee p_n \quad \forall i \in [n] \bullet \text{join-prime}(p_i)}{\min_\theta(p) = \bigvee_{i \in [n]} \min_\theta(p_i)}
 \end{array}$$

The rules from Algorithm 2 can be applied in any order without backtracking because of the uniqueness of  $\min_\theta(p)$ . Moreover, all derivations are finite since either the size of  $\theta$  decreases, or we switch from a reducible element to a finite set of irreducible elements.

**Theorem 6 (Correctness of Min).** *Algorithm 2 is sound and complete with respect to definition 7.*

Our NMIF algorithm simply combines algorithms 1 and 2.

**Algorithm 3 (Uncompromised Label Check  $\Theta \vdash \blacktriangledown \ell$ ).**

$$\frac{\theta_c \vdash \min_{\theta_i}(q) \leq p}{\langle \theta_c, \theta_i \rangle \vdash \blacktriangledown \langle p, q \rangle}$$

**Theorem 7 (Correctness of Uncompromised Label Check).** *Algorithm 3 is sound and complete with respect to theorem 4.*

<sup>7</sup> Recall that higher authority is lower in the authority lattice, thus the use of min as opposed to max.

## 5 Label Inference

In practical IFC systems, label inference is an important way to avoid redundant user annotations, since it is secure to infer labels of intermediate computations from their inputs and outputs. Typically, label inference is performed by solving a system of constraints over lattice elements [34, 48]. IFC systems further benefit from *bounded label polymorphism*, which allows user to write library code reusable at different security levels. In this section, we show how to do label inference directly over the algebraic model of labels using the algorithms from Section 4.

```
1 host Alice, Bob, Chuck
2 assume Alice = Bob for integrity
3 assume Bob = Chuck for integrity
4
5 fun average(a: int, b: int): int
6 {
7   return (a + b) / 2
8 }

8 fun main() {
9   val a = Alice.input
10  val b = Bob.input
11  val c = Chuck.input
12
13  val r1 = average(a, b)
14  val r2 = average(b, c)
15 }
```

**Fig. 4.** The average function is implicitly polymorphic over the labels of its arguments.

### 5.1 Bounded Label Polymorphism

As in traditional type systems, allowing code that is generic over labels increases expressiveness significantly. Existing languages like Jif [34] and Flow Caml [39] support bounded label polymorphism, which allows functions to be parameterized over labels that are bounded by specified security levels. Annotation burden on users can be further reduced by assuming information flow from function arguments to return values by default.

In fig. 4, the annotation-free polymorphic function `average` is applied in `main` to arguments with different security labels. Shown below is the same function with explicit annotations, all of which can be inferred.

```
1 fun average[X, Y, Z](a: int{X}, b: int{Y}): int{Z}
2   where (X  $\sqcup$  Y  $\sqsubseteq$  Z)
```

Label inference assigns existential *label variables* to all unlabeled expressions and creates constraints based on IFC typing rules. The constraints are then solved by a constraint solver that uses parameter bounds as delegation contexts.

In each function, polymorphic label variables are treated as label constants, so solutions of label variables are expressed by both label constants and polymorphic label variables. Type checking at call sites ensures that the parameter bounds are satisfied. As functions have their own delegation contexts, a new constraint system is solved for each function.

## 5.2 Constraint Solver

We describe a novel constraint solver that computes the minimum-semantic-authority solution to constraint systems with delegation contexts. Minimum-authority solutions are desirable because they allow systems to choose cheaper security enforcement mechanisms [3, 10, 15, 16, 47, 49].

**Label and Principal Constraints** Figure 5 gives the syntax of the label constraint language. Expressions in the constraint language include label constants and label variables, authority projections, as well as standard lattice operations. A constraint either asserts that an expression flows to another, or asserts that an expression is uncompromised. We translate constraints over labels to constraints over principals to leverage algorithms from Section 4. Figure 6 gives the syntax

L. Constants  $\ell$   
 L. Variables  $Y$   
 L. Expressions  $L ::= \ell \mid Y \mid L^\pi$   
                    $\mid L_1 \sqcup L_2 \mid L_1 \sqcap L_2$   
                    $\mid L_1 \vee L_2 \mid L_1 \wedge L_2$   
 L. Constraints  $C ::= L_1 \sqsubseteq L_2 \mid \blacktriangledown L$   
 Projections  $\pi \in \{c, i\}$

**Fig. 5.** Syntax of label constraints.

P. Constants  $p$   
 P. Variables  $Y^\pi$   
 P. Expressions  $P^\pi ::= p \mid Y^\pi$   
                    $\mid P_1^\pi \vee P_2^\pi \mid P_1^\pi \wedge P_2^\pi$   
                    $\mid p_1 \rightarrow P_2^\pi \mid \min_{\pi'}(P^{\pi'})$   
 P. Constraints  $D ::= P_1^\pi \Rightarrow^\pi P_2^\pi$

**Fig. 6.** Syntax of principal constraints.

of the principal constraint language. The syntax includes a principal variable  $Y^\pi$  for each combination of label variable  $Y$  and projection  $\pi$ . That is,  $Y^c$  represents the confidentiality of  $Y$ , and  $Y^i$  represents  $Y$ 's integrity.

We index expressions  $P^\pi$  by the component  $\pi$  they represent. This prevents expressions like  $Y_1^c \wedge Y_2^i$ , whose components are mixed. Expressions include principal constants and principal variables, as well as principal-lattice operations  $\vee$  and  $\wedge$ . The operation  $\rightarrow$  is called the *relative pseudocomplement* of the meet operation:  $p_1 \rightarrow p_2$  is defined as the minimal-authority principal  $p$  such that  $p_1 \wedge p \Rightarrow p_2$ . We use  $\rightarrow$  to solve constraints of the form  $Y^\pi \wedge p_1 \Rightarrow^\pi P_2^\pi$ . The  $\min_{\pi'}()$  operation allows mixing integrity and confidentiality components; we use it when solving for labels that must be uncompromised. Principal constraints have the form  $P_1^\pi \Rightarrow^\pi P_2^\pi$ , which stands for  $\theta_\pi \models P_1^\pi \leq P_2^\pi$ .

Figure 7 gives rules for translating label constraints to principal constraints. The definition of  $\llbracket L \rrbracket_\pi$  is a straightforward encoding of the label-lattice operations. Using  $\llbracket L \rrbracket$ , we translate a flows-to ( $\sqsubseteq$ ) constraint to two acts-for ( $\Rightarrow$ ) constraints, one for each label component. The constraint  $\blacktriangledown L$  follows from algorithm 3.

**Solving Principal Constraints** Our constraint solver requires the left-hand side of each constraint to be atomic (a constant or a variable), that is, constraints of the form  $p_1 \Rightarrow^\pi P_2^\pi$  and  $Y_1^\pi \Rightarrow^\pi P_2^\pi$ . We exhaustively apply the equational

$$\boxed{\llbracket C \rrbracket = D_1, \dots, D_n}$$

$$\llbracket L_1 \sqsubseteq L_2 \rrbracket = \llbracket L_2 \rrbracket_c \Rightarrow^c \llbracket L_1 \rrbracket_c, \llbracket L_1 \rrbracket_i \Rightarrow^i \llbracket L_2 \rrbracket_i \quad \llbracket \blacktriangledown L \rrbracket = \llbracket L \rrbracket_i \Rightarrow^i \text{min}_c(\llbracket L \rrbracket_c)$$

$$\boxed{\llbracket L \rrbracket_\pi = P^\pi}$$

$$\begin{aligned} \llbracket L_1 \sqcup L_2 \rrbracket_\pi &= \llbracket (L_1 \wedge L_2)^c \wedge (L_1 \vee L_2)^i \rrbracket_\pi & \llbracket \langle p, q \rangle \rrbracket_\pi &= \begin{cases} p & \text{if } \pi = c \\ q & \text{if } \pi = i \end{cases} \\ \llbracket L_1 \sqcap L_2 \rrbracket_\pi &= \llbracket (L_1 \vee L_2)^c \wedge (L_1 \wedge L_2)^i \rrbracket_\pi & \llbracket Y \rrbracket_\pi &= Y^\pi \\ \llbracket L_1 \vee L_2 \rrbracket_\pi &= \llbracket L_1 \rrbracket_\pi \vee \llbracket L_2 \rrbracket_\pi & \llbracket L^{\pi'} \rrbracket_\pi &= \begin{cases} \llbracket L \rrbracket_\pi & \text{if } \pi = \pi' \\ \top & \text{if } \pi \neq \pi' \end{cases} \\ \llbracket L_1 \wedge L_2 \rrbracket_\pi &= \llbracket L_1 \rrbracket_\pi \wedge \llbracket L_2 \rrbracket_\pi \end{aligned}$$

**Fig. 7.** Translating label constraints to principal constraints.

axioms of Heyting algebra (e.g., associativity, absorption, distributivity, etc.) until no left-hand side of any constraint can be further simplified. As equational axioms are syntactic rewrites, it does not change the constraint system over the underlying algebra. Moreover, this process always terminates, and ensures that the left-hand side of each constraint either is atomic or contains a meet ( $\wedge$ ).<sup>8</sup>

Constraint solving fails if the left-hand side of any constraint contains a meet, since such systems do not have unique solutions. For example, the system  $Y_1 \wedge Y_2 \Rightarrow \text{Alice}$  has no minimal solution: we can assign  $\{Y_1 \mapsto \text{Alice}, Y_2 \mapsto \top\}$  or  $\{Y_1 \mapsto \top, Y_2 \mapsto \text{Alice}\}$ , but neither solution is better than the other. Compiler implementations need to either restrict the syntax of polymorphic constraints, or report errors during constraint solving.

Once the simplification process succeeds, we extend the algorithm of Rehof and Mogensen [35] for iteratively solving semi-lattice constraints. We initialize all principal variables to  $\top$ , and use unsatisfied constraints to update variables repeatedly until a fixed point is reached, using the rule:

$$\text{given } Y^\pi \Rightarrow^\pi P^\pi, \quad \text{set } Y^\pi := Y^\pi \wedge \text{current-value}(\Theta, P^\pi),$$

where  $\text{current-value}(\Theta, P^\pi)$  is the value of  $P^\pi$  according to the current assignment.

Note that constraints that have constants  $p$  on the left-hand side are ignored during the fixed point computation. Once a fixed-point solution is reached, we perform the following check for each constraint with a constant left-hand side:

$$\text{given } p \Rightarrow^\pi P^\pi, \quad \text{check } \theta_\pi \models p \leq \text{current-value}(\Theta, P^\pi).$$

We have a minimal-authority solution if all such constraints are satisfied; otherwise, there is no valid solution.

<sup>8</sup> Translation rules in fig. 7 never generate constraints with  $\rightarrow$  or  $\text{min}_\pi(\cdot)$  on the left-hand side, and the constraint simplification eliminate all joins ( $\vee$ ) on the left.

### 5.3 Implementation

We modified the parser of the Viaduct compiler [3] with the delegation syntax, and extended its static analysis procedure with our label inference algorithm.<sup>9</sup>

Because the original Viaduct constraint solver can only make syntactic comparison ( $\Rightarrow$ ) between labels, it instantiates polymorphic variables with constant principal names. Therefore, Viaduct has to run a *specialization procedure* to create a monomorphic copy of a function at each call site. In nested function calls, the number of monomorphic functions created grows exponentially with the depth of calls. Recursive calls need to be handled explicitly to ensure termination.

Thanks to delegation contexts, label inference no longer requires monomorphic functions and can be done in one pass. For each function call site, the specialization procedure memoizes the argument labels, and avoid creating duplicates of monomorphic functions that are instantiated with the same polymorphic argument labels. Since all specialized functions have different labels, our procedure creates a minimal number of monomorphic functions. Recursive functions no longer need to be separately handled because the signatures of newly created monotonic copies eventually reach a fixed point, which are memoized.

## 6 Related Work

Expressive trust delegation has been widely investigated in access control and capability systems [25, 38], where delegation is called *role hierarchy* [38]. Such systems support role adoption, which is a form of dynamic delegation. However, access control systems generally do not connect to information flow security conditions. Many prior IFC systems have delegation abstractions compatible with our framework, but they either lack support for asymmetric delegation [3, 9, 45], or do not explore its effect over hyperproperties [4, 26, 32, 33, 46].

Our inference algorithm differs from prior implementations of syntax-directed IFC label inference [3, 34, 48] by operating directly over the underlying algebra. This algebraic approach enhances extensibility: adding principals or delegations does not invalidate existing analysis. Dolan [14] proposes an inference algorithm for an expressive algebraic type system with subtyping, type constructors, and recursive function types. However, their algorithm produces contrived representations of typing schemes, limiting its practicality. In contrast, our label system balances expressiveness with a simple and effective inference algorithm.

FLAM [4] proposes a label model and defines *robust authorization* to address security vulnerabilities arising from dynamic delegation. However, robust authorization is a proof-theoretic definition that lacks semantics. Arden and Myers [5] propose the FLAC calculus, based on the FLAM label model, and prove robust declassification for a language that downgrades using dynamic delegation. However, FLAC has no attacker semantics. Cecchetti et al. [9] define NMIF, but their system cannot explicitly express delegations between atomic principals.

---

<sup>9</sup> Code available at: <https://github.com/apl-cornell/viaduct>.

## 7 Conclusion and Future Directions

We present an algebraic semantic framework for IFC labels that models asymmetric delegation, along with sound and complete algorithms that enforce security conditions and inference security annotations. Our approach provides a solid foundation for building modular, expressive and extensible IFC systems.

**Acknowledgments.** This work was supported by the National Science Foundation under NSF grant 1704788. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. We also thank Suraj Kanniwadi and Ethan Cecchetti for discussions and feedback.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## Bibliography

- [1] Abadi, M.: Access control in a core calculus of dependency. In: 11<sup>th</sup> ACM SIGPLAN Int'l Conf. on Functional Programming. pp. 263–273. ACM, New York, NY, USA (2006). <https://doi.org/10.1145/1159803.1159839>
- [2] Abadi, M.: Variations in access control logic. In: van der Meyden, R., van der Torre, L. (eds.) Deontic Logic in Computer Science, Lecture Notes in Computer Science, vol. 5076, pp. 96–109. Springer Berlin Heidelberg (2008). [https://doi.org/10.1007/978-3-540-70525-3\\_9](https://doi.org/10.1007/978-3-540-70525-3_9)
- [3] Acay, C., Recto, R., Gancher, J., Myers, A., Shi, E.: Viaduct: An extensible, optimizing compiler for secure distributed programs. In: 42<sup>nd</sup> ACM SIGPLAN Conf. on Programming Language Design and Implementation (PLDI). pp. 740–755. ACM (Jun 2021). <https://doi.org/10.1145/3453483.3454074>
- [4] Arden, O., Liu, J., Myers, A.C.: Flow-limited authorization. In: 28<sup>th</sup> IEEE Computer Security Foundations Symp. (CSF). pp. 569–583 (Jul 2015). <https://doi.org/10.1109/CSF.2015.42>
- [5] Arden, O., Myers, A.C.: A calculus for flow-limited authorization. In: 29<sup>th</sup> IEEE Computer Security Foundations Symp. (CSF). pp. 135–147 (Jun 2016). <https://doi.org/10.1109/CSF.2016.17>
- [6] Askarov, A., Chong, S.: Learning is change in knowledge: Knowledge-based security for dynamic policies. In: 2012 IEEE 25th Computer Security Foundations Symposium. pp. 308–322 (2012). <https://doi.org/10.1109/CSF.2012.31>
- [7] Biba, K.J.: Integrity considerations for secure computer systems. Tech. Rep. ESD-TR-76-372, USAF Electronic Systems Division, Bedford, MA (Apr 1977), <https://ban.ai/multics/doc/a039324.pdf>, (Also available through National Technical Information Service, Springfield Va., NTIS AD-A039324.)
- [8] Blok, W.J., Pigozzi, D.: Algebraizable Logics (1989)
- [9] Cecchetti, E., Myers, A.C., Arden, O.: Nonmalleable information flow control. In: 24<sup>th</sup> ACM Conf. on Computer and Communications Security (CCS). pp. 1875–1891. ACM (Oct 2017). <https://doi.org/10.1145/3133956.3134054>
- [10] Chong, S., Liu, J., Myers, A.C., Qi, X., Vikram, K., Zheng, L., Zheng, X.: Secure web applications via automatic partitioning. In: 21<sup>st</sup> ACM Symp. on Operating System Principles (SOSP). pp. 31–44 (Oct 2007). <https://doi.org/10.1145/1323293.1294265>
- [11] Chong, S., Vikram, K., Myers, A.C.: SIF: Enforcing confidentiality and integrity in web applications. In: 16<sup>th</sup> USENIX Security Symp. (Aug 2007), <http://www.cs.cornell.edu/andru/papers/sif.pdf>
- [12] Clarkson, M.R., Schneider, F.B.: Hyperproperties. In: 21<sup>st</sup> IEEE Computer Security Foundations Symp. (CSF). pp. 51–65 (Jun 2008). <https://doi.org/10.1109/CSF.2008.7>
- [13] Denning, D.E.: A lattice model of secure information flow. *Comm. of the ACM* **19**(5), 236–243 (1976). <https://doi.org/10.1145/360051.360056>
- [14] Dolan, S.: Algebraic subtyping: Distinguished Dissertation 2017. BCS Learning & Development Ltd, Swindon, GBR (2017)

- [15] Fournet, C., le Guernic, G., Rezk, T.: A security-preserving compiler for distributed programs: From information-flow policies to cryptographic mechanisms. In: 16<sup>th</sup> ACM Conf. on Computer and Communications Security (CCS). pp. 432–441 (Nov 2009), <https://doi.org/10.1145/1653662.1653715>
- [16] Fournet, C., Rezk, T.: Cryptographically sound implementations for typed information-flow security. In: 35<sup>th</sup> ACM Symp. on Principles of Programming Languages (POPL). pp. 323–335 (Jan 2008), <https://doi.org/10.1145/1328438.1328478>
- [17] Garg, D., Pfenning, F.: Non-interference in constructive authorization logic. In: 19<sup>th</sup> IEEE Computer Security Foundations Workshop (CSFW) (2006). <https://doi.org/10.1109/CSFW.2006.18>
- [18] Goguen, J.A., Meseguer, J.: Security policies and security models. In: IEEE Symp. on Security and Privacy. pp. 11–20 (Apr 1982). <https://doi.org/10.1109/SP.1982.10014>
- [19] Goguen, J.A., Meseguer, J.: Unwinding and inference control. In: IEEE Symp. on Security and Privacy. pp. 75–86 (Apr 1984). <https://doi.org/10.1109/SP.1984.10019>
- [20] Hirsch, A.K., Amorim, P.H.A.d., Cecchetti, E., Tate, R., Arden, O.: First-order logic for flow-limited authorization. In: 2020 IEEE 33rd Computer Security Foundations Symposium (CSF). pp. 123–138 (2020). <https://doi.org/10.1109/CSF49147.2020.00017>
- [21] Hirsch, A.K., Clarkson, M.R.: Belief semantics of authorization logic. CCS ’13, Association for Computing Machinery, New York, NY, USA (2013). <https://doi.org/10.1145/2508859.2516667>
- [22] Hirt, M., Maurer, U.M.: Player simulation and general adversary structures in perfect multiparty computation. *J. Cryptol.* **13**(1), 31–60 (2000). <https://doi.org/10.1007/S001459910003>
- [23] Kozyri, E., Chong, S., Myers, A.C.: Expressing information flow properties. *Foundations and Trends in Privacy and Security* **3**(1), 1–102 (2022). <https://doi.org/10.1561/3300000008>
- [24] Landauer, J., Redmond, T.: A lattice of information. In: 6<sup>th</sup> IEEE Computer Security Foundations Workshop (CSFW). pp. 65–70. IEEE Computer Society Press (Jun 1993). <https://doi.org/10.1109/CSFW.1993.246638>
- [25] Li, N., Grosf, B.N., Feigenbaum, J.: Delegation logic: A logic-based approach to distributed authorization. *ACM Transactions on Information and System Security (TISSEC)* **6**(1), 128–171 (2003). <https://doi.org/10.1145/605434.605438>
- [26] Li, P., Zhang, D.: Towards a general-purpose dynamic information flow policy (2021), <https://arxiv.org/abs/2109.08096>
- [27] Matetic, S., Schneider, M., Miller, A., Juels, A., Capkun, S.: Delegatee: Brokered delegation using trusted execution environments. In: *USENIX Security Symposium*. pp. 1387–1403 (2018)
- [28] Montagu, B., Pierce, B.C., Pollack, R.: A theory of information-flow labels. In: 26<sup>th</sup> IEEE Computer Security Foundations Symp. (CSF). pp. 3–17 (Jun 2013). <https://doi.org/10.1109/CSF.2013.8>

- [29] Myers, A.C.: Mostly-Static Decentralized Information Flow Control. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA (Jan 1999)
- [30] Myers, A.C., Liskov, B.: A decentralized model for information flow control. In: 16<sup>th</sup> ACM Symp. on Operating System Principles (SOSP). pp. 129–142 (Oct 1997). <https://doi.org/10.1145/268998.266669>
- [31] Myers, A.C., Liskov, B.: Protecting privacy using the decentralized label model. *ACM Transactions on Software Engineering and Methodology* **9**(4), 410–442 (Oct 2000). <https://doi.org/10.1145/363516.363526>
- [32] Myers, A.C., Sabelfeld, A., Zdancewic, S.: Enforcing robust declassification. In: 17<sup>th</sup> IEEE Computer Security Foundations Workshop (CSFW). pp. 172–186 (Jun 2004). <https://doi.org/10.1109/CSFW.2004.9>
- [33] Myers, A.C., Sabelfeld, A., Zdancewic, S.: Enforcing robust declassification and qualified robustness. *Journal of Computer Security* **14**(2), 157–196 (2006). <https://doi.org/10.3233/JCS-2006-14203>
- [34] Myers, A.C., Zheng, L., Zdancewic, S., Chong, S., Nystrom, N.: Jif 3.0: Java information flow (Jul 2006), <http://www.cs.cornell.edu/jif>, software release, <http://www.cs.cornell.edu/jif>
- [35] Rehof, J., Mogensen, T.A.: Tractable constraints in finite semilattices. In: 3rd International Symposium on Static Analysis. pp. 285–300. No. 1145 in *Lecture Notes in Computer Science*, Springer-Verlag (Sep 1996). [https://doi.org/10.1007/3-540-61739-6\\_48](https://doi.org/10.1007/3-540-61739-6_48)
- [36] Rutherford, D.E.: *Introduction to Lattice Theory*. Oliver and Boyd (1965)
- [37] Sabelfeld, A., Myers, A.C.: Language-based information-flow security. *IEEE Journal on Selected Areas in Communications* **21**(1), 5–19 (Jan 2003). <https://doi.org/10.1109/JSAC.2002.806121>
- [38] Sandhu, R.S.: Role hierarchies and constraints for lattice-based access controls. In: 4<sup>th</sup> European Symp. on Research in Computer Security (ESORICS) (Sep 1996)
- [39] Simonet, V.: *The Flow Caml System: documentation and user’s manual*. Technical Report 0282, Institut National de Recherche en Informatique et en Automatique (INRIA) (Jul 2003)
- [40] Soloviev, M., Balliu, M., Guanciale, R.: Security properties through the lens of modal logic (2023). <https://doi.org/10.1109/csf61375.2024.00009>
- [41] Stefan, D., Russo, A., Mazières, D., Mitchell, J.C.: Disjunction category labels. In: Laud, P. (ed.) *Proceedings of the 16th Nordic conference on Information Security Technology for Applications*. *Lecture Notes in Computer Science*, vol. 7161, pp. 223–239. Springer (2011). [https://doi.org/10.1007/978-3-642-29615-4\\_16](https://doi.org/10.1007/978-3-642-29615-4_16)
- [42] Stone, M.H.: Topological representations of distributive lattices and brouwerian logics. *Časopis pro pěstování matematiky a fyziky* **067**(1), 1–25 (1938). <https://doi.org/10.21136/CPMF.1938.124080>
- [43] Swamy, N., Hicks, M., Tse, S., Zdancewic, S.: Managing policy updates in security-typed languages. In: 19<sup>th</sup> IEEE Computer Security Foundations Workshop (CSFW). pp. 202–216 (Jul 2006). <https://doi.org/10.1109/CSFW.2006.17>

- [44] Yao, A.C.: Protocols for secure computations. In: 23<sup>rd</sup> annual IEEE Symposium on Foundations of Computer Science. pp. 160–164 (1982). <https://doi.org/10.1109/SFCS.1982.38>
- [45] Zagieboylo, D., Suh, G.E., Myers, A.C.: Using information flow to design an ISA that controls timing channels. In: 32<sup>nd</sup> IEEE Computer Security Foundations Symp. (CSF) (Jun 2019). <https://doi.org/10.1109/CSF.2019.00026>
- [46] Zdancewic, S., Myers, A.C.: Robust declassification. In: 14<sup>th</sup> IEEE Computer Security Foundations Workshop (CSFW). pp. 15–23 (Jun 2001). <https://doi.org/10.1109/CSFW.2001.930133>
- [47] Zdancewic, S., Zheng, L., Nystrom, N., Myers, A.C.: Secure program partitioning. *ACM Trans. on Computer Systems* **20**(3), 283–328 (Aug 2002). <https://doi.org/10.1145/566340.566343>
- [48] Zhang, D., Myers, A.C., Vytiniotis, D., Peyton Jones, S.: SHErrLoc: A static holistic error locator. *ACM Trans. on Programming Languages and Systems* **39**(4), 18 (Aug 2017), <http://dl.acm.org/citation.cfm?id=3121137>
- [49] Zheng, L., Chong, S., Myers, A.C., Zdancewic, S.: Using replication and partitioning to build secure distributed systems. In: IEEE Symp. on Security and Privacy. pp. 236–250 (May 2003). <https://doi.org/10.1109/SECPRI.2003.1199340>
- [50] Zheng, L., Myers, A.C.: A language-based approach to secure quorum replication. In: 9<sup>th</sup> ACM SIGPLAN Workshop on Programming Languages and Analysis for Security (PLAS) (Aug 2014). <https://doi.org/10.1145/2637113.2637117>

## A Details for Section 3 (Semantic Framework)

**Definition 8 (Closures).** Define upward and downward closed sets as follows:

$$\uparrow p = \{q \in \mathbb{P} \mid p \Rightarrow q\} \quad \downarrow p = \{q \in \mathbb{P} \mid q \Rightarrow p\}$$

**Definition 9 (Equivalence Classes).** The equivalence class of a principal with respect to a delegation context is:

$$[p]_\theta = \{q \in \mathbb{P} \mid (\theta \models p \leq q) \wedge (\theta \models q \leq p)\}$$

We extend this notation to denote the equivalence closure of a set:

$$[P]_\theta = \bigcup_{p \in P} [p]_\theta$$

*Remark 1.* We denote the set of equivalence classes using standard notation:

$$P/\theta = P/\equiv_\theta = \{[p]_\theta \mid p \in P\}$$

for  $P \subseteq \mathbb{P}$  (we might have  $P = \mathbb{P}$ ). This is different from  $[P]_\theta \subseteq \mathbb{P}$ .

*Remark 2.* Note that  $[p]_\theta \in \mathbb{P}/\theta$  and  $\uparrow[p]_\theta \subseteq \mathbb{P}/\theta$ . In particular,  $\uparrow[p]_\theta$  is a set of sets of principals, not a set of principals.

**Theorem 8 (Prime Ideal/Filter).** Let  $L$  be a distributive lattice,  $I \subseteq L$  an ideal, and  $F \subseteq L$  a filter such that  $I \cap F = \emptyset$ . Then there exists a prime ideal  $P$  and a prime filter  $Q$  such that

$$I \subseteq P \quad F \subseteq Q \quad P \cap Q = \emptyset$$

**Theorem 1 (Algebraic Model).** The algebraic model of the principal lattice  $\mathbb{P}$  under delegation context  $\theta$  is the quotient lattice  $\mathbb{P}/\equiv_\theta$ .

*Proof.* Follows from the definition of Lindenbaum–Tarski Algebra of propositional intuitionistic logic. The Lindenbaum algebra of theory  $\theta$ ,  $\mathbb{P}/\equiv_\theta$ , is the initial algebra of the category of Heyting algebras that are consistent with  $\theta$ . In human language, there is a lattice homomorphism from  $\mathbb{P}/\equiv_\theta$  to every algebraic model  $\mathbb{J}$  of  $\theta$ .

**Theorem 2 (Attacker Model).**  $\mathbb{A}_\theta$  is the set of prime filters of  $\mathbb{P}/\equiv_\theta$ .

*Proof.* Since each attacker is a truth assignment of some Heyting algebra  $\mathbb{J}$ , it can be characterized by a homomorphism from  $\mathbb{J}$  to the two-point lattice  $2 = \{\top, \perp\}$ . Therefore, each attacker is uniquely characterized by a lattice homomorphism from  $\mathbb{P}/\equiv_\theta$  to  $2$ . We can then define each attacker by the kernel of its characterizing homomorphism. The set of such kernels are precisely the set of prime filters of  $\mathbb{P}/\equiv_\theta$ .

**Theorem 4.**  $\langle \theta_c, \theta_i \rangle \models \blacktriangledown \langle p, q \rangle \iff \exists r \in \mathbb{P}. (\theta_i \models q \leq r) \wedge (\theta_c \models r \leq p)$ .

*Proof.* We prove each direction separately.

- Case  $\implies$ . Assume, for contradiction, that  $q \notin [[\downarrow p]_{\theta_c}]_{\theta_i}$ . The sets  $\uparrow[q]_{\theta_i}$  and  $[\downarrow p]_{\theta_c}/\theta_i$  are disjoint and form a filter/ideal pair of  $\mathbb{P}/\theta_i$ . Thus, theorem 8 gives a prime filter  $I_0$  of  $\mathbb{P}/\theta_i$  with  $\uparrow[q]_{\theta_i} \subseteq I_0$  and  $I_0 \cap ([\downarrow p]_{\theta_c}/\theta_i) = \emptyset$ . Let  $I = \bigcup I_0$ . Note that  $I$  and  $[[\downarrow p]_{\theta_c}]_{\theta_i}$  are disjoint and  $[\downarrow p]_{\theta_c} \subseteq [[\downarrow p]_{\theta_c}]_{\theta_i}$ , so  $I$  and  $[\downarrow p]_{\theta_c}$  are disjoint as well. This in turn implies  $I/\theta_c$  and  $\downarrow p/\theta_c$  are disjoint. Moreover, these sets form a filter/ideal pair in  $\mathbb{P}/\theta_c$ , so theorem 8 gives a prime filter  $C_0$  of  $\mathbb{P}/\theta_c$  with  $I/\theta_c \subseteq C_0$  and  $C_0 \cap (\downarrow p/\theta_c) = \emptyset$ . Now, define  $C = \bigcup C_0$  and note that

$$I = \bigcup (I/\theta_c) \subseteq \bigcup C_0 = C$$

Since  $I_0$  is an ideal of  $\mathbb{P}/\theta_i$ , we have  $I \models \theta_i$ . Since  $C_0$  is an ideal of  $\mathbb{P}/\theta_c$ , we have  $C \models \theta_c$ . Combining these three results, we have  $\langle C, I \rangle \in \mathbf{V}_{|\langle \theta_c, \theta_i \rangle}$ . Because  $\langle \theta_c, \theta_i \rangle \models \blacktriangledown \langle p, q \rangle$  and  $q \in \uparrow[q]_{\theta_i} \subseteq I_0 \implies q \in I$ , we have  $p \in C$ . However,  $p \in [\downarrow p]_{\theta_c}$ , which contradicts the fact that  $C_0$  and  $\downarrow p/\theta_c$  are disjoint. Thus, we must have  $q \in [[\downarrow p]_{\theta_c}]_{\theta_i}$ .

Finally,  $q \in [[\downarrow p]_{\theta_c}]_{\theta_i}$  means there exists  $r \in [\downarrow p]_{\theta_c}$  such that  $q \in [r]_{\theta_i}$ . By definition, we have  $\theta_i \models q \leq r$ . Moreover,  $r \in [\downarrow p]_{\theta_c}$  means there exists  $p' \in \downarrow p$  such that  $r \in [p']_{\theta_c}$ . By definition,  $\theta_c \models r \leq p'$ , and theorem 2 with  $p' \in \downarrow p$  gives  $\theta_c \models p' \leq p$ . By transitivity, we have  $\theta_c \models r \leq p$ .

- Case  $\impliedby$ . Assume there exists  $r$  such that  $\theta_i \models q \leq r$  and  $\theta_c \models r \leq p$ . We claim  $\langle p, q \rangle$  is public whenever it is untrusted. Formally, let  $\langle C, I \rangle \in \mathbf{V}$  such that  $\langle C, I \rangle \models \langle \theta_c, \theta_i \rangle$ , and assume  $q \in I$ . We need to show  $p \in C$ . Since  $I \models \theta_i$ ,  $\theta_i \models q \leq r$ , and  $q \in I$ , we have  $r \in I$ . Moreover,  $I \subseteq C$  (definition 5), so  $r \in C$ . Finally, since  $C \models \theta_c$ ,  $\theta_c \models r \leq p$ , and  $r \in C$ , we have  $p \in C$  as desired.

## B Details for Section 4 (Algorithms)

**Lemma 1.** *If  $\theta_1 \models p \leq q$ , then  $(\theta_1, \theta_2) \models p \leq q$  for any  $\theta_2$ .*

*Proof.* Immediate since  $\mathbb{A}_{|\theta_1, \theta_2} \subseteq \mathbb{A}_{|\theta_1}$ .

**Corollary 1.** *If  $p \Rightarrow q$ , then  $\theta \models p \leq q$  for any  $\theta$ .*

*Proof.* We have  $\cdot \models p \leq q$  by theorem 2. The result then follows by lemma 1.

**Lemma 2.** *If  $\cdot \models p \leq q$ , then  $p \Rightarrow q$ .*

*Proof.* Assume, for contradiction, that  $p \not\Rightarrow q$ . The set  $\uparrow p = \{p' \in \mathbb{P} \mid p \Rightarrow p'\}$  is a filter of  $\mathbb{P}$ , and  $\downarrow q = \{q' \in \mathbb{P} \mid q' \Rightarrow q\}$  is an ideal of  $\mathbb{P}$ . Moreover,  $\uparrow p$  and  $\downarrow q$  are disjoint since  $p \not\Rightarrow q$ . By theorem 8, there exists a prime filter  $P \supseteq \uparrow p$  that is also disjoint from  $\downarrow q$ . This means  $p \in P$  and  $q \notin P$ . However,  $P \in \mathbf{A}$  by theorem 2 (attackers are prime filters) and  $P \models \cdot$  trivially, which contradicts the assumption  $\cdot \models p \leq q$ .

**Lemma 3.** *If  $(\theta, p_1 \Rightarrow q_1) \models p_2 \leq q_2$ , then  $\theta \models p_2 \wedge q_1 \leq q_2$ .*

*Proof.* Let  $A$  be an attacker such that  $A \models \theta$  and  $p_2 \wedge q_1 \in A$ . We need to show  $q_2 \in A$ . Since  $p_2 \wedge q_1 \in A$  and  $A$  is upward-closed (theorem 2), we have  $p_2, q_1 \in A$ . From  $A \models \theta$  and  $q_1 \in A$ , we get  $A \models \theta, p_1 \Rightarrow q_1$ . Using this, and the fact that  $p_2 \in A$ , we invoke our primary assumption to get  $q_2 \in A$ .

**Lemma 4.** *If  $(\theta, p_1 \Rightarrow q_1) \models p_2 \leq q_2$ , then  $\theta \models p_2 \leq q_2 \vee p_1$ .*

*Proof.* Let  $A$  be an attacker such that  $A \models \theta$  and  $p_2 \in A$ . We need to show  $q_2 \vee p_1 \in A$ . If  $p_1 \in A$ , then  $q_2 \vee p_1 \in A$  since  $A$  is upward-closed (theorem 2), so assume  $p_1 \notin A$ . Then,  $A \models \theta, p_1 \Rightarrow q_1$  and we assumed  $p_2 \in A$ , so we can invoke our primary assumption to get  $q_2 \in A$ . Since  $A$  is upward-closed, this gives  $q_2 \vee p_1 \in A$ .

**Theorem 5 (Correctness of Acts-for).** *Algorithm 1 is sound and complete with respect to definition 2:  $\theta \vdash p \leq q \iff \theta \models p \leq q$ .*

*Proof.* We prove soundness by induction on the derivation of  $\theta \vdash p \leq q$ .

- Case  $\mathbb{P}$ -AXIOM.. We have  $p \Rightarrow q$  by inversion on the derivation. The result follows from corollary 1.
- Case  $\mathbb{P}$ -DELEGATION.. We have  $\theta = (\theta', p' \Rightarrow q')$ ,  $\theta' \vdash p \wedge q' \leq q$ , and  $\theta' \vdash p \leq q \vee p'$ . Induction gives  $\theta' \models p \wedge q' \leq q$  and  $\theta' \models p \leq q \vee p'$ . Let  $A$  be an attacker such that  $A \models \theta', p' \Rightarrow q'$  and  $p \in A$ . We need to show  $q \in A$ . If  $p' \in A$ , then  $q' \in A$  since  $A \models p' \Rightarrow q'$ . Since  $p, q' \in A$ , theorem 2 gives  $p \wedge q' \in A$ . The first induction hypothesis then gives  $q \in A$ . Otherwise,  $p' \notin A$ . The second induction hypothesis gives  $q \vee p' \in A$ . Theorem 2 then implies  $q \in A$  since we cannot have  $p' \notin A$  and  $q \notin A$  but  $q \vee p' \in A$ .

We prove completeness by induction on  $\theta$ .

- Case  $\theta = \cdot$ . Assume  $\cdot \models p \leq q$ . Then  $p \Rightarrow q$  by lemma 2. Applying rule  $\mathbb{P}$ -AXIOM, we get  $\cdot \vdash p \leq q$  as desired.
- Case  $\theta = \theta', p' \Rightarrow q'$ . Assume  $\theta', p' \Rightarrow q' \models p \leq q$ . By lemmas 3 and 4, we get  $\theta' \models p \wedge q' \leq q$  and  $\theta' \models p \leq q \vee p'$ . By induction, we get  $\theta' \vdash p \wedge q' \leq q$  and  $\theta' \vdash p \leq q \vee p'$ . We can now apply rule  $\mathbb{P}$ -DELEGATION to get  $\theta', p' \Rightarrow q' \vdash p \leq q$  as desired.

**Lemma 5.** *If  $\theta \models p_2 \wedge q_1 \leq q_2$  and  $\theta \models p_2 \leq p_1$ , then  $(\theta, p_1 \Rightarrow q_1) \models p_2 \leq q_2$ .*

*Proof.* Follows trivially from theorem 2 since attackers are closed under  $\wedge$ .

**Theorem 6 (Correctness of Min).** *Algorithm 2 is sound and complete with respect to definition 7.*

*Proof.* Minimal elements are unique and algorithm 2 always terminates, so it suffices to prove soundness. More concretely, whenever algorithm 2 derives  $\min_\theta(p) = p'$ , we need to show  $\theta \models p \leq q \iff p' \Rightarrow q$  for all  $q$ .

Let  $q \in \mathbb{P}$ . We proceed by induction on the derivation.

- Case MIN-BASE. We have  $\theta = (p_1 \Rightarrow q_1, \dots, p_n \Rightarrow q_n)$ ,  $\min_\theta(p) = p$ ,  $\text{join-prime}(p)$ , and  $\forall i \in [n]. p \not\Rightarrow p_i$ .
  - Case  $\Rightarrow$ . Assume  $\theta \models p \leq q$ . We apply lemma 4  $n$  times and lemma 2 once to get  $p \Rightarrow q \vee p_1 \vee \dots \vee p_n$ . Since  $\text{join-prime}(p)$  and  $p \not\Rightarrow p_i$  for all  $i$ , it must be the case that  $p \Rightarrow q$  as desired.
  - Case  $\Leftarrow$ . Immediate by corollary 1.
- Case MIN-PICK. We have  $\theta = (\theta', p' \Rightarrow q')$ ,  $\min_\theta(p) = \min_{\theta'}(p \wedge q')$ ,  $\text{join-prime}(p)$ , and  $p \Rightarrow p'$ .
  - Case  $\Rightarrow$ . Assume  $\theta \models p \leq q$ . We need to show  $\min_{\theta'}(p \wedge q') \Rightarrow q$ . Lemma 3 gives  $\theta' \models p \wedge q' \leq q$ . We can then apply the induction hypothesis to get the desired result.
  - Case  $\Leftarrow$ . Assume  $\min_{\theta'}(p \wedge q') \Rightarrow q$ . We need to show  $\theta \models p \leq q$ . The induction hypothesis gives  $\theta' \models p \wedge q' \leq q$ . Corollary 1 gives  $\theta' \models p \leq p'$ . Combining these with lemma 5 gives the desired result.
- Case MIN-FACTOR. We have  $p = p_1 \vee \dots \vee p_n$ ,  $\min_\theta(p) = \bigvee_{i \in [n]} \min_\theta(p_i)$ , and  $\forall i \in [n]. \text{join-prime}(p_i)$ .
  - Case  $\Rightarrow$ . Assume  $\theta \models p \leq q$ . We need to show  $\bigvee_{i \in [n]} \min_\theta(p_i) \Rightarrow q$ . By the definition of  $\vee$ , we have  $\theta \models p_i \leq q$  for all  $i \in [n]$ . The induction hypotheses then give  $\min_\theta(p_i) \Rightarrow q$  for all  $i \in [n]$ . Finally, we use the definition of  $\vee$  to get the desired result.
  - Case  $\Leftarrow$ . Assume  $\bigvee_{i \in [n]} \min_\theta(p_i) \Rightarrow q$ . We need to show  $\theta \models p \leq q$ . By the definition of  $\vee$ , we have  $\min_\theta(p_i) \Rightarrow q$  for all  $i \in [n]$ . The induction hypotheses then give  $\theta \models p_i \leq q$  for all  $i \in [n]$ . Finally, we use the definition of  $\vee$  to get  $\theta \models p_1 \vee \dots \vee p_n \leq q$ .

**Theorem 7 (Correctness of Uncompromised Label Check).** *Algorithm 3 is sound and complete with respect to theorem 4.*

*Proof.* We show soundness and completeness separately.

- Case Soundness. Assume  $\theta_c \vdash \min_{\theta_i}(q) \leq p$ . Pick  $r = \min_{\theta_i}(q)$ . By definition 7 (instantiated with  $\min_{\theta_i}(q) \Rightarrow \min_{\theta_i}(q)$ ),  $\theta_i \models q \leq \min_{\theta_i}(q)$ . By theorem 5,  $\theta_c \models \min_{\theta_i}(q) \leq p$ .
- Case Completeness. Assume there exists  $r \in \mathbb{P}$  such that  $\theta_i \models q \leq r$  and  $\theta_c \models r \leq p$ . By definition 7,  $\min_{\theta_i}(q) \Rightarrow r$ . By theorem 2,  $\theta_c \models \min_{\theta_i}(q) \leq r$ . By transitivity,  $\theta_c \models \min_{\theta_i}(q) \leq p$ . Finally, by theorem 5,  $\theta_c \vdash \min_{\theta_i}(q) \leq p$ .