

First-Order Logic for Flow-Limited Authorization

Andrew K. Hirsch
MPI-SWS*
akhirsch@mpi-sws.org

Pedro H. Azevedo de Amorim
Cornell University
pamorim@cs.cornell.edu

Ethan Cecchetti
Cornell University
ethan@cs.cornell.edu

Ross Tate
Cornell University
ross@cs.cornell.edu

Owen Arden
University of California, Santa Cruz
owen@soe.ucsc.edu

Abstract

We present the Flow-Limited Authorization First-Order Logic (FLAFOL), a logic for reasoning about authorization decisions in the presence of information-flow policies. We formalize the FLAFOL proof system, characterize its proof-theoretic properties, and develop its security guarantees. In particular, FLAFOL is the first logic to provide a non-interference guarantee while supporting all connectives of first-order logic. Furthermore, this guarantee is the first to combine the notions of non-interference from both authorization logic and information-flow systems. All the theorems in this paper are proven in Coq.

1 Introduction

Distributed systems often make authorization decisions based on private data, which a public decision might leak. Preventing such leakage requires nontrivial reasoning about the interaction between information flow and authorization policies [Bec10, ALM15, AM16]. In particular, the justification for an authorization decision can violate information-flow policies. To understand this concern, consider a social network where Bob can say that only his friends may view his photos, and that furthermore only his friends may know the contents of his friend list. If Alice is not on Bob's friend list and she is denied access to one of his photos, the denial leaks Bob's private information: that Alice is not on Bob's friend list. Worse, if Alice can indirectly determine what other principals are permitted to see Bob's photos, she could completely enumerate the friend list.

Reasoning about the interaction between information flow and authorization policies is challenging for several reasons. First, authorization logics and information-flow systems use different notions of trust. Information-flow systems tend to focus on tracking data dependencies by representing information-security policies as *labels* on data. They then represent trust as a *flows-to* relation between labels, which determines when one piece of data may safely influence another. In contrast, authorization logics tend to directly encode *delegations* between principals as a *speaks-for* relation. Such delegations are often all-or-nothing, where a delegating principal trusts any statements made by the trusted principal, although some logics (e.g., [HK00, BFG10, SWS11]) support restricting delegations to specific statements. Flows-to relations implicitly encode delegations while speaks-for relations implicitly encode permitted flows. To understand *how*, we must understand how these disparate notions of trust interact.

Both forms of trust serve to selectively constrain the communication that system components rely on to make secure authorization decisions. For example, in the social network example above, suppose Bob's security settings are recorded on server X , and his photos are stored on server Y . When Alice tries to view Bob's photo, server Y communicates with server X to determine if Alice is permitted to do so. Modeling this

*Work done while author was at Cornell University

communication is important because (1) the servers that Y communicates with influence its authorization decisions, and (2) communication can leak private information.

Describing the information security of authorization decisions such as the one above requires modifying typical authorization policies to include information flow. Information-flow systems are excellent at tracking when and what information one principal communicates to another, specifically by transferring data from one label to another. It is less clear when communications occur in authorization logics. A common approach [SWS11, LABW91, Aba06] simply models Alice delegating trust to Bob as Alice importing all of Bob’s beliefs.

Authorization logics *do*, however, excel at reasoning about beliefs. Authorization logics allow us to write *Alice says* φ , meaning that Alice believes formula φ . This *says* statement is itself a formula, so we can reason about what Bob believes Alice believes by nesting *says* formulae. Information flow, in contrast, has no notion of belief, and so cannot reason about principals’ beliefs about each others’ beliefs.

In order to express authorization policies, not only does one need the ability to express trust and communication, but also a battery of propositions and logical connectives. Any tool that combines authorization and information flow should be capable of expressing enough logical connectives to reason about real-world policies. First-order logic seems to be a sweet spot of expressive power: it can encode most authorization policies, but it is still simple enough to have clean semantics. For instance, Nexus [SWS11, SBR⁺11]—a distributed operating system that uses authorization logic directly in its authorization mechanisms—can encode all of its authorization policies using first-order logic.¹

Finally, evaluating any attempt to combine authorization and information flow policies must examine the resulting security guarantees. Both authorization logics and information-flow systems have a security property called *non-interference*. Information-flow systems view non-interference as standard, while authorization logics often view it as desirable but unobtainable. Although the two formulations look quite different, both make guarantees limiting how one component of a system can influence—i.e., interfere with—another. In authorization logics, this takes the form “Alice’s beliefs can only impact the provability of Bob’s beliefs if Bob trusts Alice.” In information-flow systems—which are mostly defined over programs—changing the value of an input variable x can only change the value of an output variable y when the label of x flows to the label of y .

Both of these notions of non-interference are important. Consider again the example where Bob’s friend list is private but Alice attempts to view his photo. Because Bob’s friend list is private, changing the list should not affect Alice’s beliefs. For instance, Alice should not be affected by Bob adding or removing Cathy. To enforce this, whether or not Cathy is Bob’s friend must not affect the set of Bob’s beliefs that Alice *may* learn. This requires authorization-logic non-interference, since Bob’s beliefs should not affect Alice’s beliefs unless they communicate. It also, however, requires information-flow non-interference, since the privacy of Bob’s belief is why he is unwilling to communicate.

Gluing together both ideas of non-interference requires understanding the connection between their notions of trust. As we have discussed, this connection is difficult to formulate, making the non-interference combination harder still.

Our goal in this work is to provide a logic that supports reasoning about both information flow and authorization policies by combining their models of trust to obtain the advantages of both. To this end, we present the *Flow-Limited Authorization First-Order Logic* (FLAFOL), which

- provides a notion of trust between principals that can vary depending on information-flow labels,
- clearly denotes points where communication occurs,
- uses *says* formulae to reason about principals’ beliefs, including their beliefs about others’ beliefs,

¹The Nexus Authorization Logic is actually a monadic second-order logic, but this is used only to encode *speaksfor*; their examples only use first-order quantification [SWS11].

- is expressive enough to encode real-world authorization policies, and
- provides a strong security guarantee which combines both authorization-logic and information-flow non-interference.

We additionally aim to clarify the foundations of flow-limited authorization (introduced by Arden et al. [ALM15]). We therefore strive to keep FLAFOL’s model of principals, labels, and communication as simple as possible. For example, unlike previous work, we do not require that labels form a lattice.

A final contribution is an implementation of FLAFOL in Coq [Coq04] and formal proofs of all theorems in this paper.² Together these consists of 18,384 lines of Coq code. For more details, see Appendix C.

We are, of course, not the first to recognize the important interaction of information-flow policies with authorization, but all prior work in this area is missing at least one important feature. The three projects that have done the most to combine authorization and information flow are FLAM [ALM15], SecPAL⁺ [BFG10, Bec10], and AURA [JVM⁺08, JZ09]. FLAM models trust using information flow, AURA uses DCC [ABHR99, Aba06], a propositional authorization logic, and SecPAL⁺ places information flow labels on principal-based trust policies, but does not attempt to reason about the combination at all. Neither FLAM nor SecPAL⁺ can reason about nested beliefs, and both are significantly restricted in what logical forms are allowed. Finally, FLAM’s security guarantees are non-standard and difficult to compare to other languages (see Section 8), while AURA relies on DCC’s non-interference guarantee which does not apply on any trust relationships outside of those assumed in the static lattice.

The rest of this paper is organized as follows: In Section 2 we discuss three running examples. This also serves as an intuitive introduction to FLAFOL. In Section 3 we show how FLAFOL’s parameterization allows it to model real systems. In Section 4 we detail the FLAFOL proof rules. In Section 5 we discuss the proof theory of FLAFOL, proving important meta-level theorems, including consistency and cut elimination. In Section 6 we provide FLAFOL’s non-interference theorem. We discuss related work in Section 8, and finally we conclude in Section 9.

2 FLAFOL By Example

We now examine several examples of authorization policies and how FLAFOL expresses them. This will serve as a gentle introduction to the main ideas of FLAFOL, and introduce notation and running examples we use throughout the paper.

We explore three main examples in this section:

1. Viewing pictures on social media
2. Sanitizing data inputs to prevent SQL injection attacks
3. Providing a hospital bill in the presence of reinsurance

Each setting has different requirements, such as defining the meaning of labels in its own way. The ability of FLAFOL to adapt to each demonstrates its expressive power. In a new setting, it is often convenient—even necessary—to define constants, functions, and relations beyond those baked into FLAFOL. FLAFOL supports this by being parameterized over such definitions and having a security guarantee which holds for any parameterization. We use such symbols freely in our examples to express our intent clearly. Formally, FLAFOL interprets them using standard proof-theoretic techniques, as we see in Section 3.

Notably, FLAFOL does not allow computation on terms, so the meaning of functions and constants are axiomatized via FLAFOL formulae. This allows principals to disagree on how functions behave, which can be useful in modeling situations where each principal has their own view of some piece of data.

²The Coq code is available at <https://github.com/FLAFOL/flafol-coq>.

2.1 Viewing Pictures on Social Media

We begin by reconsidering in more detail the example from Section 1 where Alice requests to view Bob’s picture on a social-media service. This service allows Bob to set privacy policies, and Bob made his pictures visible only to his friends. When Alice makes her request, the service can check if she is authorized by scanning Bob’s friend list. If she is on the list and the photo is available, it shows her the photo. If she is *not* on Bob’s friend list, it shows her HTTP 403: Forbidden.

Bob may choose who belongs in the role of “friend.” Following the lead of other authorization logics, FLAFOL represents Bob believing that Alice is his friend as **Bob says** $\text{IsFriend}(\text{Alice})$. Since **says** statements can encompass any formula, we can express the fact that Bob believes that Alice is *not* his friend as **Bob says** $\neg\text{IsFriend}(\text{Alice})$.

We interpret these statements as Bob’s *beliefs*. This reflects the fact that Bob could be wrong, in the sense that he may affirm formulae with provable negations. There is no requirement that Bob believes all true things nor that Bob only believe true things (see Section 4), so holding an incorrect belief does not require Bob to believe **False**. Note that because **False** allows us to prove anything, a principal who *does* believe **False** will affirm every statement.

Now imagine that, as in Section 1, the social-media service allows Bob to set a privacy policy on his friend list as well. As before, Bob can restrict his friend list so that only his friends may learn its contents. In order to discuss such a policy in FLAFOL, we need a way to express that Bob’s friend list is private. Since, formally, his friend list is a series of beliefs about who his friends are, we must express the privacy of those beliefs. We view this as giving each belief a *label* describing Bob’s policy about who may learn that belief. Syntactically, we attach this label to the **says** connective. For example, Bob may use the label **Friends** to represent the information-security policy “I will share this with only my friends.”

If he attaches this policy to the beliefs representing his friend list, there is no way to securely prove either **Bob says** _{ℓ} $\text{IsFriend}(\text{Alice})$ or **Bob says** _{ℓ} $\neg\text{IsFriend}(\text{Alice})$ when ℓ is less restrictive than **Friends**. To see why, imagine what happens when Alice makes her request. If she is on Bob’s friend list, she may again see the photo. However, if she is not, showing her an HTTP 403 page would leak Bob’s private information; Alice would learn that she is not on Bob’s friend list, something Bob only shared with his friends. Since FLAFOL’s security guarantee (Theorem 7) shows that every FLAFOL proof is secure, neither option is provable in FLAFOL. Clearly Bob needs to define a more permissive policy on his friend list.

If Bob’s friend list were public, simply checking the list would be enough to prove either of the above statements. FLAFOL can easily express this by labeling each of Bob’s beliefs about IsFriend as **Public**. Another, more subtle, change would be to say that every principal can find out whether *they* are on Bob’s friend list, but only Bob’s friends can see the rest of the list. FLAFOL can also express this policy and prove it decidable, but doing so will require significant infrastructure using the technology we will build in Sections 3 and 4. We show how to express this policy in Appendix A.1.

This example demonstrates how naively reasoning about authorization with information flow can cause leaks, and how FLAFOL can help reason about those beliefs, leading to enforceable policies that capture the intent of system developers.

2.2 Integrity Tracking to Prevent SQL Injection

For our second example, imagine a stateful web application. It takes requests, updates its database, and returns web pages. In order to avoid SQL injection attacks, the system will only update its database based on high-integrity input. However, it marks all web request inputs as low integrity, representing the fact that they may contain attacks. The server has a sanitization function **San** that will neutralize attacks, so when it encounters a low-integrity input, it is willing to sanitize that input and endorse the result.

FLAFOL’s support for arbitrary implications allows it to easily encode such endorsements. Let the

predicate $\text{DBInput}(x)$ mean that a value x —possibly taken from a web request—is a database input. When a user makes a request with database input x , we can thus represent it as **System says**_{LInt} $\text{DBInput}(x)$. Here LInt represents low-integrity beliefs. We represent the system’s willingness to endorse any sanitized input as:

$$\text{System says}_{\text{LInt}} \text{DBInput}(x) \rightarrow \text{System says}_{\text{HInt}} \text{DBInput}(\text{San}(x))$$

This example shows the power of arbitrary implications for expressing authorization and information-flow policies. It also, however, demonstrates their dangers, since unconstrained downgrades can allow information to flow in unintended ways. In Section 6 we will discuss how non-interference (Theorem 7) adapts to these downgrades by weakening its guarantees.

2.3 Hospital Bills Calculation and Reinsurance

Imagine now that Alice finds herself in the hospital. Luckily her employer provides health insurance, but they have just switched companies. Now she has two unexpired insurance cards, and she cannot figure out which one is valid. Thus, either of two insurers, I_1 and I_2 , may be paying.

Imagine further that Bob’s job is to create a correct hospital bill for Alice. He uses the label ℓ_H to determine both who may learn the contents of Alice’s bill and who may help determine them. That is, ℓ_H expresses both a confidentiality policy and an integrity policy. Bob believes that Alice’s insurer may help determine the contents of Alice’s bill, since they can decide what they are willing to pay for Alice’s surgery.

Bob knows that I_2 has a reinsurance contract with I_1 . This means that if Alice is insured with I_2 and the surgery is very expensive, I_1 will pay some of the bill. Thus, I_1 may help determine the contents of Alice’s hospital bill, even if I_2 turns out to be her current insurer.

Bob is willing to accept Alice’s insurance cards as evidence that she is insured by either I_1 or I_2 , which we can express as **Bob says** _{ℓ_H} $(\text{CanWrite}(I_1, \ell_H) \vee \text{CanWrite}(I_2, \ell_H))$. Because Bob knows about I_2 ’s reinsurance contract with I_1 , he knows that if I_2 helps determine the contents of Alice’s bill, they will delegate some of their power to I_1 , which we express as **Bob says** _{ℓ_H} $(I_2 \text{ says}_{\ell_H} \text{CanWrite}(I_1, \ell_H))$.

Bob’s beliefs allow him to prove that I_1 may help determine the contents of Alice’s bill, since by assuming the previous two statements we can prove that **Bob says** _{ℓ_H} $\text{CanWrite}(I_1, \ell_H)$. There are two possible cases: if Bob already believes that I_1 can help determine the contents of Alice’s bill, we are done. Otherwise, Bob believes that I_2 can help determine the contents of Alice’s bill, and so Bob is willing to let I_2 delegate their power. Since he knows that they will delegate their power to I_1 , he knows that I_1 can help determine the contents of Alice’s bill in this case as well. This covers all of the cases, so we can conclude that **Bob says** _{ℓ_H} $\text{CanWrite}(I_1, \ell_H)$.

We think of Bob as performing this proof, since it is entirely about Bob’s beliefs. From this point of view, Bob’s ability to reason about I_2 ’s beliefs appears to be Bob *simulating* I_2 . This ability of one principal to simulate another provides the key intuition to understand the *generalized principal*, a fundamental construct in the formal presentation of FLAFOL (see Section 3).

We also note that Bob used I_2 ’s beliefs in this proof, even though he does not necessarily trust I_2 . However, he *might* trust it if it turns out to be Alice’s insurer. Because Bob trusts I_2 in part of the proof but not in general, we refer to this as *discoverable trust*. FLAFOL’s ability to handle discoverable trust makes reasoning about its security properties more difficult, as we see in Section 6.

This example shows how disjunctions can be used to express policies when principals do not know the state of the world. It also demonstrates how disjunctions make it difficult to know how information can flow at any point in time, since we may discover new statements of trust under one branch of a disjunction. FLAFOL’s non-interference theorem adapts to this by considering all declarations of trust that could possibly be discovered in a given context.

2.4 Further Adapting FLAFOL

All of the above examples use information-flow labels to express confidentiality policies, integrity policies, or both. While confidentiality and integrity are mainstay features of information flow tracking, information-flow labels can also express other properties. For instance, MixT [MM18] describes how to use information-flow labels to create safe transactions across databases with different consistency models, and the work of Zheng and Myers [ZM05] uses information-flow labels to provide availability guarantees. FLAFOL allows such alternative interpretations of labels by using an abstract *permission model* to give meaning to labels.

By default, the permissions gain meaning only through their behavior in context, but they are able to encode and reason about a wide variety of authorization mechanisms. In Section 3, we see how FLAFOL can be used to reason about capabilities, and in Appendix B we discuss a model closer to military classification.

3 Using FLAFOL

In this section, we examine how to use FLAFOL to reason about real systems. To do this, we look at a fictional verified-distributed-systems designer Dana. She wants to formally prove that confused-deputy attacks are impossible in her capability-based system with copyable, delegatable read capabilities. Dana employs a six-step process to reason about her system in FLAFOL:

1. Decide on a set \mathcal{S} of *sorts* of data she wants to represent.
2. Choose a set \mathcal{F} of *function symbols* representing operations in the system, and give those operations types.
3. Choose a set \mathcal{R} of *relation symbols* representing atomic facts to reason about, and give the relations types.
4. Develop axioms that encode meaning for these relationships.
5. Specify meta-level theorems stating her desired properties.
6. Prove that those meta-level theorems hold.

Sorts. First, Dana decides on what sorts of data she wants to represent. We can think of *sort* as the logic word for “type.” FLAFOL is defined with respect to a set \mathcal{S} of sorts that must include at least **Label** and **Principal**, but may contain more. Dana wants to reason about capability tokens that grant read access to data, so she also includes a sort named **Token**.

Dana uses the **Principal** sort to represent system principals, but conceptually divides the **Label** sort into **Confidentiality** and **Integrity**, two sorts which she also adds. Each **Confidentiality** value defines a confidentiality policy which may be applied to many pieces of data. A capability (which is always public itself) grants read access to data governed by one or more such policies. She uses the **Integrity** sort to represent integrity policies on tokens themselves. We will see below how she can enforce **Label** = **Confidentiality** \times **Integrity**.

Function Symbols. Dana next decides on operations she wants to reason about. This is also her chance to define constants using nullary operations. Formally, FLAFOL is defined with respect to an arbitrary set \mathcal{F} of *function symbols*. Each function comes equipped with a *signature*, or type, expressing when it can be applied.

Dana considers what information she needs about a given token. She needs a way to determine which confidentiality level a token grants permission to read, the integrity of that token, and which principal is the token’s *root of authority*—that is, who created the token. She thus creates three function symbols:

$$\begin{aligned}\text{TknConf} &: \text{Token} \rightarrow \text{Confidentiality} \\ \text{IntegOfTkn} &: \text{Token} \rightarrow \text{Integrity} \\ \text{RootOfAuth} &: \text{Token} \rightarrow \text{Principal}\end{aligned}$$

She also needs to be able to determine the integrity that a principal commands, so she includes a function symbol $\text{IntegOf} : \text{Principal} \rightarrow \text{Integrity}$. Finally, since a token can potentially be transferred to anyone in her system, she creates a constant $\text{Public} : \text{Confidentiality}$ to represent this.

Dana wants to enforce that labels are pairs of confidentiality and integrity. She therefore creates two “projection” function symbols π_C and π_I , and a third pair symbol $(-, -)$ with the following signatures:

$$\begin{aligned}\pi_C &: \text{Label} \rightarrow \text{Confidentiality} \\ \pi_I &: \text{Label} \rightarrow \text{Integrity} \\ (-, -) &: \text{Confidentiality} \rightarrow \text{Integrity} \rightarrow \text{Label}\end{aligned}$$

The first two ensure that labels contain a confidentiality and an integrity, while pairing allows creation of labels from a confidentiality with an integrity. This makes labels pairs of confidentiality and integrity. Dana also adds axioms corresponding to the η and β laws for pairs.

Relation Symbols. Dana can now choose relations representing facts that she wants to reason about. Along with sorts and functions, FLAFOL is defined with respect to a set \mathcal{R} of *relation symbols*, allowing it to reason about more facts. The set \mathcal{R} must include at least flows-to (\sqsubseteq), CanRead , and CanWrite , but may contain more. We call these required relations *permissions* because they define the trust relationships governing communication. The relation $\ell \sqsubseteq \ell'$ means information with label ℓ can affect information with label ℓ' , $\text{CanRead}(p, \ell)$ means that principal p may learn beliefs with label ℓ , and $\text{CanWrite}(p, \ell)$ means p may influence beliefs with label ℓ .

Dana is able to use these relations to define the permissions her capability tokens grant. She also includes a fourth relation in \mathcal{R} , $\text{HasToken}(\text{Principal}, \text{Token})$, defining token possession: if $\text{HasToken}(p, t)$, then principal p has (a copy of) token t .

Axioms. Dana describes the behavior of her system with axioms that use the sorts, functions, and relations she defined above. These should be *consistent*, in the sense that they do not allow a derivation of False . Theorem 2 in Section 5.1 gives conditions under which all of the axioms that we will discuss in this section are consistent.

Dana uses three main axioms: one describing how tokens may be copied and delegated, one describing when one principal may read another’s beliefs, and one describing when a principal may affect another’s beliefs. She may use more axioms if she likes—e.g., to capture principals’ beliefs about permitted flows between labels.

Dana’s first axiom allows any principal to copy any capability it holds and give that copy to another principal:

$$\forall q : \text{Principal}. \forall t : \text{Token}. \left(\begin{array}{l} \exists p : \text{Principal}. \text{HasToken}(p, t) \\ \wedge p \text{ says}_{(\text{Public}, \text{IntegOfTkn}(t))} \text{HasToken}(q, t) \end{array} \right) \rightarrow \text{HasToken}(q, t)$$

This says that, for principals p and q , if p holds a read capability token t , p can pass t to q . To do so, p must affirm that q has t at a public label with the integrity of the token. Note that the use of Public here means Dana’s system must allow everyone to learn whenever one principal copies a token and passes it to another.

Dana’s second axiom defines when a principal p allows q to read a belief of p ’s labeled ℓ . First, p checks that q has a token, and that p believes that the token gives read access to something at least as confidential as ℓ . Second, p checks to make sure that the token’s root authority may influence this belief:

$$\forall q : \text{Principal}. \forall \ell : \text{Label}. \forall p : \text{Principal}. \forall \ell' : \text{Label}. \left(\begin{array}{l} \exists t : \text{Token}. \text{HasToken}(q, t) \\ \wedge p \text{ says}_{\ell'} \pi_C(\ell) \sqsubseteq \text{TknConf}(t) \\ \wedge p \text{ says}_{\ell'} \text{CanWrite}(\text{RootOfAuth}(t), \ell') \end{array} \right) \rightarrow p \text{ says}_{\ell'} (\text{CanRead}(q, \ell))$$

Sorts	σ	::=	Label Principal ...
Labels	ℓ		
Principals	p, q, r		
Functions	f	::=	...
Relations	R	::=	CanRead(Principal, Label) CanWrite(Principal, Label) Label \sqsubseteq Label ...
σ -terms	t	::=	x $f(t_1, \dots, t_n)$
Formulae	φ, ψ, χ	::=	$R(t_1, \dots, t_n)$ True False $\varphi \wedge \psi$ $\varphi \vee \psi$ $\varphi \rightarrow \psi$ $\forall x:\sigma. \varphi$ $\exists x:\sigma. \varphi$ $p \text{ says}_\ell \varphi$
Generalized Principals	g	::=	$\langle \rangle$ $g \cdot p(\ell)$

Figure 1: FLAFOL Syntax

More formally, it says that if q holds some token t and p believes both that t grants read permissions for ℓ 's confidentiality and that the root of authority for t can influence p 's beliefs at ℓ' , then p will allow q to read ℓ . This defines what it means for a principal (p here) to believe that a token grants read access to their data. Dana now needs to make sure that whenever a read access is granted in her system, not only does the principal who gets read access have a token, but that the principal who owns the data does indeed believe that the token grants read access to that data.

Finally, her third axiom states that one principal p believes that another principal q , can write a label ℓ if p believes that the integrity of q flows to the integrity of ℓ :

$$\forall q:\text{Principal}. \forall \ell:\text{Label}. \forall p:\text{Principal}. \forall \ell':\text{Label}. \\ p \text{ says}_{\ell'} (\text{IntegOf}(q) \sqsubseteq \pi_I(\ell)) \rightarrow p \text{ says}_{\ell'} (\text{CanWrite}(q, \ell))$$

Dana then needs to make sure that write accesses are only granted to principals with high enough integrity.

Metatheoretic Properties. Dana has now created a model of her system, so she can use it to state and prove properties of her system as meta-theorems. Luckily, Rajani, Garg, and Rezk [RGR16] have shown that information-flow integrity tracking with a non-interference result is sufficient to avoid confused deputy attacks with capability systems. Therefore Theorem 7 provides the guarantees she needs.

FLAFOL Syntax. This example demonstrates FLAFOL's flexibility as a powerful tool for reasoning about authorization mechanisms in the presence of information-flow policies. We saw that, since FLAFOL is defined with respect to the three sets \mathcal{S} , \mathcal{F} , and \mathcal{R} , it can express the key components of a system. This parameterized definition gives rise to the formal FLAFOL syntax in Figure 1.

In order to use the function and relation symbols and incorporate axioms, FLAFOL allows proofs to occur in a context. FLAFOL additionally includes rules requiring flows-to to be reflexive and transitive, placing a preorder on the Label sort,³ and requiring CanRead and CanWrite to respect a form of variance. If $\ell_1 \sqsubseteq \ell_2$ and Alice can read data A with label ℓ_2 , then she may learn information about data with label ℓ_1 used to calculate A . This means she should also be able to read data with label ℓ_1 . Thus, CanRead must (contravariantly) respect the preorder on labels. Similarly, if Alice can help determine some piece of data B

³Many information-flow tools require their labels to form a lattice. We find that a preorder is sufficient for FLAFOL's design and guarantees, so we decline to impose additional structure. In Section 5.1 we show that enforcing a lattice structure is both simple and logically consistent.

$$\begin{array}{c}
\text{FLOWSTOREFL} \frac{}{\Gamma \vdash \ell \sqsubseteq \ell @ g} \qquad \text{FLOWSTOTRANS} \frac{\Gamma \vdash \ell_1 \sqsubseteq \ell_2 @ g \quad \Gamma \vdash \ell_2 \sqsubseteq \ell_3 @ g}{\Gamma \vdash \ell_1 \sqsubseteq \ell_3 @ g} \\
\text{CRVAR} \frac{\Gamma \vdash \text{CanRead}(p, \ell_2) @ g \quad \Gamma \vdash \ell_1 \sqsubseteq \ell_2 @ g}{\Gamma \vdash \text{CanRead}(p, \ell_1) @ g} \\
\text{CWVAR} \frac{\Gamma \vdash \text{CanWrite}(p, \ell_1) @ g \quad \Gamma \vdash \ell_1 \sqsubseteq \ell_2 @ g}{\Gamma \vdash \text{CanWrite}(p, \ell_2) @ g}
\end{array}$$

Figure 2: Permission Rules

labeled with ℓ_1 , she can influence any data labeled with ℓ_2 that is calculated from B , so Alice should be able to help determine data labeled at ℓ_2 . Thus, **CanWrite** must (covariantly) respect the preorder on labels.

Figure 2 presents these rules formally. We give the proof rules in the form of a sequent calculus. The trailing $@ g$ represents *who* affirms that formula in the proof, similarly to how **says** formulae represent who affirms a statement at the object level. Unlike **says** formulae, these meta-level objects—which we call *generalized principals*—encode arbitrary reasoners, including possibly-simulated principals.

Recall from Section 2.3 that we can think of some proofs as being performed by principals if those proofs entirely involve that principal’s beliefs. In that example, Bob reasoned about his belief that another principal, the insurer I_2 , trusted a third principal, the insurer I_3 . We think of this ability to reason about the beliefs of others as the ability to *simulate* other principals. In fact, because principals’ beliefs are segmented by labels, principals can have multiple simulations of the same other principal.

This suggests that FLAFOL captures the reasoning of principals *at some level of simulation*. A generalized principal is a stack of principal/label pairs, representing a stack of simulators and simulations. The empty stack, written $\langle \rangle$, represents *ground truth*. A stack with one more level, written $g \cdot p \langle \ell \rangle$, represents the beliefs of p at level ℓ according to the generalized principal g . Figure 1 contains the formal grammar for generalized principals.

4 Proof System

So far, we have discussed the intuitions behind FLAFOL and its syntax. Here we introduce FLAFOL formally. Unfortunately, we cannot examine every aspect of FLAFOL’s formal presentation in detail, though interested readers should see Appendix D. Instead, we discuss the most novel and most security-relevant aspects of FLAFOL’s design.

FLAFOL sequents are of the form $\Gamma \vdash \varphi @ g$, where Γ is a context containing beliefs. This means that the FLAFOL proof system manipulates beliefs, as described in Section 3. Readers familiar with sequent calculus may recognize that FLAFOL is intuitionistic, as there is only one belief on the right side of the turnstile.⁴

Sequent calculus rules tend to manipulate beliefs either on the left or the right side of the turnstile. For instance, consider the FLAFOL rules for disjunctions:

$$\text{ORL} \frac{\Gamma, \varphi @ g \vdash \chi @ g' \quad \Gamma, \psi @ g \vdash \chi @ g'}{\Gamma, (\varphi \vee \psi @ g) \vdash \chi @ g'} \qquad \text{ORR1} \frac{\Gamma \vdash \varphi @ g}{\Gamma \vdash \varphi \vee \psi @ g} \qquad \text{ORR2} \frac{\Gamma \vdash \psi @ g}{\Gamma \vdash \varphi \vee \psi @ g}$$

⁴Recall that we argued in Section 2.1 that reasoning about authorization and information-flow security together is naturally intuitionistic, since we cannot securely conclude φ or $\neg\varphi$ in some naturally-occurring contexts.

$$\begin{array}{c}
\text{FALSEL} \frac{}{\Gamma, \text{False} @ g \vdash \varphi @ g \cdot g'} \\
\text{IMPL} \frac{\Gamma \vdash \varphi @ \langle \rangle \quad \Gamma, \psi @ g \vdash \chi @ g'}{\Gamma, (\varphi \rightarrow \psi @ g) \vdash \chi @ g'} \qquad \text{IMPR} \frac{\Gamma, \varphi @ \langle \rangle \vdash \psi @ g}{\Gamma \vdash \varphi \rightarrow \psi @ g} \\
\text{SAYSL} \frac{\Gamma, \varphi @ g \cdot p \langle \ell \rangle \vdash \psi @ g'}{\Gamma, p \text{ says}_\ell \varphi @ g \vdash \psi @ g'} \qquad \text{SAYSR} \frac{\Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle}{\Gamma \vdash p \text{ says}_\ell \varphi @ g} \\
\text{VARR} \frac{\Gamma \vdash \varphi @ g \cdot p \langle \ell' \rangle \cdot g' \quad \Gamma \vdash \ell' \sqsubseteq \ell @ g \cdot p \langle \ell \rangle}{\Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle \cdot g'} \qquad \text{FWDR} \frac{\Gamma \vdash \text{CanRead}(g, \ell) @ g \cdot p \langle \ell \rangle \quad \Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle \cdot g' \quad \Gamma \vdash \text{CanWrite}(p, \ell) @ g \cdot q \langle \ell \rangle}{\Gamma \vdash \varphi @ g \cdot q \langle \ell \rangle \cdot g'}
\end{array}$$

Figure 3: Selected FLAFOL Proof Rules

We find it easiest to read left rules “up” and right rules “down.” With this reading, the ORL rule tells us how to use an assumption of the form $\varphi \vee \psi @ g$ in order to prove a belief $\chi @ g'$ by performing case analysis. That is, ORL tells us how to prove $\chi @ g'$ assuming $\varphi \vee \psi @ g$ if we can prove that $\chi @ g'$ assuming $\varphi @ g$ and separately assuming $\psi @ g$.

The ORR1 rule takes a proof of $\varphi @ g$ and uses it to prove $\varphi \vee \psi @ g$. The ORR2 rule is symmetric, so it takes a proof of $\psi @ g$ and uses it to prove $\varphi \vee \psi @ g$.⁵

Note that these rules (along with the **says** rules discussed below) allow **says** to distribute over disjunctions. That is, given $p \text{ says}_\ell (\varphi \vee \psi)$, we can prove $(p \text{ says}_\ell \varphi) \vee (p \text{ says}_\ell \psi)$. In an intuitionistic logic like FLAFOL, disjunctions must be a proof of one side or the other. The proof that **says** distributes over \vee then says that if p has evidence of either φ or ψ , then p can examine this evidence to discover whether it is evidence of φ or of ψ .

Most of the rules of FLAFOL are standard rules for first-order logic, but with generalized principals included to indicate who believes each formula. For instance, the rules for disjunctions above were likely familiar to those who know sequent calculus.

Figure 3 contains FLAFOL rules selected for discussion. The first, FALSEL, tells us how to use **False** as an assumption. In standard intuitionistic first-order logic, this is simply the principle of Ex Falso: if we assume **False**, we can prove anything. In FLAFOL, a generalized principal who assumes false is willing to affirm any formula. This includes statements about other principals, so FALSEL extends the generalized principal arbitrarily. We use $g \cdot g'$ as notation for extending the generalized principal g with a list of principal-label pairs, denoted g' .

The implication rules IMPR and IMPL interpret the premise of an implication as ground truth, while the generalized principal who believes the implication believes the consequent. In particular, this means that **says** statements do not distribute over implication as one might expect, i.e., $p \text{ says}_\ell (\varphi \rightarrow \psi)$ does not imply that $(p \text{ says}_\ell \varphi) \rightarrow (p \text{ says}_\ell \psi)$. Instead, $p \text{ says}_\ell (\varphi \rightarrow \psi)$ implies $\varphi \rightarrow (p \text{ says}_\ell \psi)$. We can thus think of implications as *conditional* knowledge. That is, if a generalized principal g believes $\varphi \rightarrow \psi$, then g believes ψ conditional on φ being true about the system.

We can still form implications about generalized principals’ beliefs, but we must insert appropriate **says** statements into the premise to do so. In Section 5.5, we discuss how this semantics is necessary for both our proof theoretic and our security results.

The next two rules of Figure 3, SAYSR and SAYSL, are the only rules which specifically manipulate

⁵For readers interested in learning more about sequent calculus, we recommend MIT’s interactive tool for teaching sequent calculus as a tutorial [Yan12].

says formulae. Essentially, generalized principals allow us to delete the **says** part of a formula while not forgetting who said it. Thus, generalized principals allow us to define sequent calculus rules once for every possible reasoner.

The final rules, VARR and FWDR, define communication in FLAFOL. Both manipulate beliefs on the right and have corresponding left rules, which act contravariantly and can be found in Appendix D.

Information-flow communication is provided by the variance rule VARR. This can be thought of like the variance rules used in subtyping. Most systems with information-flow labels do not have explicit variance rules, but instead manipulate relevant labels in every rule. By adding an explicit variance rule, we not only simplify every other FLAFOL rule, we also remove the need for the label join and meet operators that are usually used to perform the label manipulations. Others have noted that adding explicit variance rules simplifies the design of the rest of the system [VSI96, Alg18], but it remains an unusual choice.

The forwarding rule FWDR provides authorization-logic-style communication. In FLAFOL, p can forward a belief at label ℓ to q if:

- p is willing to send its beliefs at label ℓ to q , denoted $p \text{ says}_\ell \text{ CanRead}(q, \ell)$, and
- q is willing to allow p to determine its beliefs at label ℓ , denoted $q \text{ says}_\ell \text{ CanWrite}(p, \ell)$.

After establishing this trust, p can package up its belief and send it to q , who will believe it at the same label.

5 Proof Theory

In this section, we evaluate FLAFOL's logical design. We show that FLAFOL has the standard sequent calculus properties of (positive) consistency and cut elimination and discuss fundamental limitations that inform our unusual implication semantics. We also develop a new proof-theoretic tool, *compatible supercontexts*, for use in our non-interference theorem in Section 6.

5.1 Consistency

One of the most important properties about a logic is consistency, meaning it is impossible to prove **False**. This is not possible in an arbitrary context, since one could always assume **False**. One standard solution is to limit the theorem to the empty context. By examining the FLAFOL proof rules, however, we see that it is only possible to prove **False** by assumption or by Ex Falso. Either method requires that **False** already be on the left-hand side of the turnstile, so if **False** can never get there, then it should be impossible to prove.

To understand when **False** can appear on the left-hand side of the turnstile, we note that formulae on the left tend to stay on the left and formulae on the right tend to stay on the right. The only exception is the implication rules IMPL and IMPR which move the premise of the implication to the other side. The fact that no proof rule allows us to change either side of the sequent arbitrarily gives useful structure to proofs. To handle implications, however, we must keep track of their nesting structure, which we do by considering *signed formulae*. We call a formula in a sequent *positive* if it appears on the right side of the turnstile and *negative* if it appears on the left. If φ is positive we write φ^+ , and if φ is negative we write φ^- . Figure 4 contains the rules for the *signed subformula relation*.

Note that every subformula of a signed formula has a unique sign. If a subformula appears by itself in a sequent during a proof, then which side of the turnstile it is on is determined by its sign. This structure results in the following formal property.

Theorem 1 (Left Signed-Subformula Property). *If $\Gamma \vdash \varphi @ g_1$ appears in a proof of $\Delta \vdash \psi @ g_2$, then for all $\chi_1 @ g_3 \in \Gamma$, either (1) $\chi_1^- \leq \psi^+$ or (2) there is some $\chi_2 @ g_4 \in \Delta$ such that $\chi_1^- \leq \chi_2^-$.*

This proof follows by induction on the FLAFOL proof rules.

$$\begin{array}{c}
s \in \{+, -\} \qquad \bar{+} = - \qquad \bar{-} = + \\
\frac{\overline{\varphi^s \leq \varphi^s}}{\overline{\varphi^s \leq \varphi^s}} \quad \frac{\overline{\varphi^s \leq \psi^{s'}} \quad \overline{\psi^{s'} \leq \chi^{s''}}}{\overline{\varphi^s \leq \chi^{s''}}} \quad \overline{\varphi^s \leq (\varphi \vee \psi)^s} \quad \overline{\psi^s \leq (\varphi \vee \psi)^s} \quad \overline{\varphi^s \leq (\varphi \wedge \psi)^s} \\
\overline{\psi^s \leq (\varphi \wedge \psi)^s} \quad \overline{\varphi^s \leq (\varphi \rightarrow \psi)^s} \quad \overline{\psi^s \leq (\varphi \rightarrow \psi)^s} \quad \overline{(\varphi[x \mapsto t])^- \leq (\forall x:\sigma. \varphi)^-} \\
\overline{\varphi^+ \leq (\forall x:\sigma. \varphi)^+} \quad \overline{\varphi^- \leq (\exists x:\sigma. \varphi)^-} \quad \overline{(\varphi[x \mapsto t])^+ \leq (\exists x:\sigma. \varphi)^+} \quad \overline{\varphi^s \leq (p \text{ says}_\ell \varphi)^s}
\end{array}$$

Figure 4: Signed Subformula Relation

Many logics also have a similar *right* signed-subformula property. FLAFOL does not enjoy that property since $\Gamma \vdash \varphi @ g_1$ may be a side condition on a forward or a variance rule, and thus not related directly to ψ .

This property allows us to prove an important result about the consistency of FLAFOL.

Theorem 2 (Positive Consistency). *For any context Γ , if*

$$\text{False}^- \not\leq \varphi^- \text{ for all } \varphi @ g \in \Gamma$$

then $\Gamma \not\vdash \text{False} @ g'$.

The proof follows by induction on the FLAFOL proof rules. Note that formulae which do not contain **False** as a negative subformula are called *positive* formulae, explaining the name.

We get the result with an empty context as a corollary. This states that **False** is not a theorem of FLAFOL.

Corollary 1 (Consistency). $\not\vdash \text{False} @ g$

Theorem 2 demonstrates that a variety of useful constructs are logically consistent. For instance, we can add a lattice structure to FLAFOL's labels. We can define join (\sqcup) and meet (\sqcap) as binary function symbols on labels and \top and \perp as label constants. Then we can simply place the lattice axioms (e.g., $\forall \ell: \text{Label}. \ell \sqsubseteq \top$) in our context to achieve the desired result. Since none of the lattice axioms include **False**, Theorem 2 ensures that they are consistent additions to the logic.

5.2 Simulation

In (multi-)modal logics, we are interested in modeling *perfect* reasoners. That is, reasoners should reason correctly based on their assumed beliefs; if their assumed beliefs were true, then all of their derived beliefs would be as well.

In most logics (which do not have generalized principals, this is axiomatized as a rule in the system, written as follows:

$$\frac{\Gamma \vdash \varphi}{p \text{ says}_\ell \Gamma \vdash p \text{ says}_\ell \varphi}$$

Here, $p \text{ says}_\ell \Gamma$ refers to a copy of Γ with $p \text{ says}_\ell$ in front of every formula in Γ . In such logics, this is the main rule for manipulating says statements. However, this requires removing all beliefs that are not those of p at level ℓ in a context before using this rule to reason as p at level ℓ .

FLAFOL instead uses the **says** introduction rules in Section 4, which allows us to retain the beliefs of other principals and of p at other labels, making it easier to discuss communication. This difference causes no harm. FLAFOL reasoners are still be perfect reasoners, which we show by proving a theorem analogous

to the above rule. We refer to this as the *simulation* theorem, since it says that p is correctly simulating the world in its head.

Adopting the above rule directly fails for two reasons. The first is that our belief syntax pushes **says** statements into generalized principals, so we must place the new principal-label pair at the beginning of the generalized principal instead of on the formula. The second is that the semantics of implications in FLAFOL mean that $p \text{ says}_\ell (\varphi \rightarrow \psi)$ has different semantics from $(p \text{ says}_\ell \varphi) \rightarrow (p \text{ says}_\ell \psi)$. To address this concern, we define the \odot operator:

$$p\langle\ell\rangle \odot \varphi \triangleq \begin{cases} (p \text{ says}_\ell (p\langle\ell\rangle \odot \psi)) \rightarrow (p\langle\ell\rangle \odot \chi) & \varphi = \psi \rightarrow \chi \\ (p\langle\ell\rangle \odot \psi) \wedge (p\langle\ell\rangle \odot \chi) & \varphi = \psi \wedge \chi \\ (p\langle\ell\rangle \odot \psi) \vee (p\langle\ell\rangle \odot \chi) & \varphi = \psi \vee \chi \\ \forall x:\sigma. (p\langle\ell\rangle \odot \psi) & \varphi = \forall x:\sigma. \psi \\ \exists x:\sigma. (p\langle\ell\rangle \odot \psi) & \varphi = \exists x:\sigma. \psi \\ \varphi & \text{otherwise} \end{cases}$$

This essentially “repairs” implications to have the right **says** statements in front of the premise.

Because FLAFOL can move **says** statements into generalized principals, we need to lift the operator to beliefs. In doing so, we must place $p\langle\ell\rangle$ at the *beginning* of the generalized principal, leading to the following definition:

$$p\langle\ell\rangle \odot (\varphi @ \langle \rangle \cdot g') \triangleq (p\langle\ell\rangle \odot \varphi) @ \langle \rangle \cdot p\langle\ell\rangle \cdot g',$$

From there we can lift the operator to contexts as well.

$$p\langle\ell\rangle \odot \Gamma \triangleq \begin{cases} \cdot & \Gamma = \cdot \\ (p\langle\ell\rangle \odot \Gamma'), p\langle\ell\rangle \odot (\varphi @ g) & \Gamma = \Gamma', \varphi @ g \end{cases}$$

With this definition in hand, we can now state the simulation theorem in full:

Theorem 3 (Simulation). *The following rule is admissible:*

$$\frac{\Gamma \vdash \varphi @ g}{(p\langle\ell\rangle \odot \Gamma) \vdash p\langle\ell\rangle \odot (\varphi @ g)}$$

5.3 Compatible Supercontexts

To prove Theorem 2 we needed to consider the possible locations of *formulae* within a sequent, but in Section 6 we will need to reason about the possible locations of *beliefs*. To enable this, we introduce the concept of a *compatible supercontext* (CSC). Informally, the CSCs of a sequent are those contexts that contain all of the information in the current context, along with any counterfactual information that can be considered during a proof. Intuitively, the rules ORL and IMPL allow a generalized principal to consider such information by using either side of a disjunction or the conclusion of an implication. If it is possible to consider such a counterfactual, there is a CSC which contains it. We use the syntax $\Delta \ll \Gamma \vdash \varphi @ g$ to denote that Δ is a CSC of the sequent $\Gamma \vdash \varphi @ g$. Figure 5 contains selected rules for CSCs. The full CSC relation can be found in Appendix E.

Since all of the information in Γ has already been discovered by the generalized principal who believes that information, we require that $\Gamma \ll \Gamma \vdash \varphi @ g$ with CSCREFL.

If we can discover two sets of information, we can discover everything in the union of those sets using CSCUNION. This rule feels different from the others, since it axiomatizes certain *properties* of CSCs. We conjecture that there is an alternative presentation of CSCs where we can prove this rule.

$$\begin{array}{c}
\text{CSCREFL} \frac{}{\Gamma \ll \Gamma \vdash \varphi @ g} \qquad \text{CSCUNION} \frac{\Delta_1 \ll \Gamma \vdash \varphi @ g \quad \Delta_2 \ll \Gamma \vdash \varphi @ g}{\Delta_1 \cup \Delta_2 \ll \Gamma \vdash \varphi @ g} \\
\text{CSCORL1} \frac{\Delta \ll \Gamma, \varphi @ g \vdash \chi @ g'}{\Delta \ll \Gamma, (\varphi \vee \psi @ g) \vdash \chi @ g'} \qquad \text{CSCIMPR} \frac{\Delta \ll \Gamma, \varphi @ \langle \rangle \vdash \psi @ g}{\Delta \ll \Gamma \vdash \varphi \rightarrow \psi @ g}
\end{array}$$

Figure 5: Selected Rules for Compatible Supercontexts

The rest of the rules for CSCs essentially follow the proof rules, so that any belief added to the context during a proof can be added to a CSC. For instance CSCORL1 and CSCORL2 allow either branch of an assumed disjunction to be added to a CSC, following the two branches of the ORL rule of FLAFOL.

If a context appears in a proof of a sequent, then it is a CSC of that sequent. We refer to this as the *compatible-supercontext property* (CSC property).

Theorem 4 (CSC Property). *If $\Delta \vdash \psi @ g'$ appears in a proof of $\Gamma \vdash \varphi @ g$, then $\Delta \ll \Gamma \vdash \varphi @ g$.*

5.4 Cut Elimination

In constructing a proof, it is often useful to create a lemma, prove it separately, and use it in the main proof. If we both prove and use the lemma in the same context, the main proof follows in that context as well. We can formalize this via the following rule:

$$\text{CUT} \frac{\Gamma \vdash \varphi @ g_1 \quad \Gamma, \varphi @ g_1 \vdash \psi @ g_2}{\Gamma \vdash \psi @ g_2}$$

This rule is enormously powerful. It allows us to not only create lemmata to use in a proof, but also simply prove things whose other proofs are complicated and non-obvious. For instance, consider the rule

$$\text{UNSAYSR} \frac{\Gamma \vdash p \text{ says}_\ell \varphi @ g}{\Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle}$$

We can show that this rule is *admissible*—meaning any sequent provable with this rule is provable without it—by cutting a proof of the sequent $\Gamma \vdash p \text{ says}_\ell \varphi @ g$ with the following proof:⁶

$$\text{SAYS�} \frac{\text{Ax} \frac{}{\varphi @ g \cdot p \langle \ell \rangle \vdash \varphi @ g \cdot p \langle \ell \rangle}}{p \text{ says}_\ell \varphi @ g \vdash \varphi @ g \cdot p \langle \ell \rangle}$$

However, the CUT rule allows an arbitrary formula to appear on both sides of the turnstile in a proof. That formula may not even be a subformula of anything in the sequent at the root of the proof-tree! This would seemingly destroy the CSC property that FLAFOL enjoys, and which we rely on in order to prove FLAFOL's security results. As is standard in sequent calculus proof theory, we show that CUT can be admitted, allowing FLAFOL the proof power of CUT while maintaining the analytic power of the CSC property.

Theorem 5 (Cut Elimination). *The CUT rule is admissible.*

⁶Not only can UNSAYSR be proven without CUT (as can all FLAFOL proofs), it is actually important for proving cut elimination. See the Coq code.

$$\begin{array}{c}
\text{AX} \frac{}{\varphi @ g \cdot p \langle \ell \rangle \vdash \varphi @ g \cdot p \langle \ell \rangle} \quad \text{AX} \frac{}{\psi @ g \cdot p \langle \ell \rangle \vdash \psi @ g \cdot p \langle \ell \rangle} \\
\text{SAYS L} \frac{}{p \text{ says}_\ell \varphi @ g \vdash \varphi @ g \cdot p \langle \ell \rangle} \quad \text{SAYS R} \frac{}{\psi @ g \cdot p \langle \ell \rangle \vdash p \text{ says}_\ell \psi @ g} \\
\text{IMPL}' \frac{}{(\varphi \rightarrow \psi) @ g \cdot p \langle \ell \rangle, p \text{ says}_\ell \varphi @ g \vdash p \text{ says}_\ell \psi @ g} \\
\text{IMPR}' \frac{}{(\varphi \rightarrow \psi) @ g \cdot p \langle \ell \rangle \vdash (p \text{ says}_\ell \varphi) \rightarrow (p \text{ says}_\ell \psi) @ g} \\
\text{SAYS L} \frac{}{p \text{ says}_\ell (\varphi \rightarrow \psi) @ g \vdash (p \text{ says}_\ell \varphi) \rightarrow (p \text{ says}_\ell \psi) @ g}
\end{array}$$

Figure 6: Proof that IMPL' and IMPR' allow **says** to distribute over implication.

$$\begin{array}{c}
\text{AX} \frac{}{\varphi @ g \cdot p \langle \ell \rangle \vdash \varphi @ g \cdot p \langle \ell \rangle} \quad \text{AX} \frac{}{\psi @ g \cdot p \langle \ell \rangle \vdash \psi @ g \cdot p \langle \ell \rangle} \\
\text{SAYS R} \frac{}{\varphi @ g \cdot p \langle \ell \rangle \vdash p \text{ says}_\ell \varphi @ g} \quad \text{SAYS L} \frac{}{p \text{ says}_\ell \psi @ g \vdash \psi @ g \cdot p \langle \ell \rangle} \\
\text{IMPL}' \frac{}{(p \text{ says}_\ell \varphi) \rightarrow (p \text{ says}_\ell \psi) @ g, \varphi @ g \cdot p \langle \ell \rangle \vdash \psi @ g \cdot p \langle \ell \rangle} \\
\text{IMPR}' \frac{}{(p \text{ says}_\ell \varphi) \rightarrow (p \text{ says}_\ell \psi) @ g \vdash (\varphi \rightarrow \psi) @ g \cdot p \langle \ell \rangle} \\
\text{SAYS R} \frac{}{(p \text{ says}_\ell \varphi) \rightarrow (p \text{ says}_\ell \psi) @ g \vdash p \text{ says}_\ell (\varphi \rightarrow \psi) @ g}
\end{array}$$

Figure 7: Proof that IMPL' and IMPR' allow **says** to undistribute over implication.

To prove Theorem 5, we first normalize each FLAFOL proof and then induct on the formula φ followed by each proof in turn. Both of these inductions are very involved. Appendix C contains more details.

This theorem is one of the key theorems of proof theory [Tak87, GLT89]. Frank Pfenning has called it “[t]he central property of sequent calculi” [Pfe95]. From the propositions-as-types perspective, cut elimination is preservation of types under substitution.

5.5 Implications and Communication

Recall from Section 4 how we interpret implication formulae such as **Alice says** $_\ell (\varphi \rightarrow \psi)$: if φ is true about the system, then Alice believes ψ at label ℓ . We can now see why we use this interpretation of implication. In particular, we consider replacing IMPL and IMPR with the following rules:

$$\text{IMPL}' \frac{\Gamma \vdash \varphi @ g \quad \Gamma, \psi @ g \vdash \chi @ g'}{\Gamma, (\varphi \rightarrow \psi) @ g \vdash \chi @ g'} \quad \text{IMPR}' \frac{\Gamma, \varphi @ g \vdash \psi @ g}{\Gamma \vdash \varphi \rightarrow \psi @ g}$$

Doing so allows us to prove that **says** distributes over implications, as we can see in Figure 6. It also allows us to prove that **says** *un-distributes* over implication, as we see in Figure 7. While IMPL' , IMPR' , and the **says** distribution results may all appear sensible, they actually cause security bugs and make cut elimination impossible.

To see why, imagine that there are three principals of interest: Alice, Bob, and Cathy, and three labels: ℓ_P , ℓ_S , and ℓ_{TS} , representing **Public**, **Secret**, and **TopSecret**, respectively. (We use the shorter names to make our formal proofs easier to read.) Anybody in the system can read public data (i.e., data labeled with ℓ_P). Alice and Cathy believe all three principals of interest can read secret data (i.e., data labeled with ℓ_S), but Bob is unsure of the security clearances and will only send public data to other principals. Alice and Cathy also have top secret clearance, but Bob does not, so he *cannot* read data labeled at ℓ_{TS} . We can formalize these permission policies in the following context:

$$\begin{array}{c}
\text{AX} \frac{}{\Gamma, \varphi @ \text{Alice}\langle l_{\text{TS}} \rangle \vdash \varphi @ \text{Alice}\langle l_{\text{TS}} \rangle} \\
\text{IMPL}' \frac{\Gamma, \varphi @ \text{Alice}\langle l_{\text{TS}} \rangle, \psi @ \text{Alice}\langle l_{\text{TS}} \rangle \vdash \psi @ \text{Alice}\langle l_{\text{TS}} \rangle}{\Gamma, (\varphi \rightarrow \psi) @ \text{Alice}\langle l_{\text{TS}} \rangle, \varphi @ \text{Alice}\langle l_{\text{TS}} \rangle \vdash \psi @ \text{Alice}\langle l_{\text{TS}} \rangle} \text{AX} \\
\text{FWDL}^\dagger \frac{\Gamma, (\varphi \rightarrow \psi) @ \text{Alice}\langle l_{\text{TS}} \rangle, \varphi @ \text{Alice}\langle l_{\text{TS}} \rangle \vdash \psi @ \text{Alice}\langle l_{\text{TS}} \rangle}{\Gamma, (\varphi \rightarrow \psi) @ \text{Alice}\langle l_{\text{TS}} \rangle, \varphi @ \text{Cathy}\langle l_{\text{TS}} \rangle \vdash \psi @ \text{Alice}\langle l_{\text{TS}} \rangle}
\end{array}$$

Figure 8: Alice using Cathy’s φ and a redaction function

$$\begin{array}{c}
\text{SAYSR} \frac{\text{AX} \frac{}{\Gamma, \varphi @ \text{Bob}\langle l_{\text{S}} \rangle \vdash \varphi @ \text{Bob}\langle l_{\text{S}} \rangle}}{\Gamma, \varphi @ \text{Bob}\langle l_{\text{S}} \rangle \vdash \text{Bob says}_{l_{\text{S}}} \varphi @ \langle \rangle} \quad \frac{\text{AX} \frac{}{\Gamma, \psi @ \text{Bob}\langle l_{\text{P}} \rangle \vdash \psi @ \text{Bob}\langle l_{\text{P}} \rangle}}{\Gamma, \text{Bob says}_{l_{\text{P}}} \psi @ \langle \rangle \vdash \psi @ \text{Bob}\langle l_{\text{P}} \rangle} \text{SAYS} \\
\text{IMPL}' \frac{\Gamma, \varphi @ \text{Bob}\langle l_{\text{S}} \rangle \vdash \text{Bob says}_{l_{\text{S}}} \varphi @ \langle \rangle \quad \Gamma, \text{Bob says}_{l_{\text{P}}} \psi @ \langle \rangle \vdash \psi @ \text{Bob}\langle l_{\text{P}} \rangle}{\Gamma, (\text{Bob says}_{l_{\text{S}}} \varphi) \rightarrow (\text{Bob says}_{l_{\text{P}}} \psi) @ \langle \rangle, \varphi @ \text{Bob}\langle l_{\text{S}} \rangle \vdash \psi @ \text{Bob}\langle l_{\text{P}} \rangle} \\
\text{FWD}^\dagger \frac{\Gamma, (\text{Bob says}_{l_{\text{S}}} \varphi) \rightarrow (\text{Bob says}_{l_{\text{P}}} \psi) @ \langle \rangle, \varphi @ \text{Bob}\langle l_{\text{S}} \rangle \vdash \psi @ \text{Bob}\langle l_{\text{P}} \rangle}{\Gamma', \varphi @ \text{Bob}\langle l_{\text{S}} \rangle \vdash \psi @ \text{Alice}\langle l_{\text{P}} \rangle} \\
\text{FWDL}^\dagger \frac{\Gamma', \varphi @ \text{Bob}\langle l_{\text{S}} \rangle \vdash \psi @ \text{Alice}\langle l_{\text{P}} \rangle}{\Gamma', \varphi @ \text{Alice}\langle l_{\text{S}} \rangle \vdash \psi @ \text{Alice}\langle l_{\text{P}} \rangle} \\
\text{VARR}^\dagger \frac{\Gamma', \varphi @ \text{Alice}\langle l_{\text{S}} \rangle \vdash \psi @ \text{Alice}\langle l_{\text{P}} \rangle}{\Gamma', \varphi @ \text{Alice}\langle l_{\text{S}} \rangle \vdash \psi @ \text{Alice}\langle l_{\text{S}} \rangle} \\
\text{IMPR}' \frac{\Gamma', \varphi @ \text{Alice}\langle l_{\text{S}} \rangle \vdash \psi @ \text{Alice}\langle l_{\text{S}} \rangle}{\Gamma' \vdash (\varphi \rightarrow \psi) @ \text{Alice}\langle l_{\text{S}} \rangle} \\
\text{VARR}^\dagger \frac{\Gamma' \vdash (\varphi \rightarrow \psi) @ \text{Alice}\langle l_{\text{S}} \rangle}{\Gamma' \vdash (\varphi \rightarrow \psi) @ \text{Alice}\langle l_{\text{TS}} \rangle}
\end{array}$$

Figure 9: Proof corresponding to Alice sending φ to Bob and receiving a ψ back

$$\begin{aligned}
\Gamma = & \forall p: \text{Principal}. p \text{ says}_{l_{\text{S}}} l_{\text{P}} \sqsubseteq l_{\text{S}} @ \langle \rangle, \\
& \forall p: \text{Principal}. p \text{ says}_{l_{\text{TS}}} l_{\text{S}} \sqsubseteq l_{\text{TS}} @ \langle \rangle, \\
& \text{CanRead}(\text{Bob}, l_{\text{S}}) @ \text{Alice}\langle l_{\text{S}} \rangle, \\
& \text{CanRead}(\text{Alice}, l_{\text{TS}}) @ \text{Cathy}\langle l_{\text{TS}} \rangle, \\
& \forall p, q: \text{Principal}. p \text{ says}_{l_{\text{P}}} \text{CanRead}(q, l_{\text{P}}) @ \langle \rangle, \\
& \forall p, q: \text{Principal}. \forall \ell, \ell': \text{Label}. p \text{ says}_{\ell} \text{CanWrite}(q, \ell') @ \langle \rangle
\end{aligned}$$

Additionally, Bob serves as a redactor: given φ —which represents a document containing secret information—he can produce ψ —which represents a redacted version of the same document—performing a declassification in the process. We represent Bob’s ability by adding one belief:

$$\Gamma' = \Gamma, (\text{Bob says}_{l_{\text{S}}} \varphi) \rightarrow (\text{Bob says}_{l_{\text{P}}} \psi) @ \langle \rangle$$

Imagine further that Alice decides she wants to redact secret information from a **TopSecret** version of φ that she receives from Cathy, but leave it **TopSecret**. If she can figure out how to get an implication representing redaction, she can simply receive φ from Cathy and use the implication. This is the proof in Figure 8. For the sake of brevity and readability, we do not explicitly state side conditions that are proven straightforwardly from Γ . The rules where these side conditions should appear are marked with “ \dagger .”

While she knows how to *use* an implication representing redaction, Alice does not know how to redact φ except by giving it to Bob. Using IMPL' and IMPR' , she is able to package up the process “give Bob a secret version of φ , get back a public version of ψ , and then use variance to get a secret version of ψ ” as a belief $\varphi \rightarrow \psi @ \text{Alice}\langle l_{\text{S}} \rangle$. She can then use variance again to get a belief $\varphi \rightarrow \psi @ \text{Alice}\langle l_{\text{TS}} \rangle$. This is the proof in Figure 9. Again, we elide side conditions that are proven straightforwardly from Γ , and mark the rules where they should appear with “ \dagger .”

Cutting these two proofs together gives Alice what she wants: a **TopSecret** version of ψ . However, this

cut is not possible to eliminate! Examining this through a propositions-as-types lens tells us why: one of Alice or Cathy must send a `TopSecret` version of φ to Bob, which neither is willing to do.

6 Non-Interference

Both authorization logics and information flow systems have important security properties called *non-interference* [Den76, GM82, GP06]. On the face, these two notions of non-interference look very different, but their core intuitions are the same. Both statements aim to prevent one belief or piece of data from interfering with another—even *indirectly*—unless the security policies permit an influence. Authorization logics traditionally define trust relationships between principals and non-interference requires that p 's beliefs affect the provability of q 's beliefs only when q trusts p . Information flow control systems generally specify policies as labels on program data and use the label flows-to relation to constrain how inputs can affect outputs. For non-interference to hold, changing an input with label ℓ_1 can only alter an output with label ℓ_2 if $\ell_1 \sqsubseteq \ell_2$.

FLAFOL views both trust between principals and flows between labels as ways to constrain communication of beliefs. The forward rules model an authorization-logic-style sending of beliefs from one principal to another based on their trust relationships. The label variance rules model a single principal transferring beliefs between labels based on the flow relationship between them. By reasoning about generalized principals, which include both the principal and the label, we are able to capture both at the same time. The result (Theorem 7) mirrors the structure of existing authorization logic non-interference statements [GP06, Aba06]. No similar theorem reasons about information flow or applies to policies combining discoverable trust and logical disjunction. Theorem 7 does both.

6.1 Trust in FLAFOL

Building a notion of trust on generalized principals requires us to consider both the trust of the underlying (regular) principals and label flows. The explicit label flow relation (\sqsubseteq) cleanly captures restrictions on changing labels. Trust between principals requires more care. Alice may trust Bob with public data, but that does not mean she trusts him with secret data. Similarly, Alice may believe that Bob can influence low integrity data without believing Bob is authorized to influence high integrity data. This need to trust principals differently at different labels leads us to define our trust in terms of the two permission relations: $\text{CanRead}(p, \ell)$ and $\text{CanWrite}(p, \ell)$.

We group label flows and principal trust together in a meta-level statement relating generalized principals. As this relation is the fundamental notion of trust in FLAFOL, we follow existing authorization logic literature and call it *speaks for*.

The speaks-for relation captures any way that one generalized principal's beliefs can be safely transferred to another. This can happen through flow relationships ($g \cdot p\langle\ell\rangle$ speaks for $g \cdot p\langle\ell'\rangle$ if $\ell \sqsubseteq \ell'$), forwarding ($g \cdot p\langle\ell\rangle$ speaks for $g \cdot q\langle\ell\rangle$ if p can forward beliefs at ℓ to q), and introspection ($g \cdot p\langle\ell\rangle$ speaks for $g \cdot p\langle\ell\rangle \cdot p\langle\ell\rangle$ and vice versa). We formalize speaks-for with the rules in Figure 10.

To validate this notion of trust, we note that existing authorization logics often define speaks-for as an atomic relation and create trust by requiring that, if p speaks for q , then p 's beliefs can be transferred to q . As our speaks-for relation exactly mirrors FLAFOL's rules for communication, it enjoys this same property.

Theorem 6 (Speaks-For Elimination). *The following rule is admissible in FLAFOL:*

$$\text{ELIMSF} \frac{\Gamma \vdash \varphi @ g_1 \quad \Gamma \vdash g_1 \text{ SF } g_2}{\Gamma \vdash \varphi @ g_2}$$

$$\begin{array}{c}
\text{REFLSF} \frac{}{\Gamma \vdash g \text{ SF } g} \qquad \text{EXTSF} \frac{\Gamma \vdash g_1 \text{ SF } g_2}{\Gamma \vdash g_1 \cdot p\langle \ell \rangle \text{ SF } g_2 \cdot p\langle \ell \rangle} \\
\text{SELFLSF} \frac{}{\Gamma \vdash g \cdot p\langle \ell \rangle \text{ SF } g \cdot p\langle \ell \rangle \cdot p\langle \ell \rangle} \qquad \text{SELFRSF} \frac{}{\Gamma \vdash g \cdot p\langle \ell \rangle \cdot p\langle \ell \rangle \text{ SF } g \cdot p\langle \ell \rangle} \\
\text{VARSF} \frac{\Gamma \vdash \ell \sqsubseteq \ell' @ g \cdot p\langle \ell' \rangle}{\Gamma \vdash g \cdot p\langle \ell \rangle \text{ SF } g \cdot p\langle \ell' \rangle} \qquad \text{FWDSF} \frac{\Gamma \vdash \text{CanRead}(q, \ell) @ g \cdot p\langle \ell \rangle \quad \Gamma \vdash \text{CanWrite}(p, \ell) @ g \cdot q\langle \ell \rangle}{\Gamma \vdash g \cdot p\langle \ell \rangle \text{ SF } g \cdot q\langle \ell \rangle} \\
\text{TRANSSF} \frac{\Gamma \vdash g_1 \text{ SF } g_2 \quad \Gamma \vdash g_2 \text{ SF } g_3}{\Gamma \vdash g_1 \text{ SF } g_3}
\end{array}$$

Figure 10: The rules defining `says` for.

With this notion of trust we can begin structuring a non-interference statement. We might like to say that beliefs of g_1 can only influence beliefs of g_2 if $\Gamma \vdash g_1 \text{ SF } g_2$, or formally: if $\Gamma, (\varphi @ g_1) \vdash \psi @ g_2$ is provable, then either $\Gamma \vdash \psi @ g_2$ is provable or $\Gamma \vdash g_1 \text{ SF } g_2$. Unfortunately, this statement is false for three critical reasons: `says` statements, implication, and the combination of discoverable trust and disjunctions.

6.2 Says Statements and Non-Interference

The first way to break the proposed non-interference statement above is simply by moving affirmations of a statement between the formula—using `says`—and the generalized principal who believes it. For example, we can trivially prove $p \text{ says}_\ell \varphi @ \langle \rangle \vdash \varphi @ \langle \rangle \cdot p\langle \ell \rangle$, yet we cannot prove $\langle \rangle \text{ SF } \langle \rangle \cdot p\langle \ell \rangle$.

To address this case, we can view $p \text{ says}_\ell \varphi @ \langle \rangle$ as a statement that $\langle \rangle \cdot p\langle \ell \rangle$ believes φ . This insight suggests generally pushing all `says` modalities into the generalized principal. We can do this for simple formulae, but the process breaks down with conjunction and disjunction. In those cases, the different sides may have different `says` modalities, and either side could influence a belief through the different resulting generalized principals. We alleviate this concern by considering a *set* of generalized principals referenced in a given belief. We build this set using an operator \mathcal{G} :

$$\mathcal{G}(\chi @ g) \triangleq \begin{cases} \mathcal{G}(\varphi @ g \cdot p\langle \ell \rangle) & \chi = p \text{ says}_\ell \varphi \\ \mathcal{G}(\varphi @ g) \cup \mathcal{G}(\psi @ g) & \chi = \varphi \wedge \psi \text{ or } \varphi \vee \psi \\ \mathcal{G}(\psi @ g) & \chi = \varphi \rightarrow \psi \\ \bigcup_{t:\sigma} \mathcal{G}(\varphi[x \mapsto t] @ g) & \chi = \forall x:\sigma. \varphi \text{ or } \exists x:\sigma. \varphi \\ \{g\} & \text{otherwise} \end{cases}$$

For implications, \mathcal{G} only considers the consequent, since only its consequent can affect the provability of a belief. For quantified formulae, a proof may substitute any term of the correct sort for the bound variable, so we must as well.

Using this new operator, we can patch the hole `says` statements created in our previous non-interference statement, producing the following: If $\Gamma, (\varphi @ g_1) \vdash \psi @ g_2$, then either $\Gamma \vdash \psi @ g_2$, or there is some $g'_1 \in \mathcal{G}(\varphi @ g_1)$, $g'_2 \in \mathcal{G}(\psi @ g_2)$, and some g''_1 such that $\Gamma \vdash g'_1 \cdot g''_1 \text{ SF } g'_2$.

Here g''_1 represents the ability of a generalized principal to ship entire simulations to other generalized principals. In particular, the forward and variance rules operate on an “active” prefix of the current generalized principal; g'_1 represents the “inactive” suffix.

The \mathcal{G} operator converts reasoning about beliefs from the object level (FLAFOL formulae) to the meta level (generalized principals). FLAFOL’s ability to freely move between the two forces us to push all such

$$\begin{array}{c}
\text{SF-CI} \frac{\Gamma \vdash g_1 \text{ SF } g_2}{\Gamma \vdash g_1 \text{ CanInfl } g_2} \qquad \text{EXTCI} \frac{\Gamma \vdash g_1 \text{ CanInfl } g_2}{\Gamma \vdash g_1 \cdot g' \text{ CanInfl } g_2 \cdot g'} \\
\text{TRANSCI} \frac{\Gamma \vdash g_1 \text{ CanInfl } g_2 \quad \Gamma \vdash g_2 \text{ CanInfl } g_3}{\Gamma \vdash g_1 \text{ CanInfl } g_3} \\
\text{IMPCI} \frac{\varphi \rightarrow \psi @ g \in \Gamma \quad g_1 \in \mathcal{G}(\varphi @ \langle \rangle) \quad g_2 \in \mathcal{G}(\psi @ g)}{\Gamma \vdash g_1 \text{ CanInfl } g_2}
\end{array}$$

Figure 11: The rules defining the *can influence* relation.

reasoning in the same direction to effectively compare the reasoner in two different beliefs. Prior authorization logics do not contain a meta-level version of **says**, meaning similar conversions do not even make sense.

6.3 Implications

While use of the \mathcal{G} function solves part of the problem with our original non-interference proposal, it does not address all of the problems. Implications can implicitly create new trust relationships, allowing beliefs of one generalized principal to affect beliefs of another, even when no speaks-for relationship exists. To understand how this can occur, we revisit our example of preventing SQL injection attacks from from Section 2.2.

Recall from Section 2.2 that a web server might treat sanitized versions of low-integrity input as high integrity. Further recall, it might represent this willingness with the following implication.

$$\text{System says}_{\text{LInt}} \text{DBInput}(x) \rightarrow \text{System says}_{\text{HInt}} \text{DBInput}(\text{San}(x))$$

In an intuitively-sensible context where **System** believes $\text{HInt} \sqsubseteq \text{LInt}$ —high integrity flows to low integrity—but not vice versa, there is no way to prove $\text{System}\langle \text{LInt} \rangle \text{ SF } \text{System}\langle \text{HInt} \rangle$. The presence of this implication, however, allows some beliefs at $\text{System}\langle \text{LInt} \rangle$ to influence beliefs at $\text{System}\langle \text{HInt} \rangle$. This influence is actually an endorsement from LInt to HInt , and our speaks-for relation explicitly does not capture such effects.

Prior work manages this trust-creating effect of implications either by claiming security only when all implications are provable [Aba06] or by explicitly using assumed implications to represent trust [GP06]. We hew closer to the latter model and make the implicit trust of implications explicit in our statement of non-interference. We therefore cannot use the speaks-for relation, so we construct a new relation between generalized principals we call *can influence*.

Intuitively, g_1 can influence g_2 —which we denote $\Gamma \vdash g_1 \text{ CanInfl } g_2$ —if either g_1 speaks for g_2 or there is an implication in Γ that allows a belief of g_1 to affect the provability of a belief of g_2 . This relation, formally defined in Figure 11, uses the \mathcal{G} operator discussed above to capture the generalized principals actually discussed by each subformula of the implication. Because FLAFOL interprets the premise of an implication as a condition whose modality is independent of the entire belief, so too does the can-influence relation. The relation is also transitive, allowing it to capture the fact that a proof may require many steps to go from a belief at g_1 to a belief at g_2 .

Simply taking our attempted non-interference statement from above and replacing speaks-for with can-influence allows us to straightforwardly capture the effect of implications on trust within the system.

While this change may appear small, it results in a highly conservative estimate of possible influence. Implications are precise statements that can allow usually-disallowed information flows under very particular circumstances. Unfortunately, because our non-interference statement only considers the generalized

principals involved, not the entire beliefs, it cannot represent the same level of precision. A single precise implication added to a context can therefore relate whole classes of previously-unrelated generalized principals, eliminating the ability for non-interference to say anything about their relative security. A similar lack of precision in information flow non-interference statements has resulted in long lines of research on how to precisely model or safely restrict declassification and endorsement [ZM01, SM04, MS04, LZ05, SS05, MSZ06, CM08, AM11, WBK⁺15, CMA17].

6.4 Discovering Trust with Disjunctions

The \mathcal{G} operator and can-influence relation address difficulties from both **says** formulae and implications, but our statement of non-interference still does not account for the combination of disjunctions and the ability to discover trust relationships. To understand the effect of these two features in combination, recall the reinsurance example from Section 2.3. Bob can derive $\text{CanWrite}(I_1, \ell_H)$ if he already believes both $\text{CanWrite}(I_1, \ell_H) \vee \text{CanWrite}(I_2, \ell_H)$ and $I_2 \text{ says}_{\ell_H} \text{CanWrite}(I_1, \ell_H)$.

We clearly cannot remove either of Bob's beliefs and still prove the result. Our desired theorem statement would thus require that $\text{Bob}\langle \ell_H \rangle \cdot I_2\langle \ell_H \rangle$ can influence $\text{Bob}\langle \ell_H \rangle$, which there is no way to prove. The reason the sequent is still provable, as we noted in Section 2.3, is that Bob can *discover* trust in I_2 when he branches on an Or statement, which then allows I_2 to influence Bob. In this branch, we can prove $\text{Bob}\langle \ell_H \rangle \cdot I_2\langle \ell_H \rangle \text{ SF } \text{Bob}\langle \ell_H \rangle \cdot \text{Bob}\langle \ell_H \rangle$, which then speaks for $\text{Bob}\langle \ell_H \rangle$.

To handle such assumptions, we cannot simply consider the context in which we are proving a sequent; we must consider any context that can appear in the proof of that sequent. We developed the notion of compatible supercontexts in Section 5.3 for exactly this purpose. Indeed, if we replace Γ with an appropriate CSC when checking the potential influence of generalized principals, we remove the last barrier to a true non-interference theorem.

6.5 Formal Non-Interference

The techniques above allow us to modify our attempted non-interference statement into a theorem that holds.

Theorem 7 (Non-Interference). *For all contexts Γ and beliefs $\varphi @ g_1$ and $\psi @ g_2$, if*

$$\Gamma, \varphi @ g_1 \vdash \psi @ g_2,$$

then either (1) $\Gamma \vdash \psi @ g_2$, or (2) there is some $\Delta \ll \Gamma, \varphi @ g_1 \vdash \psi @ g_2, g'_1 \in \mathcal{G}(\varphi @ g_1), g'_2 \in \mathcal{G}(\psi @ g_2)$, and g''_1 such that $\Delta \vdash g'_1 \cdot g''_1 \text{ CanInfl } g'_2$.

The proof of this theorem follows by induction on the proof of $\Gamma, \varphi @ g_1 \vdash \psi @ g_2$. For each proof rule, we argue that either $\varphi @ g_1$ is unnecessary for all premises or we can extend an influence from one or more subproofs to an influence from $\varphi @ g_1$ to $\psi @ g_2$.

This theorem limits when a belief $\varphi @ g_1$ can be necessary to prove $\psi @ g_2$ in context Γ , much like other authorization logic non-interference statements [GP06, Aba06]. As we mentioned above, however, it is the first such non-interference statement for any authorization logic supporting all first-order connectives and discoverable trust. Moreover, it describes how FLAFOL mitigates both:

- communication between principals, through **CanRead** and **CanWrite** statements, and
- movement of information between security levels represented by information flow labels, via flows-to statements.

The **CanInfl** relation seems to make our non-interference statement much less precise than we would like. After all, implications precisely specify what beliefs can be declassified or endorsed, whereas **CanInfl**

conservatively assumes any beliefs can move between the relevant generalized principals. This lack of precision serves a purpose. It allows us to reason about any implications, including those that arbitrarily change principals and labels, something which other no authorization logics have done before. It is therefore worth noting that, when all of the implications in the context are provable, the theorem holds *even if you replace CanInfl with SF everywhere*. The same proof works, with some simple repair in the IMPL case.

Another complaint of imprecision applies to compatible supercontexts. Specifically, if any principal assumes $\varphi \vee \neg\varphi$ for any formula φ , then there is a CSC in which that principal has assumed both, even though these are arrived at through mutually-exclusive choices. Since CSCs have been added in order to allow disjunctions and discoverable trust to co-exist, it is good to know that if we disallow either, CSCs are not required for non-interference. That is, if there are no disjunctions in the context, then we can always instantiate the Δ in Theorem 7 with $\Gamma, \varphi @ g_1$. Similarly, if every permission that is provable in any CSC of $\Gamma, \varphi @ g_1 \vdash \psi @ g_2$ is provable under $\Gamma, \varphi @ g_1$, then we can again always instantiate Δ with $\Gamma, \varphi @ g_1$.

Together, these points demonstrate that there are only two types of poorly-behaved formulae that force the imprecision in Theorem 7. This further shows that our non-interference result is no less precise than those of other authorization logics in the absence of such formulae. We add imprecision only when needed to allow our statement to apply to more proofs.

To see how Theorem 7 corresponds to traditional non-interference results for information flow, consider a setting where every principal agrees on the same label ordering, and where there are no implications corresponding to declassifications or endorsements. Then any two contexts Γ and Γ' which disagree only on beliefs labeled above some ℓ can prove exactly the same things at label ℓ — $\Gamma \vdash \varphi @ g \cdot p(\ell)$ if and only if $\Gamma' \vdash \varphi @ g \cdot p(\ell)$ —since Theorem 7 allows us to delete all of the beliefs on which they disagree. If we view contexts as inputs, as in a propositions-as-types interpretation, then this says that changing high inputs cannot change low results.

7 Future Work

FLAFOL is already very powerful, but it suggests numerous avenues for future work.

First, FLAFOL only disallows *direct* flows of information in proofs, but checking proofs can cause communication and potentially leak information. Importantly, eliminating cuts in proofs can *increase* the information leaked during proof-checking because eliminating cuts can reduce the uncertainty about which discoveries can be made during a proof. This is disturbing, since we would like to be able to perform sound security analyses on proofs with cut; system designers should not need to understand the very complicated cut-elimination proof. The *program counter* mechanism used by information flow control systems like Fabric [LAGM17] and FLAM [ALM15] seems to prevent similar leaks. Incorporating program counter labels to limit communication in FLAFOL proofs could eliminate these leaks in FLAFOL as well.

This improvement also widens the range of programs that can safely use FLAFOL. Justifications for authorization need to be found as well as checked. From the point of view of an authorization logic, this corresponds to proof search. Searching for an authorization proof in a distributed system, however, may require communication between principals, potentially leaking why they are searching for this proof in the first place. One avenue forward embeds FLAFOL in a language with information-flow types, and runs proof search in that language. This would guarantee that the proof search does not leak data assuming FLAFOL proofs do not leak data when checked.

We have developed new techniques to reason about authorization-logic proofs in order to prove non-interference for FLAFOL. These reasoning principles could be expanded and used in other logics. For instance, using the tools developed in Section 6, we should be able to give non-interference proofs for logics like NAL [SWS11] and FOCAL [HC13] which reason about implication and disjunction. We should also be able to add disjunction and implication to logics like DCC [Aba06, ABHR99] while still providing a

non-interference theorem.

Another avenue of further work would understand better how **says** statements can interact with other logical connectives. For instance, one might want to model a principal who cannot observe whether they are holding evidence of φ or of ψ . For instance, we might want to model a principal p who receives an encrypted message containing a bit b . Then p knows that either $b = 0$ or $b = 1$, but p has no way to determine which. Thus, while p **says** _{ℓ} $(b = 0 \vee b = 1)$, we should not be able to show that $(p$ **says** _{ℓ} $b = 0) \vee (p$ **says** _{ℓ} $b = 1)$. A NuPRL-like “squash” operator, which prevents evidence from being used [Cal98], could model this, but further research is needed for FLAFOL to reason about the security of such protocols.

A similar avenue for future work involves exploring ways to allow **says** to distribute over implications while remaining coherent. One potential approach would be to confine most reasoning to a single generalized principal, but this would restrict implications so that the principal who believes them cannot communicate in their proof. The consequences of such a restriction on modeling real-world systems are unclear.

Finally, it would be nice to reason about the *temporal* components of authorization; this is one place where work on information flow far outstrips that on authorization logic [ZM01, SM04, MS04, LZ05, SS05, MSZ06, CM08, AM11, WBK⁺15, CMA17]. Trust relationships may change over time, allowing or disallowing communication pathways. Understanding how this changes which authorizations should be provable, and how this affects information-flow policies, is a rich area for exploration.

8 Related Work

Prior work in information flow and authorization logics has explored the connection between the two. The Decentralized Label Model [ML98, ML00] includes a notion of ownership in information flow policies specifying who may authorize exceptions to the policy. The Flow-Limited Authorization Model (FLAM) [ALM15] was the first logic to directly consider the effects of data confidentiality and integrity on trust relationships between principals. Prior work on Rx [SHTZ06] and RTI [BWW08] enforced language-based information flow policies via *roles* whose membership were protected with confidentiality and integrity labels. By contrast, FLAFOL is a formal authorization logic containing every first-order connective.

Decentralized Label Model. The Decentralized Label Model (DLM) [ML98, ML00] is a model for expressing information flow labels in a decentralized system. Its labels contain two components, confidentiality and integrity, which are each specified as a set of principals who may read or write the data, respectively. FLAFOL separates principals and labels by making them independent sorts that are related by the **CanRead** and **CanWrite** relations. This allows system designers much more freedom in determining the semantics of principals and labels. For instance, DLM labels cannot represent availability.

DLM labels also include a notion of ownership, but it only specifies who may authorize exceptions to the policy. FLAFOL has no built-in ownership notion, nor does it allow specific exceptions to policies.

Moreover, DLM assumes that labels form a global static lattice. As we have discussed in detail, FLAFOL does not make this assumption. In particular, FLAFOL labels need not be static, since they can be the result of functions. Second, they need not be a lattice, but merely a partial order. Finally, the order of FLAFOL labels need not be global, since different (generalized) principals may have very different ideas of the order. This generality allows FLAFOL to be extremely expressive, as we saw in Section 3.

Flow-Limited Authorization Model. The Flow-Limited Authorization Model (FLAM) [ALM15] was the first information-flow label model to directly consider the interaction between information flow and authorization. FLAM does not, however, provide a full authorization logic. It lays out important rules for reasoning about communication in systems with discoverable trust relationships where principals may disagree on those relationships. It also restricts participation in a proof using a program counter label to help full systems remain secure in contexts where merely checking a proof may leak data. FLAM, however,

provides no means to directly express authorization policies other than one principal trusting another. It has no first-order connectives or quantifiers and no way for one principal to reason about another’s beliefs.

FLAM also takes the principal-label connection a step beyond the DLM and represents principals directly as a combination of confidentiality and integrity labels. This view restricts FLAM from reasoning about labels with policies other than confidentiality and integrity, since they might necessitate subtle changes to FLAM’s reasoning rules. FLAFOL’s `CanRead` and `CanWrite` relations abstract out how different label components may interact, allowing each system to specify appropriate restrictions given the meaning of its labels.

Unifying principals and labels also undermines FLAM’s effectiveness as an authorization logic. It is often convenient to construct complex policies from simpler ones, such as a policy protecting Alice’s confidentiality and Bob’s integrity. FLAM regards such a compound policy as a principal, breaking the connection between formal principals and system entities. While FLAFOL can certainly represent these policies, doing so does not force a reasoner to break this connection.

FLAM additionally does not provide a non-interference guarantee, instead offering a guarantee called *robust authorization*. In FLAM, each fact has a label representing its confidentiality and integrity and is stored on a node, which is itself represented by a label. If a node c believes a derived fact at label ℓ , robust authorization says:

- The label of every fact used in the derivation flows to ℓ ,
- Every node in the derivation may control whether the derivation took place,
- c is allowed to learn every fact used in the derivation, and
- For each node n involved in the derivation, c will listen to n at ℓ and n will talk to c at ℓ .

FLAFOL’s non-interference theorem gives similar guarantees. In particular, our non-interference theorem shows that the label of every belief used in a derivation (without implications) flows to the label of the derived belief. Moreover, for each belief $\varphi @ g$ used in a derivation without implications, the generalized principal who believes the conclusion must (transitively) trust g .

However, FLAFOL does not have any notion of who may control whether a derivation takes place. We are able to achieve FLAFOL’s security guarantee without the restrictions imposed by FLAM’s program counter label.

DCC and FLAC. The Dependency Core Calculus (DCC) [ABHR99] is a small functional core calculus designed to capture dependencies within programs, including information flows. It uses a monadic structure to represent labels at the level of types and enforce standard information flow typing constraints. Abadi also reinterpreted DCC’s type system as an authorization logic [Aba06], but used the modalities created by the monadic structure to represent principals’ beliefs. This technique allows DCC to reason about either information flow or authorization, but not both at the same time. DCC does provide a non-interference property, but it employs a static external lattice to express trust.

The Flow-Limited Authorization Calculus (FLAC) [AM16] builds a computational model for FLAM by extending Polymorphic DCC [Aba06] with discoverable trust relationships. It uses DCC’s information-flow interpretation and FLAM’s discoverable trust rules to bound information flows and how they can affect trust assumptions.

Because FLAC incorporates DCC’s computational model, we can view its type system as a propositional logic that reasons about discoverable trust. Since the logic is based on System F, it contains some elements of second-order logic by supporting universal quantification over types, but lacks any existential quantification. Critically, FLAC programs execute only on a single machine with no notion of communication. This means that, unlike both FLAM and FLAFOL, it does not allow reasoning about the interaction between different system components with different trust assumptions, and thus does not form a full authorization logic. It can only reason about how data may influence trust assumptions and resulting decisions within a single

component. DFLATE [GCA19] extends FLAC with channels that support a limited form of communication.

FLAC provides strong information security guarantees for computations defined in the language. The local nature of these computations, however, means these assurances apply only to local reasoning on a single host. FLAFOL, by contrast, provides strong security guarantees in a fully distributed context when reasoning about differing beliefs within a system. It does not yet have an associated programming model, but developing one would be interesting future work.

Other Authorization Logics. Becker [Bec12] explores preventing probing attacks, authorization queries which leak secret information, in Datalog-based authorization logics like DKAL [GN08] and SecPAL [BFG10]. In SecPAL⁺ [Bec10], Becker proposes a new *can listen to* operator, similar to FLAFOL’s *CanRead* permission, that expresses who is permitted to learn specific statements. However, *can listen to* expresses permissions on specific statements, not labels as *CanRead* does. Moreover, FLAFOL tracks dependencies between statements using these labels, so the security consequences of adding a new permission are more explicit.

Garg and Pfenning [GP06] present an authorization logic and a non-interference result that ensures untrusted principals cannot influence the truth of statements made by other principals. Garg and Pfenning, however, support a more limited set of logical connectives than FLAFOL, use only implications to encode trust, and do not reason directly about information flow.

Finally, AURA [JVM⁺08, JZ09] embeds DCC into a language with dependent types to explore how authorization logic interacts with programs. They inherit their non-interference result directly from DCC, but they express first-order properties by combining other programming language constructs with DCC. This makes it unclear what guarantees the theorem provides. Jia and Zdancewic encode information-flow labels into AURA as principals and develop a non-interference theorem in the style of information-flow systems [JZ09]. This setup unfortunately makes it impossible for principals to disagree about the meaning of labels, since the labels themselves define their properties.

9 Conclusion

We have introduced FLAFOL, a first-order logic which combines notions of trust from both authorization and information flow. It provides a concrete model of communication that respects this combination and gives principals the ability to reason about each other’s differing opinions, including differing opinions about trust. FLAFOL has a powerful non-interference theorem that navigates this complexity, a top-tier result for authorization logics. It is, moreover, the most complete first-order logic with such a guarantee.

Acknowledgments

We would first like to thank our anonymous reviewers for their insightful comments and helpful suggestions. Andrew Myers and Jed Liu provided early feedback on the design of FLAFOL. Discussion with Deepak Garg gave us insight into how to introduce FLAFOL, while Phokion G. Kolatis pointed us to some related work. We would finally like to thank Coşku Acay, Arthur Azevedo de Amorim, Eric Campbell, Dietrich Geisler, Elisavet Kozyri, Tom Magrino, Matthew Milano, Andrew Morgan, and Drew Zagieboylo for their valuable help editing this paper.

Funding for this work was provided by NSF grant #1704788, NSF CAREER grant #1750060, and a National Defense Science and Engineering Graduate Fellowship. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and may not reflect those of these sponsors.

References

- [Aba06] Martín Abadi. Access control in a core calculus of dependency. In *11th ACM SIGPLAN Int'l Conf. on Functional Programming*, pages 263–273, New York, NY, USA, 2006. ACM.
- [ABHR99] Martín Abadi, Anindya Banerjee, Nevin Heintze, and Jon Riecke. A core calculus of dependency. In *26th ACM Symp. on Principles of Programming Languages (POPL)*, pages 147–160, January 1999.
- [Alg18] Maximilian Alghed. Short paper: A perspective on the dependency core calculus. *PLAS*, October 2018.
- [ALM15] Owen Arden, Jed Liu, and Andrew C. Myers. Flow-limited authorization. In *28th IEEE Symp. on Computer Security Foundations (CSF)*, pages 569–583, July 2015.
- [AM11] Aslan Askarov and Andrew C. Myers. Attacker control and impact for confidentiality and integrity. *Logical Methods in Computer Science*, 7(3), September 2011.
- [AM16] Owen Arden and Andrew C. Myers. A calculus for flow-limited authorization. In *29th IEEE Symp. on Computer Security Foundations (CSF)*, pages 135–147, June 2016.
- [Bec10] Moritz Y. Becker. Information flow in credential systems. In *23rd IEEE Symp. on Computer Security Foundations (CSF)*, pages 171–185. IEEE, 2010.
- [Bec12] Moritz Y Becker. Information flow in trust management systems. *Journal of Computer Security*, 20(6):677–708, 2012.
- [BFG10] Moritz Y Becker, Cédric Fournet, and Andrew D Gordon. SecPAL: Design and semantics of a decentralized authorization language. *Journal of Computer Security*, 18(4):619–665, 2010.
- [BWW08] Sruthi Bandhakavi, William Winsborough, and Marianne Winslett. A trust management approach for flexible policy management in security-typed languages. In *Computer Security Foundations Symposium, 2008*, pages 33–47, 2008.
- [Cal98] James L. Caldwell. Classical propositional decidability via nuprl proof extraction. *TPHOLs*, 1998.
- [Cha12] Arthur Charguéraud. The locally nameless representation. *Journal of Automated Reasoning*, 49(3):363–408, October 2012.
- [CM08] Stephen Chong and Andrew C. Myers. End-to-end enforcement of erasure and declassification. In *IEEE Symp. on Computer Security Foundations (CSF)*, pages 98–111, June 2008.
- [CMA17] Ethan Cecchetti, Andrew C. Myers, and Owen Arden. Nonmalleable information flow control. In *24th ACM Conf. on Computer and Communications Security (CCS)*, pages 1875–1891, October 2017.
- [Coq04] Coq development team. *The Coq proof assistant reference manual*. LogiCal Project, 2004. Version 8.0.
- [Den76] Dorothy E. Denning. A lattice model of secure information flow. *Comm. of the ACM*, 19(5):236–243, 1976.

- [GCA19] Anitha Gollamudi, Stephen Chong, and Owen Arden. Information flow control for distributed trusted execution environments. In *32nd IEEE Symp. on Computer Security Foundations (CSF)*, 2019.
- [GLT89] Jean-Yves Girard, Yves Lafont, and Paul Taylor. *Proofs and Types*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1989.
- [GM82] Joseph A. Goguen and Jose Meseguer. Security policies and security models. In *IEEE Symp. on Security and Privacy*, pages 11–20, April 1982.
- [GN08] Yuri Gurevich and Itay Neeman. DKAL: Distributed-knowledge authorization language. In *IEEE Symp. on Computer Security Foundations (CSF)*, pages 149–162. IEEE, 2008.
- [GP06] Deepak Garg and Frank Pfenning. Non-interference in constructive authorization logic. In *19th IEEE Computer Security Foundations Workshop (CSFW)*, New Jersey, USA, 2006. IEEE.
- [HC13] Andrew K. Hirsch and Michael R. Clarkson. Belief semantics of authorization logic. *CCS*, pages 561–572, November 2013.
- [HK00] Jon Howell and David Kotz. A formal semantics for SPKI. In *ESORICS 2000*, volume 1895 of *Lecture Notes in Computer Science*, pages 140–158. Springer Berlin Heidelberg, 2000.
- [JVM⁺08] Limin Jia, Jeffrey A. Vaughan, Karl Mazurak, Jianzhou Zhao, Luke Zarko, Joseph Schorr, and Steve Zdancewic. AURA: A programming language for authorization and audit. In *13th ACM SIGPLAN Int’l Conf. on Functional Programming*, September 2008.
- [JZ09] Limin Jia and Steve Zdancewic. Encoding information flow in AURA. *PLAS*, pages 17–29, June 2009.
- [LABW91] Butler Lampson, Martín Abadi, Michael Burrows, and Edward Wobber. Authentication in distributed systems: Theory and practice. In *13th ACM Symp. on Operating System Principles (SOSP)*, pages 165–182, October 1991.
- [LAGM17] Jed Liu, Owen Arden, Michael D. George, and Andrew C. Myers. Fabric: Building open distributed systems securely by construction. *J. Computer Security*, 25(4–5):319–321, May 2017.
- [LZ05] Peng Li and Steve Zdancewic. Downgrading policies and relaxed noninterference. In *32nd ACM Symp. on Principles of Programming Languages (POPL)*, Long Beach, CA, January 2005.
- [ML98] Andrew C. Myers and Barbara Liskov. Complete, safe information flow with decentralized labels. In *IEEE Symp. on Security and Privacy*, pages 186–197, May 1998.
- [ML00] Andrew C. Myers and Barbara Liskov. Protecting privacy using the decentralized label model. *ACM Transactions on Software Engineering and Methodology*, 9(4):410–442, October 2000.
- [MM18] Matthew P. Milano and Andrew C. Myers. MixT: A language for mixing consistency in geodistributed transactions. In *39th ACM SIGPLAN Conf. on Programming Language Design and Implementation (PLDI)*, June 2018.
- [MS04] Heiko Mantel and David Sands. Controlled Declassification based on Intransitive Noninterference. In *2nd ASIAN Symposium on Programming Languages and Systems, APLAS 2004*, LNCS 3303, pages 129–145, Taipei, Taiwan, November 2004. Springer-Verlag.

- [MSZ06] Andrew C. Myers, Andrei Sabelfeld, and Steve Zdancewic. Enforcing robust declassification and qualified robustness. *Journal of Computer Security*, 14(2):157–196, 2006.
- [Pfe95] Frank Pfenning. Structural cut elimination. *LICS*, pages 156–166, June 1995.
- [RGR16] Vincent Rajani, Deepak Garg, and Tamara Rezk. On access control, capabilities, their equivalence, and confused deputy attacks. *CSF*, June 2016.
- [SBR⁺11] Emin Gün Sirer, Willem De Bruijin, Patrick Reynolds, Alan Shieh, Kevin Walsh, Dan Williams, and Fred B. Schneider. Logical attestation: An authorization architecture for trustworthy computing. In *11th ACM Symp. on Operating System Principles (SOSP)*, 2011.
- [SHTZ06] Nikhil Swamy, Michael Hicks, Stephen Tse, and Steve Zdancewic. Managing policy updates in security-typed languages. In *19th IEEE Computer Security Foundations Workshop (CSFW)*, pages 202–216, July 2006.
- [SM04] Andrei Sabelfeld and Andrew C. Myers. A model for delimited release. In *2003 International Symposium on Software Security*, number 3233 in Lecture Notes in Computer Science, pages 174–191. Springer-Verlag, 2004.
- [SS05] Andrei Sabelfeld and David Sands. Dimensions and principles of declassification. In *18th IEEE Computer Security Foundations Workshop (CSFW)*, pages 255–269, June 2005.
- [SWS11] Fred B. Schneider, Kevin Walsh, and Emin Gün Sirer. Nexus Authorization Logic (NAL): Design rationale and applications. *ACM Trans. Inf. Syst. Secur.*, 14(1):8:1–8:28, June 2011.
- [Tak87] Gaisi Takeuti. *Proof Theory*. Dover Books on Mathematics. Dover Books, 1987. Second Edition, republished by Dover Books in 2013. Originally published by North-Holland, Amsterdam.
- [VSI96] Dennis Volpano, Geoffrey Smith, and Cynthia Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3):167–187, 1996.
- [WBK⁺15] Lucas Wayne, Pablo Buiras, Dan King, Stephen Chong, and Alejandro Russo. It’s my privilege: Controlling downgrading in DC-labels. In *Proceedings of the 11th International Workshop on Security and Trust Management*, September 2015.
- [Yan12] Edward Z. Yang. Logitext, 2012. Accessed February 19, 2019.
- [ZM01] Steve Zdancewic and Andrew C. Myers. Robust declassification. In *14th IEEE Computer Security Foundations Workshop (CSFW)*, pages 15–23, June 2001.
- [ZM05] Lantian Zheng and Andrew C. Myers. End-to-end availability policies and noninterference. In *18th IEEE Computer Security Foundations Workshop (CSFW)*, pages 272–286, June 2005.

A Formalizing the Examples

In Section 2, we discussed several examples of authorization policies that interact with information-flow policies in non-trivial ways. While we used these to introduce FLAFOL’s features and syntax, in Sections 2.1 and 2.3 we elided some technical details in order to simplify the presentation. In this section, we formalize these examples, making it clear how FLAFOL can represent each of the policies described in Section 2.

A.1 Viewing Pictures on Social Media

Recall the example in Section 2.1: Bob has uploaded a picture to a social-media account along with a policy that only those on a friend list that he maintains on that account may view the photo. Moreover, he has a policy that only his friends may know who is on his friend list.

We mentioned that we can represent Bob’s friend list as a collection of beliefs of the form $\text{Bob says}_{\text{Friends}} \text{IsFriend}(p)$. We can now discuss these beliefs in more detail. Bob’s friend list contains a finite number of principals; let L be the set of principals on the list. We can then represent Bob’s friend list with the following FLAFOL assumptions:

$$\Gamma_L = \{ \text{Bob says}_{\text{Friends}} \text{IsFriend}(p) \mid p \in L \} \cup \left\{ \begin{array}{l} \forall q: \text{Principal}. \\ \text{Bob says}_{\text{Friends}} \text{IsFriend}(q) \\ \rightarrow \left(\bigvee_{p \in L} \text{Bob says}_{\text{Friends}} q = p \right) \end{array} \right\}$$

(For brevity, we omit “@ $\langle \rangle$ ” on all beliefs at $\langle \rangle$.)

This set of beliefs, along with some beliefs about equality (such as reflexivity, decidability of equality on principals, and Bob’s belief that if two principals are equal, then one is a friend if the other is) are enough to determine Bob’s friend list. In particular, one can show that

$$\Gamma_L, \Gamma \vdash \forall p: \text{Principal}. \left(\text{Bob says}_{\text{Friends}} \text{IsFriend}(p) \vee \text{Bob says}_{\text{Friends}} \neg \text{IsFriend}(p) \right).$$

This proof is conceptually simple, but quite tedious, so we elide it here. This is a decision procedure since any FLAFOL proof of $\Gamma \vdash \varphi \vee \psi @ g$ can be transformed into a proof of either $\Gamma \vdash \varphi @ g$ or $\Gamma \vdash \psi @ g$.⁷

Recall that we created the label `Friends` in order to represent Bob’s policy “I will only share this with my friends.” However, we never showed how to connect this with Bob’s friend list, expressed as the relation $\text{IsFriend}(\text{Principal})$. In order to make this connection, we use a bi-implication

$$\forall p: \text{Principal}. \left(\text{Bob says}_{\text{Friends}} \text{IsFriend}(p) \leftrightarrow \text{Bob says}_{\text{Friends}} \text{CanRead}(p, \text{Friends}) \right)$$

It might seem strange to have an implication from $\text{CanRead}(p, \text{Friends})$ to $\text{IsFriend}(p)$. After all, this seems to suggest that, if a principal can read Bob’s `Friends` label, then Bob is going to be willing to consider them a friend, even if they were not on his friend list. However, note that in Γ_L we have a belief that says that if a principal is a friend, then they are one of the principals on Bob’s friend list. Thus, the troublesome implication above actually suggests that only those principals in L can read the label `Friends`.

While we can decide whether Bob believes that somebody is on his friend list—and therefore whether they can read things labeled `Friends`—when Alice tries to look at Bob’s picture the system needs to be able to tell Alice whether she is allowed to do so or not. We informally argued earlier that this was impossible. To see why, imagine that there is a label ℓ that we know Alice can read. Since we don’t know if Alice can read things of label `Friends`, we assume that $\Gamma_L, \Gamma \not\vdash \text{Friends} \sqsubseteq \ell @ \langle \rangle \cdot \text{Bob}(\ell)$. But then it is not the case that $\langle \rangle \cdot \text{Bob}(\text{Friends}) \text{ SF } \langle \rangle \cdot \text{Bob}(\ell)$, so the same holds for CanInfl in the absence of irrelevant implications. Now if we have any proof of

$$\Gamma_L, \Gamma \vdash \forall p: \text{Principal}. \left(\text{Bob says}_{\ell} \text{IsFriend}(p) \vee \text{Bob says}_{\ell} (\neg \text{IsFriend}(p)) \right)$$

by Theorem 7, we could remove Bob’s friend list and get a proof under only Γ , which is clearly impossible.

⁷This has been proven in Coq as part of the cut-elimination proof.

As we discussed in Section 2.1, this suggests that Bob should choose a more-permissive policy for his friend list. One possibility is for Bob to label his friend list publicly, but this is not a very satisfying solution. Another possibility is for Bob to allow any principal p to know whether p is on the list, but to not allow any principal p that is not on the list to know the status of any other principal. However, we did not discuss how to represent this policy in FLAFOL.

One simple way to represent this policy is using an implication. That is, we can assume the following:

$$\begin{aligned} & \forall p: \text{Principal}. \\ & \left(\text{Bob says}_{\text{Friends}} \text{IsFriend}(p) \right. \\ & \quad \left. \rightarrow p \text{ says}_{\text{Friends}} (\text{Bob says}_{\text{Friends}} \text{IsFriend}(p)) \right) \\ & \wedge \left(\text{Bob says}_{\text{Friends}} \neg \text{IsFriend}(p) \right. \\ & \quad \left. \rightarrow p \text{ says}_{\text{Friends}} (\text{Bob says}_{\text{Friends}} \neg \text{IsFriend}(p)) \right) \end{aligned}$$

This allows p to know whether p is on Bob's friend list. However, it is rather unsatisfying, because it destroys all of the security guarantees of Theorem 7. After all, we can now show that $\langle \rangle \cdot \text{Bob} \langle \text{Friends} \rangle \text{CanInfl} \langle \rangle \cdot p \langle \text{Friends} \rangle$ for any p .

We can create a more-subtle version of this policy which still enjoys the guarantees of Theorem 7. To do this, Bob labels his belief about whether or not a principal p is on his friend list at a label $f(p)$ that p may read and Bob's friends may read, but no one else. We thus create a function symbol $f : \text{Principal} \rightarrow \text{Label}$ and use it to define Bob's friend list. Now we can re-define Γ_L as follows:

$$\Gamma_L = \{ \text{Bob says}_{f(p)} \text{IsFriend}(p) \mid p \in L \} \cup \left\{ \begin{array}{l} \forall q: \text{Principal}. \\ \text{Bob says}_{f(q)} \text{IsFriend}(q) \\ \rightarrow \left(\bigvee_{p \in L} \text{Bob says}_{f(q)} q = p \right) \end{array} \right\}$$

If Γ contains axioms about equality, then

$$\Gamma_L, \Gamma \vdash \forall p: \text{Principal}. \left(\text{Bob says}_{f(p)} \text{IsFriend}(p) \right. \\ \left. \vee \text{Bob says}_{f(p)} \neg \text{IsFriend}(p) \right).$$

We now need to formalize the statement that p and Bob's friends may read $f(p)$, but nobody else. We do this with the following assumptions Γ_f .

$$\Gamma_f = \left\{ \begin{array}{l} \forall p: \text{Principal}. \text{Bob says}_{f(p)} \text{CanRead}(p, f(p)), \\ \forall p: \text{Principal}. \text{Bob says}_{\text{Friends}} f(p) \sqsubseteq \text{Friends}, \\ \forall p: \text{Principal}. \left(\text{Bob says}_{f(p)} \text{IsFriend}(p) \right. \\ \left. \leftrightarrow \text{Bob says}_{f(p)} \text{CanRead}(p, \text{Friends}) \right) \end{array} \right\}$$

These rules allow Bob to forward to p the results of the above decision procedure on $\text{Bob says}_{f(p)} \text{IsFriend}(p)$ if p will listen—which requires $p \text{ says}_{f(p)} \text{CanWrite}(\text{Bob}, f(p))$. Similarly, if $p \in L$, then Bob will forward whether or not q is Bob's friend for any principal q at label Friends .

A.2 Hospital Bills Calculation and Reinsurance

Recall the example from Section 2.3: Alice has two possible insurers, I_1 and I_2 . Bob is trying to figure out which will be allowed to influence Alice's hospital bill, labeled ℓ_H . He represents the fact that either I_1 or I_2 will be able to influence the bill as $\text{Bob says}_{\ell_H} (\text{CanWrite}(I_1, \ell_H) \vee \text{CanWrite}(I_2, \ell_H))$. He knows that I_2 reinsures with I_1 , which we represent as $\text{Bob says}_{\ell_H} (I_2 \text{ says}_{\ell_H} \text{CanWrite}(I_1, \ell_H))$.

$$\begin{array}{c}
\text{Ax} \frac{}{\Gamma', \varphi \vdash \varphi} \\
\frac{\frac{\Gamma', \text{CanWrite}(I_2, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H), \text{CanWrite}(I_1, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H) \vdash \varphi}{\Gamma', \text{CanWrite}(I_2, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H), \text{CanWrite}(I_1, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H) \cdot \text{Bob}(\ell_H) \vdash \varphi} \text{Ax}}{\Gamma', \text{CanWrite}(I_2, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H) \vdash \varphi} \text{SELF} \\
\text{ORL} \frac{}{\Gamma \vdash \varphi} \frac{\Gamma', \text{CanWrite}(I_2, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H) \vdash \varphi}{\Gamma \vdash \varphi} \text{FWDL}^\dagger
\end{array}$$

Figure 12: Bob’s proof that I_1 can influence the bill

In Section 2.3, we did not discuss the confidentiality requirements of this situation. In this case, since both insurers know that Bob is a hospital administrator, they are willing to talk to Bob about ℓ_H . Bob moreover knows this. Therefore, we can use $\text{Bob says}_{\ell_H} (I_2 \text{ says}_{\ell_H} \text{CanRead}(\text{Bob}, \ell_H))$, which we will need in the proof.

We can then formalize this example as a proof of the sequent $\Gamma \vdash \text{CanWrite}(I_1, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H)$ where

$$\Gamma = \left\{ \begin{array}{l} \left(\text{CanWrite}(I_1, \ell_H) \right) @ \langle \rangle \cdot \text{Bob}(\ell_H), \\ \left(\vee \text{CanWrite}(I_2, \ell_H) \right) @ \langle \rangle \cdot \text{Bob}(\ell_H), \\ \text{CanWrite}(I_1, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H) \cdot I_2(\ell_H), \\ \text{CanRead}(\text{Bob}, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H) \cdot I_2(\ell_H) \end{array} \right\}$$

Note that we have moved from **says** statements to generalized principals here. This is conceptually the same, and simply requires less work to move all of the says statements to the generalized principal in the proof. The formal proof is available in Figure 12, where we use a few pieces of shorthand for brevity and readability. First we refer to the belief $\text{CanWrite}(I_1, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H)$ as φ . Second, we let

$$\Gamma' = \left\{ \begin{array}{l} \text{CanWrite}(I_1, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H) \cdot I_2(\ell_H), \\ \text{CanRead}(\text{Bob}, \ell_H) @ \langle \rangle \cdot \text{Bob}(\ell_H) \cdot I_2(\ell_H) \end{array} \right\}$$

Finally, we do not explicitly state side conditions which are proven straightforwardly from Γ . The rules where these side conditions should appear are marked with “ \dagger ”.

B Examples of Permission Models

In Section 4 we saw how FLAFOL can be used to reason about a capabilities-based system. However, FLAFOL’s flexibility allows it to model many other kinds of systems. In this appendix, we explore modeling two other systems in FLAFOL: a simple system with no additional assumptions, and a system similar to military classification levels or FLAM’s system.

There is no particular reason for there to be some external model of permissions. The “default” permission model simply gives meaning to **CanRead** and **CanWrite** through their behavior. That is, the only properties FLAFOL assumes about **CanRead** and **CanWrite** are variance constraints, while all other properties of **CanRead** and **CanWrite** come from formulae in the context of a proof. This is appropriate in many cases. For instance, in the example of viewing photos on social media, **CanRead** and **CanWrite** have their behavior tuned by Bob’s selections on his account settings page. It is appropriate for the behaviors based on the selections to be axiomatized directly, rather than forced into some other model. Note that since we only care about confidentiality in that example, **CanWrite** can have a trivial implementation:

$$p \text{ says}_{\ell'} \text{CanWrite}(q, \ell) \leftrightarrow \text{True}.$$

FLAFOL can encode a more-concrete possible permission model by assigning every principal a label representing “which data this person is allowed to read or write.” This model appears in the real world in the U.S. military, where every person has a clearance label, and they are allowed to read documents labeled at or below their clearance. A more subtle version of this model separates reading and writing into confidentiality and integrity labels and allows every principal to have their own idea of each person’s label. This is similar to FLAM’s model, though our version is typed and does not force principals and labels to be the same.

We can formalize this by giving projection functions from both principals and labels to both confidentiality and integrity. The $\pi_{P,C}$ and $\pi_{P,I}$ projections take principals and produce confidentiality and integrity, respectively, and $\pi_{L,C}$ and $\pi_{L,I}$ do the same, but with labels as arguments. We can think of $\pi_{P,C}(p)$ as “the most confidential data that p can read,” while $\pi_{P,I}(p)$ is “the highest integrity data that p can write.” We think of $\pi_{L,C}(\ell)$ as “the confidentiality component of label ℓ ,” while $\pi_{L,I}(\ell)$ is “the integrity component of label ℓ .” With these functions, we can say that

$$\begin{aligned}
 p \text{ says}_{\ell'} \text{ CanRead}(q, \ell) &\leftrightarrow p \text{ says}_{\ell'} (\pi_{L,C}(\ell) \sqsubseteq \pi_{P,C}(q)), \\
 &\text{and} \\
 p \text{ says}_{\ell'} \text{ CanWrite}(q, \ell) &\leftrightarrow p \text{ says}_{\ell'} (\pi_{P,I}(q) \sqsubseteq \pi_{L,I}(\ell)).
 \end{aligned}$$

The reversal of the order here comes from the fact that integrity, as a flow ordering, is dual to confidentiality.

C Details of the Coq proofs

In this appendix, we give some basic guidance to the Coq code, available at <https://github.com/FLAFOL/flafol-coq>.

General Structure. In the file `Term.v` we define the term language used by FLAFOL along with its type system. In the same file the module `GroundInfo` is defined. This module takes as parameters information necessary to instantiate the recipe specified in Section 3. For instance, it assumes the existence of a type of sorts and the existence of two sorts: `Principal` and `Label`. It also assumes that fresh variables can be generated and that equality of function and relation symbols is decidable. In the file `Formula.v` we define FLAFOL formulae. To simplify proofs, we use a locally nameless representation of variables [Cha12] and binding, and we prove some basic results about this binding discipline. Note also that the definition of FLAFOL formulae is slightly different than that in the paper; rather than being part of the set \mathcal{R} , the permission relations are baked into the syntax of FLAFOL formulae directly.

We define the FLAFOL proof system in the file `Sequent.v`. There are three ways in which our Coq formalism differs from the presentation of FLAFOL in Section 4: (1) we use an equivalent presentation of the structural rules, (2) we use a slightly more general logic, and (3) we use two representations of the logic.

First, as is suggested by Pfenning [Pfe95], we drop the structural rules from the logic (`WEAKENING`, `EXCHANGE` and `CONTRACTION`), modify our rules so that they never erase anything from the context and we prove that the removed rules are admissible. This makes meta-theoretic proofs simpler.

Second, the logic described in the Coq is slightly more general than the one described in the paper. In the Coq version the ground generalized principal has a label attached to it. Originally we added ground-level labels to accommodate features that we left for future work, but we do not need them for this version of FLAFOL. To show that this is a generalization, for any FLAFOL proof without ground labels, we can simply assign the same ground label to every belief in the proof and acquire a valid proof in the Coq version.

Third, we have two representations of our logic. The first is an (untyped) term language with the appropriate typing rules, and the second is a dependent inductive type. The untyped version eases reasoning about equality, reduces compilation time, and makes proving the admissibility of weakening and substitution

easier. The typed version is easier to write automation tactics for. We have proved that both representations are equivalent.

Details of Cut Admissibility Proof. In `NormalForm.v` we define a normal form for FLAFOL proofs. The cut-elimination procedure uses normalization as an essential step.⁸ A proof is in normal form if all rules which do not manipulate formulae are higher in the proof tree than those which do. Formally, we define 2 normal forms, first and second normal form, which represent “might use formula-manipulating rules” and “will not use formula manipulating rules”, respectively. A proof is in first normal form if, when a rule which manipulates something other than a formula is used, all subproofs above that rule are in second normal form, while a proof is in second normal form if it never uses any rules which manipulate formulae. The main result in this file is that every FLAFOL proof has a normal form.

Theorem 8 (FLAFOL Normal Form). *If $\Gamma \vdash \varphi @ g$ is provable in FLAFOL, then it is provable with a proof in normal form.*

Lastly the file `Cut.v` contains the cut-elimination procedure. First we normalize both proofs. If they’re both in First Normal Form but not in Second Normal Form, we proceed as Pfenning suggests in [Pfe95]: nested triple induction on the formula being cut and on both proofs. If one of them is in Second Normal Form we use a different procedure. This procedure consists of getting the dual rule to the last rule used in the proof that is in Second Normal Form (e.g. `VARL` for the `VARR` case) and make it the last rule to the other proof. Due to the covariant-contravariant nature of these rules and their duals, this is always possible. For more details see lemmas `Cut_h1MCR` and `Cut_h2MCR` in `Cut.v`

Non-Interference. In `Speaksfor.v` we define the relations `SF`, `CanInfl`, define the function `G` and prove Theorem 6. The compatible supercontexts rules are defined in `CompatibleSuperContext.v`. Finally, the Coq proof of Non-Interference is in `Noninterference.v`; it closely follows the pen-and-paper proof sketched in Section 6.

Simulation. The file `Simulation.v` contains the definition of the function `⊙` a proof of the Simulation Theorem (Theorem 3).

D The Full FLAFOL Proof System

The full FLAFOL proof system can be found in Figure 13.

E Compatible Supercontexts

Figure 14 contains the full rules for compatible super-contexts.

⁸In the literature, “normal proof” refers to a cut-free proof, rather than a proof in FLAFOL’s normal form.

$$\begin{array}{c}
\text{AX} \frac{}{\Gamma, \varphi @ g \vdash \varphi @ g} \\
\text{CONTRACTION} \frac{\Gamma, (\varphi @ g), (\varphi @ g) \vdash \psi @ g'}{\Gamma, \varphi @ g \vdash \psi @ g'} \\
\text{FALSEL} \frac{}{\Gamma, \mathbf{False} @ g \vdash \varphi @ g \cdot g'} \\
\text{ANDL} \frac{\Gamma, (\varphi @ g), (\psi @ g) \vdash \chi @ g'}{\Gamma, (\varphi \wedge \psi @ g) \vdash \chi @ g'} \\
\text{ORL} \frac{\Gamma, \varphi @ g \vdash \chi @ g' \quad \Gamma, \psi @ g \vdash \chi @ g'}{\Gamma, (\varphi \vee \psi @ g) \vdash \chi @ g'} \\
\text{IMPL} \frac{\Gamma \vdash \varphi @ \langle \rangle \quad \Gamma, \psi @ g \vdash \chi @ g'}{\Gamma, (\varphi \rightarrow \psi @ g) \vdash \chi @ g'} \\
\text{FORALLL} \frac{\Gamma, \varphi[x \mapsto t] @ g \vdash \psi @ g'}{\Gamma, (\forall x: \sigma. \varphi @ g) \vdash \psi @ g'} \\
\text{EXISTSL} \frac{\Gamma, \varphi @ g \vdash \psi @ g' \quad x \notin \mathbf{FV}(\Gamma, \psi, g, g')}{\Gamma, (\exists x: \sigma. \varphi @ g) \vdash \psi @ g'} \\
\text{SAYSL} \frac{\Gamma, \varphi @ g \cdot p \langle \ell \rangle \vdash \psi @ g'}{\Gamma, p \mathbf{says}_\ell \varphi @ g \vdash \psi @ g'} \\
\text{SELF L} \frac{\Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot g') \vdash \psi @ g''}{\Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot p \langle \ell \rangle \cdot g') \vdash \psi @ g''} \\
\text{VARL} \frac{\Gamma, (\varphi @ g \cdot p \langle \ell' \rangle \cdot g') \vdash \psi @ g'' \quad \Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot g') \vdash \ell \sqsubseteq \ell' @ g \cdot p \langle \ell \rangle}{\Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot g') \vdash \psi @ g''} \\
\text{FWDL} \frac{\Gamma, (\varphi @ g \cdot q \langle \ell \rangle \cdot g') \vdash \chi @ g'' \quad \Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot g') \vdash \mathbf{CanRead}(q, \ell) @ g \cdot p \langle \ell \rangle \quad \Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot g') \vdash \mathbf{CanWrite}(p, \ell) @ g \cdot q \langle \ell \rangle}{\Gamma, \varphi @ g \cdot p \langle \ell \rangle \cdot g' \vdash \chi @ g''} \\
\text{FLOWSTOREFL} \frac{}{\Gamma \vdash \ell \sqsubseteq \ell @ g} \\
\text{CRVAR} \frac{\Gamma \vdash \mathbf{CanRead}(p, \ell_2) @ g \quad \Gamma \vdash \ell_1 \sqsubseteq \ell_2 @ g}{\Gamma \vdash \mathbf{CanRead}(p, \ell_1) @ g} \\
\text{CWVAR} \frac{\Gamma \vdash \mathbf{CanWrite}(p, \ell_2) @ g \quad \Gamma \vdash \ell_2 \sqsubseteq \ell_1 @ g}{\Gamma \vdash \mathbf{CanWrite}(p, \ell_1) @ g} \\
\text{WEAKENING} \frac{\Gamma \vdash \psi @ g}{\Gamma, \varphi @ g' \vdash \psi @ g} \\
\text{EXCHANGE} \frac{\Gamma, (\varphi @ g_1), (\psi @ g_2), \Gamma' \vdash \chi @ g}{\Gamma, (\psi @ g_2), (\varphi @ g_1), \Gamma' \vdash \chi @ g} \\
\text{TRUER} \frac{}{\Gamma \vdash \mathbf{True} @ g} \\
\text{ANDR} \frac{\Gamma \vdash \varphi @ g \quad \Gamma \vdash \psi @ g}{\Gamma \vdash \varphi \wedge \psi @ g} \\
\text{ORR1} \frac{\Gamma \vdash \varphi @ g}{\Gamma \vdash \varphi \vee \psi @ g} \quad \text{ORR2} \frac{\Gamma \vdash \psi @ g}{\Gamma \vdash \varphi \vee \psi @ g} \\
\text{IMPR} \frac{\Gamma, \varphi @ \langle \rangle \vdash \psi @ g}{\Gamma \vdash \varphi \rightarrow \psi @ g} \\
\text{FORALLR} \frac{\Gamma \vdash \varphi @ g \quad x \notin \mathbf{FV}(\Gamma, g)}{\Gamma \vdash \forall x: \sigma. \varphi @ g} \\
\text{EXISTSR} \frac{\Gamma \vdash \varphi[x \mapsto t] @ g}{\Gamma \vdash \exists x: \sigma. \psi @ g} \\
\text{SAYSR} \frac{\Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle}{\Gamma \vdash p \mathbf{says}_\ell \varphi @ g} \\
\text{SELF R} \frac{\Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle \cdot g'}{\Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle \cdot p \langle \ell \rangle \cdot g'} \\
\text{VARR} \frac{\Gamma \vdash \varphi @ g \cdot p \langle \ell' \rangle \cdot g' \quad \Gamma \vdash \ell' \sqsubseteq \ell @ g \cdot p \langle \ell \rangle}{\Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle \cdot g'} \\
\text{FWDR} \frac{\Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle \cdot g' \quad \Gamma \vdash \mathbf{CanRead}(q, \ell) @ g \cdot p \langle \ell \rangle \quad \Gamma \vdash \mathbf{CanWrite}(p, \ell) @ g \cdot q \langle \ell \rangle}{\Gamma \vdash \varphi @ g \cdot q \langle \ell \rangle \cdot g'}
\end{array}$$

Figure 13: Full FLAFOL Proof System

$$\begin{array}{c}
\text{CSCREFL} \frac{}{\Gamma \ll \Gamma \vdash \varphi @ g} \qquad \text{CSCUNION} \frac{\Delta_1 \ll \Gamma \vdash \varphi @ g \quad \Delta_2 \ll \Gamma \vdash \varphi @ g}{\Delta_1 \cup \Delta_2 \ll \Gamma \vdash \varphi @ g} \\
\text{CSCCONTRACTION} \frac{\Delta \ll \Gamma, (\varphi @ g), (\varphi @ g) \vdash \psi @ g'}{\Delta \ll \Gamma, \varphi @ g \vdash \psi @ g'} \\
\text{CSCEXCHANGE} \frac{\Delta \ll \Gamma, (\varphi @ g_1), (\psi @ g_2), \Gamma' \vdash \chi @ g}{\Delta \ll \Gamma, (\psi @ g_2), (\varphi @ g_1), \Gamma' \vdash \chi @ g} \qquad \text{CSCANDL} \frac{\Delta \ll \Gamma, (\varphi @ g), (\psi @ g) \vdash \chi @ g'}{\Delta \ll \Gamma, (\varphi \wedge \psi @ g) \vdash \chi @ g'} \\
\text{CSCANDR1} \frac{\Delta \ll \Gamma \vdash \varphi @ g}{\Delta \ll \Gamma \vdash \varphi \wedge \psi @ g} \qquad \text{CSCANDR2} \frac{\Delta \ll \Gamma \vdash \psi @ g}{\Delta \ll \Gamma \vdash \varphi \wedge \psi @ g} \\
\text{CSCORL1} \frac{\Delta \ll \Gamma, \varphi @ g \vdash \chi @ g'}{\Delta \ll \Gamma, (\varphi \vee \psi @ g) \vdash \chi @ g'} \qquad \text{CSCORL2} \frac{\Delta \ll \Gamma, \psi @ g \vdash \chi @ g'}{\Delta \ll \Gamma, (\varphi \vee \psi @ g) \vdash \chi @ g'} \\
\text{CSCORR1} \frac{\Delta \ll \Gamma \vdash \varphi @ g}{\Delta \ll \Gamma \vdash \varphi \vee \psi @ g} \qquad \text{CSCORR2} \frac{\Delta \ll \Gamma \vdash \psi @ g}{\Delta \ll \Gamma \vdash \varphi \vee \psi @ g} \\
\text{CSCIMPL1} \frac{\Delta \ll \Gamma, \psi @ g \vdash \chi @ g'}{\Delta \ll \Gamma, (\varphi \rightarrow \psi @ g) \vdash \chi @ g'} \qquad \text{CSCIMPL2} \frac{\Delta \ll \Gamma \vdash \varphi @ \langle \rangle}{\Delta \ll \Gamma, (\varphi \rightarrow \psi @ g) \vdash \chi @ g'} \\
\text{CSCIMPR} \frac{\Delta \ll \Gamma, \varphi @ \langle \rangle \vdash \psi @ g}{\Delta \ll \Gamma \vdash \varphi \rightarrow \psi @ g} \qquad \text{CSCFORALL} \frac{\Delta \ll \Gamma, \varphi[x \mapsto t] @ g \vdash \psi @ g'}{\Delta \ll \Gamma, (\forall x:\sigma. \varphi @ g) \vdash \psi @ g'} \\
\text{CSCFORALLR} \frac{\Delta \ll \Gamma \vdash \varphi @ g \quad x \notin \text{FV}(\Gamma, g)}{\Delta \ll \Gamma \vdash \forall x:\sigma. \varphi @ g} \\
\text{CSCEXISTSL} \frac{\Delta \ll \Gamma, \varphi @ g \vdash \psi @ g' \quad x \notin \text{FV}(\Gamma, \psi, g, g')}{\Delta \ll \Gamma, (\exists x:\sigma. \varphi @ g) \vdash \psi @ g'} \qquad \text{CSCEXISTSR} \frac{\Delta \ll \Gamma \vdash \varphi[x \mapsto t] @ g}{\Delta \ll \Gamma \vdash \exists x:\sigma. \varphi @ g} \\
\text{CSCSAYSL} \frac{\Delta \ll \Gamma, \varphi @ g \cdot p \langle \ell \rangle \vdash \psi @ g'}{\Delta \ll \Gamma, p \text{ says}_\ell \varphi @ g \vdash \psi @ g'} \qquad \text{CSCSAYSR} \frac{\Delta \ll \Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle}{\Delta \ll \Gamma \vdash p \text{ says}_\ell \varphi @ g} \\
\text{CSCSELF L} \frac{\Delta \ll \Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot g') \vdash \psi @ g''}{\Delta \ll \Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot p \langle \ell \rangle \cdot g') \vdash \psi @ g''} \qquad \text{CSCSELF R} \frac{\Delta \ll \Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle \cdot g'}{\Delta \ll \Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle \cdot p \langle \ell \rangle \cdot g'} \\
\text{CSCVARL} \frac{\Delta \ll \Gamma, (\varphi @ g \cdot p \langle \ell' \rangle \cdot g') \vdash \psi @ g'' \quad \Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot g') \vdash \ell \sqsubseteq \ell' @ g \cdot p \langle \ell \rangle}{\Delta \ll \Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot g') \vdash \psi @ g''} \\
\text{CSCVARR} \frac{\Delta \ll \Gamma \vdash \varphi @ g \cdot p \langle \ell' \rangle \cdot g' \quad \Gamma \vdash \ell' \sqsubseteq \ell @ g \cdot p \langle \ell \rangle}{\Delta \ll \Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle \cdot g'} \\
\text{CSCFwDL} \frac{\Delta \ll \Gamma, (\varphi @ g \cdot q \langle \ell \rangle \cdot g') \vdash \chi @ g'' \quad \Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot g') \vdash \text{CanRead}(q, \ell) @ g \cdot p \langle \ell \rangle}{\Gamma, (\varphi @ g \cdot p \langle \ell \rangle \cdot g') \vdash \text{CanWrite}(p, \ell) @ g \cdot q \langle \ell \rangle} \\
\text{CSCFwDR} \frac{\Delta \ll \Gamma \vdash \varphi @ g \cdot p \langle \ell \rangle \cdot g' \quad \Gamma \vdash \text{CanRead}(q, \ell) @ g \cdot p \langle \ell \rangle \quad \Gamma \vdash \text{CanWrite}(p, \ell) @ g \cdot q \langle \ell \rangle}{\Delta \ll \Gamma \vdash \varphi @ g \cdot q \langle \ell \rangle \cdot g'}
\end{array}$$

Figure 14: Compatible Supercontext Rules