

RIF: Reactive Information Flow Labels*

Elisavet Kozyri
ekozyri@cs.cornell.edu

Fred B. Schneider
fbs@cs.cornell.edu

Department of Computer Science
Cornell University
Ithaca, New York 14853

April 8, 2019

Abstract

Restrictions that a *reactive information flow* (RIF) label imposes on a value are determined by the sequence of operations used to derive that value. This allows declassification, endorsement, and other forms of reclassification to be supported in a uniform way. *Piecewise noninterference* (PWNI) is introduced as a fitting security policy, because noninterference is not suitable. A type system is given for static enforcement of PWNI in programs that associate *checkable* classes of RIF labels with variables. Two checkable classes of RIF labels are described: *RIF automata* are general-purpose and based on finite-state automata; *κ -labels* concern confidentiality in programs that use cryptographic operations.

*Supported in part by AFOSR grant F9550-16-0250 and NSF grant 1642120. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of these organizations or the U.S. Government.

1 Introduction

In Denning’s initial work [20] and in much that has followed—both for confidentiality and integrity—the set of restrictions assigned to the output of an operation is the union of the restrictions associated with its inputs. But by ignoring the operator and the values of inputs, that approach can be too conservative. The most general formulation of *flow-derived restrictions* would assign restrictions to the output of an operation $op(x_1, x_2, \dots, x_n)$ according to operator op , its inputs x_1, x_2, \dots, x_n , and the restrictions associated with those inputs. That output might warrant fewer restrictions, additional restrictions, or an incomparable set of restrictions than are associated with its inputs.

- With an operation that computes the winner of an election, the inputs are votes and the output is the majority. Each input is secret to the principal casting that vote, whereas the output ought to be readable by any principal. So the output should be associated with fewer restrictions than the inputs.
- A conference-management system matches papers to reviewers, where that matching is generated by a non-deterministic computation. The inputs—a list of reviewers and a list of submissions—can be read by the entire program committee, but conflicts of interest dictate that only a subset of the program committee learn which reviewers are assigned to any given paper. So outputs are associated with more stringent restrictions than inputs.

Previous work on information flow control—declassification and erasure policies [15], information flow locks [10, 9], typed declassification [23], expressions for declassification (for confidentiality) and endorsement (for integrity) [41, 43], and capability-based mechanisms for downgrading security policies [33, 46, 52]—does not support arbitrary changes in restrictions linked with specific operations. Other approaches (e.g., [29, 38, 44, 45, 49]) allow such changes but only between two levels of labels (e.g., **public** and **secret**).

Reactive information flow (RIF) labels, which we introduce in §2, seeks to address these limitations by allowing stronger, weaker, or incomparable restrictions to be associated with the output of an operation, as determined by the operator and the restrictions on inputs. Piecewise noninterference (PWNI) described in §3 then extends classical noninterference in a way that

handles changes to restrictions that RIF labels support. A type system is given in §4 to support static enforcement of PWNI for certain classes of RIF labels. Examples of those classes include RIF automata and κ -labels, which are described in §5. And §6 discusses how previous work relates to RIF labels and PWNI.

2 RIF Labels

Restrictions. Restrictions are assumed to be elements of a join semilattice $\langle R, \sqcup_R, \sqsubseteq_R \rangle$.¹ For confidentiality, an element of R might identify which principals are allowed to read some value, either by enumerating that set of principals or by giving the name (e.g., **public**, **secret**, etc.) for such a set; for integrity, it might identify the set of principals allowed to write that value. Other kinds of restrictions might be specified using elements in R that are a set of programs, or the name for a task or set of tasks.

Relation $r \sqsubseteq_R r'$ is satisfied if compliance with restrictions $r' \in R$ implies compliance with restrictions $r \in R$, so r' is *at least as strong* as r or, equivalently, r is *at least as weak* as r' . When elements of R denote sets of principals, then for confidentiality we would define $r \sqcup_R r'$ to be $r \cap r'$ and define $r \sqsubseteq_R r'$ to hold if and only if $r \supseteq r'$ holds—if r' allows a principal to read then so does r , but r might also allow other principals to read, too. And, for integrity, we would define $r \sqcup_R r'$ to be $r \cup r'$ and define $r \sqsubseteq_R r'$ to hold if and only if $r \subseteq r'$ holds.

Reclassifiers and RIF Labels. A *reclassifier* abstracts how an operation changes the restrictions for an argument.² To associate reclassifiers with operations, we extend a language that defines *ordinary expressions*: variables and terms $op(\mathcal{E}_1, \dots, \mathcal{E}_n)$, where op is an operator and each \mathcal{E}_i is an ordinary expression. For ordinary expressions $\mathcal{E}_1, \dots, \mathcal{E}_n$,

$$[op(\mathcal{E}_1, \dots, \mathcal{E}_n)]_{f_1, f_2, \dots, f_n} \tag{1}$$

defines a *reclassifying expression*. It specifies that reclassifier f_i identifies how the restrictions associated with the value of ordinary expression \mathcal{E}_i should

¹A partial order would suffice for the theory of RIF labels developed in this section. A join semilattice becomes useful for certain instantiations of RIF labels, as illustrated in §5.

²The term *reclassify* was used in Denning’s thesis [20] to describe an operation that changes the restrictions imposed on objects.

be changed for constructing the restrictions associated with the value produced by $op(\mathcal{E}_1, \dots, \mathcal{E}_n)$. Notation $[op(\mathcal{E}_1, \dots, \mathcal{E}_n)]_f$ is used as a shorthand for $[op(\mathcal{E}_1, \dots, \mathcal{E}_n)]_{f, f, \dots, f}$, and we sometimes abbreviate $[op(\mathcal{E}_1, \dots, \mathcal{E}_n)]_{f_1, f_2, \dots, f_n}$ by $[op(\mathcal{E}_1, \dots, \mathcal{E}_n)]_{\bar{f}}$ or simply $[\mathcal{E}]_{\bar{f}}$ if the elided specifics are irrelevant.

When reclassifying expressions are used to compute values, then sequences of reclassifiers offer abstract descriptions for the series of operations that have been applied to values as program execution proceeds. Such a sequence of reclassifiers then provides a basis for determining the set of restrictions associated with a computed value. For example, consider the following program.

$$w_1 := [div(x_1, x_2)]_{\bar{f}}; \quad y_2 := [mod(w_1, w_2)]_{f'}; \quad z_3 := [add(y_1, y_2, y_3)]_{f''}$$

Here, restrictions on the value stored in z_3 are derived from:

- the restrictions on the value stored in x_1 and in x_2 , changed according to sequence $f'f''$ because each of these values flows to w_1 through f , then w_1 flows to y_2 through f' , and finally y_2 flows to z_3 through f'' .
- the restrictions on the value stored in w_2 , changed according to sequence $f'f''$, and
- the restrictions on the values stored in y_1 and y_3 , changed according to f'' .

A RIF label specifies a set of restrictions for an associated value v as well as for values derived by executing operations on v . Formally, a RIF label maps sequences of reclassifiers to elements of underlying join semilattice $\langle R, \sqcup_R, \sqsubseteq_R \rangle$ of restrictions. This mapping is specified for a set Λ of RIF labels by giving a set \mathcal{F} of reclassifiers along with two functions \mathcal{R} and \mathcal{T} .

- \mathcal{R} maps $\lambda \in \Lambda$ to the restriction $r \in R$ that λ currently imposes:

$$\mathcal{R}: \Lambda \rightarrow R \tag{2}$$

- \mathcal{T} maps $\lambda \in \Lambda$ and each reclassifier $f \in \mathcal{F}$ to a RIF label that should be associated with the value produced by an operation f abstracts:

$$\mathcal{T}: \Lambda \times \mathcal{F} \rightarrow \Lambda \tag{3}$$

The restrictions associated with the value of a reclassifying expression (1) will incorporate restrictions $\mathcal{R}(\mathcal{T}(\lambda_i, f_i))$ derived from RIF label λ_i being associated with the value of expression \mathcal{E}_i .

\mathcal{T} is extended to a finite sequence³ $F \in \mathcal{F}^*$ of reclassifiers in the usual

³As is conventional, \mathcal{F}^* denotes the set of finite sequences of elements in \mathcal{F} .

way, with empty sequence ϵ of reclassifiers considered an element of every set \mathcal{F} of reclassifiers (where it serves as an identity reclassifier).

$$\mathcal{T}(\lambda, \epsilon) \triangleq \lambda \quad (4)$$

$$\mathcal{T}(\lambda, Ff) \triangleq \mathcal{T}(\mathcal{T}(\lambda, F), f) \quad (5)$$

Classes of RIF Labels. A class of RIF labels can serve as the basis for a join semilattice $\langle \Lambda, \sqcup_\Lambda, \sqsubseteq_\Lambda \rangle$, where the cardinality of Λ can be infinite. The *join* operator \sqcup_Λ is used for combining RIF labels; the *restrictiveness* relation \sqsubseteq_Λ specifies whether one RIF label is *at least as restrictive as* another. We posit that Λ includes elements \perp and \top such that for all $\lambda \in \Lambda$: $\perp \sqsubseteq_\Lambda \lambda$ and $\lambda \sqsubseteq_\Lambda \top$ hold.

Since, by definition, $\lambda \sqsubseteq_\Lambda \lambda \sqcup_\Lambda \lambda'$ and $\lambda' \sqsubseteq_\Lambda \lambda \sqcup_\Lambda \lambda'$ holds in a join semilattice, a combination of RIF labels is at least as restrictive as any of its constituents. It seems reasonable (although it is not required) that \sqcup_Λ and \sqcup_R would be related in a way that ensures restrictions imposed by $\lambda \sqcup_\Lambda \lambda'$ are the same as the combined restrictions imposed by individual RIF labels λ and λ' .⁴

$$\mathcal{R}(\lambda \sqcup_\Lambda \lambda') = \mathcal{R}(\lambda) \sqcup_R \mathcal{R}(\lambda') \quad (6)$$

The definition of restrictiveness relation \sqsubseteq_Λ depends on partial order \sqsubseteq_R for set R of restrictions, \mathcal{R} , and \mathcal{T} :

$$\lambda \sqsubseteq_\Lambda \lambda' \triangleq (\forall F \in \mathcal{F}^*: \mathcal{R}(\mathcal{T}(\lambda, F)) \sqsubseteq_R \mathcal{R}(\mathcal{T}(\lambda', F))) \quad (7)$$

So if $\lambda \sqsubseteq_\Lambda \lambda'$ holds, then

- current restrictions $\mathcal{R}(\lambda')$ specified by λ' are at least as strong as what λ imposes because, by definition, $\epsilon \in \mathcal{F}$ and thus $\epsilon \in \mathcal{F}^*$, $\mathcal{R}(\mathcal{T}(\lambda, \epsilon)) = \mathcal{R}(\lambda)$ and $\mathcal{R}(\mathcal{T}(\lambda', \epsilon)) = \mathcal{R}(\lambda')$ hold, so (7) implies that $\mathcal{R}(\lambda) \sqsubseteq_R \mathcal{R}(\lambda')$, and
- restrictions that λ' imposes for any derived value are at least as strong as what λ imposes—if v flows to w then there is a sequence $F \in \mathcal{F}^*$ that v flows through, and (7) requires that $\mathcal{R}(\mathcal{T}(\lambda, F)) \sqsubseteq_R \mathcal{R}(\mathcal{T}(\lambda', F))$ hold.

⁴Reasonable examples do exist where (6) does not hold. But the property does hold for the example RIF labels described in this paper.

	$\mathcal{R}(\lambda) \subset \mathcal{R}(\mathcal{T}(\lambda, \mathbf{f}))$	$\mathcal{R}(\lambda) \supset \mathcal{R}(\mathcal{T}(\lambda, \mathbf{f}))$
Confidentiality	<i>declassification</i>	<i>classification</i>
Integrity	<i>deprecation</i>	<i>endorsement</i>

Figure 1: Terminology for reclassifications

The quantification over all sequences $F \in \mathcal{F}^*$ in (7) means that this definition for $\lambda \sqsubseteq_{\Lambda} \lambda'$ is conservative, since it imposes conditions for sequences F of reclassifiers that never arise in a program execution.

If $\mathcal{T}(\lambda, \mathbf{f}) \neq \lambda$ holds, then \mathbf{f} triggers a *reclassification*. In the literature, reclassification is categorized based on the restrictions the new and old labels impose—specifically, whether $\mathcal{R}(\mathcal{T}(\lambda, \mathbf{f}))$ is weaker than $\mathcal{R}(\lambda)$. Figure 1 defines specialized terminology for this categorization where $\mathcal{R}(\lambda)$ gives a set of principals that must be trusted not to divulge the value (for confidentiality) or not to have corrupted the value (for integrity). That said, this paper benefits from using a finer-grained categorization. It is based on how a new label $\mathcal{T}(\lambda, \mathbf{f})$ is related to λ : When $\mathcal{T}(\lambda, \mathbf{f}) \sqsubseteq_{\Lambda} \lambda$ holds then we say that \mathbf{f} triggers a *downgrade*, whereas when $\mathcal{T}(\lambda, \mathbf{f}) \not\sqsubseteq_{\Lambda} \lambda$ holds then we say that \mathbf{f} triggers an *upgrade*.

A *class* of RIF labels is formed by putting the pieces together:

$$\langle\langle R, \sqcup_R, \sqsubseteq_R \rangle, \langle \Lambda, \sqcup_{\Lambda}, \sqsubseteq_{\Lambda} \rangle, \mathcal{F}, \mathcal{R}, \mathcal{T} \rangle$$

This definition is silent about the existence of algorithms for computing \sqcup_{Λ} or for deciding \sqsubseteq_{Λ} . We say that a class of RIF labels is *checkable* if and only if:

- (i) There exists an algorithm for computing $\lambda \sqcup_{\Lambda} \lambda'$ for all $\lambda, \lambda' \in \Lambda$.
- (ii) There exists a sound (no false positives) if not complete (false negatives possible) test for determining whether $\lambda \sqsubseteq_{\Lambda} \lambda'$ holds for all $\lambda, \lambda' \in \Lambda$.

The type system we give in §4 is decidable if its RIF labels are from a checkable class.

3 A Security Policy for RIF Labels

Our security policy for RIF labels is obtained by extending termination-insensitive noninterference [48].⁵ Noninterference applies to programs comprising deterministic commands, so execution of a command \mathcal{C}_1 that is started in a memory \mathcal{M}_1 and terminates in a memory \mathcal{M}_N can be described by giving a finite *trace*

$$\langle \mathcal{C}_1, \mathcal{M}_1 \rangle \rightarrow \langle \mathcal{C}_2, \mathcal{M}_2 \rangle \rightarrow \cdots \rightarrow \langle \bullet, \mathcal{M}_N \rangle \quad (8)$$

where each *state* $\langle \mathcal{C}_i, \mathcal{M}_i \rangle$ gives a command \mathcal{C}_i and a memory \mathcal{M}_i . We write \bullet in the final state of a trace to signify further execution is not possible, but \bullet is not considered a command. Thus the trace for a terminating command \mathcal{C} includes at least two states: one state having command \mathcal{C} followed by one having \bullet .

Memories are represented by functions that map variables x appearing in any command to values $\mathcal{M}(x)$. These functions are then extended as usual for mapping ordinary expressions and reclassifying expressions.

$$\begin{aligned} \mathcal{M}(op(\mathcal{E}_1, \dots, \mathcal{E}_n)) &\triangleq op(\mathcal{M}(\mathcal{E}_1), \dots, \mathcal{M}(\mathcal{E}_n)) \\ \mathcal{M}([\mathcal{E}]_{f_1, f_2, \dots, f_n}) &\triangleq \mathcal{M}(\mathcal{E}) \end{aligned}$$

An operational semantics for command execution can be given by a partial function Υ , where $\Upsilon(\mathcal{C}, \mathcal{M})$ equals the finite trace corresponding to an execution of \mathcal{C} that terminates when started in memory \mathcal{M} , and $\Upsilon(\mathcal{C}, \mathcal{M})$ is undefined if that execution does not terminate.

Observations as a Threat Model. Assume fixed mapping Γ associates a RIF label $\Gamma(x) \in \Lambda$ with each variable x , where Γ is extended to handle ordinary expressions as usual

$$\Gamma(op(\mathcal{E}_1, \dots, \mathcal{E}_n)) \triangleq \Gamma(\mathcal{E}_1) \sqcup_{\Lambda} \dots \sqcup_{\Lambda} \Gamma(\mathcal{E}_n) \quad (9)$$

so that $\Gamma(op(x_1, \dots, x_n)) = \Gamma(x_1) \sqcup_{\Lambda} \dots \sqcup_{\Lambda} \Gamma(x_n)$. The RIF label associated with a reclassifying expression combines RIF labels obtained after the indicated transitions have been performed:

$$\Gamma([\mathcal{E}]_{f_1, \dots, f_n}) \triangleq \mathcal{T}(\Gamma(\mathcal{E}_1), f_1) \sqcup_{\Lambda} \dots \sqcup_{\Lambda} \mathcal{T}(\Gamma(\mathcal{E}_n), f_n) \quad (10)$$

⁵Our approach for deriving these extensions, however, could be applied to variations of noninterference intended for non-deterministic and probabilistic programs.

Γ induces equivalence classes of memories that agree on the values of variables storing values that are allowed to flow to $\lambda \in \Lambda$:

$$\mathcal{M} =_\lambda \mathcal{M}' \triangleq (\forall x: \Gamma(x) \sqsubseteq_\Lambda \lambda \Rightarrow \mathcal{M}(x) = \mathcal{M}'(x))$$

Our *threat model* is intended for analyzing systems in which principals are co-resident and able to detect changes (subject to restrictions defined by RIF labels) being made to shared memory or other resources. We formalize this threat model relative to λ by constructing an *observation*

$$\mathcal{M}_{i+1} \ominus_\lambda \mathcal{M}_i \triangleq \{\langle x, \mathcal{M}_{i+1}(x) \rangle \mid \mathcal{M}_i(x) \neq \mathcal{M}_{i+1}(x) \wedge \Gamma(x) \sqsubseteq_\Lambda \lambda\}$$

for each *transition* $\langle \mathcal{C}_i, \mathcal{M}_i \rangle \rightarrow \langle \mathcal{C}_{i+1}, \mathcal{M}_{i+1} \rangle$ that occurs. So observation $\mathcal{M}_{i+1} \ominus_\lambda \mathcal{M}_i$ gives the new value for each variable that was (i) changed by the transition and (ii) has a RIF label at most λ .

It is easy to show:

$$\begin{aligned} \mathcal{M} =_\lambda \mathcal{M}' &\Rightarrow (\mathcal{M}' \ominus_\lambda \mathcal{M} = \emptyset) \\ \lambda \sqsubseteq_\Lambda \lambda' &\Rightarrow (\mathcal{M}' \ominus_\lambda \mathcal{M}) \subseteq (\mathcal{M}' \ominus_{\lambda'} \mathcal{M}) \end{aligned}$$

For a RIF label λ , program execution described by a trace τ results in a sequence $\tau|_\lambda$ of non-empty observations that are derived from the successive transitions in τ^6 and induces an equivalence relation on traces

$$\tau =_\lambda \tau' \triangleq \tau|_\lambda = \tau'|_\lambda.$$

Sequences $\tau|_\lambda$ of observations are the basis for our threat model. Each principal p is assigned a set $\mathcal{R}(p)$ of restrictions. The threat model stipulates that a principal p sees changes to any variable x satisfying $\mathcal{R}(\Gamma(x)) \sqsubseteq_R \mathcal{R}(p)$.

⁶A formal definition of $\tau|_\lambda$ is thus given by the following, where \rightarrow is used to delimit observations in the resulting sequence.

$$\tau|_\lambda \triangleq \begin{cases} \epsilon & \text{if } \tau = \epsilon \text{ or } \tau = \langle \mathcal{C}, \mathcal{M} \rangle \\ \epsilon & \text{if } \tau = \langle \mathcal{C}, \mathcal{M} \rangle \rightarrow \langle \mathcal{C}', \mathcal{M}' \rangle \wedge \mathcal{M}' \ominus_\lambda \mathcal{M} = \emptyset \\ \mathcal{M}' \ominus_\lambda \mathcal{M} & \text{if } \tau = \langle \mathcal{C}, \mathcal{M} \rangle \rightarrow \langle \mathcal{C}', \mathcal{M}' \rangle \wedge \mathcal{M}' \ominus_\lambda \mathcal{M} \neq \emptyset \\ (\langle \mathcal{C}', \mathcal{M}' \rangle \rightarrow \tau')|_\lambda & \text{if } \tau = \langle \mathcal{C}, \mathcal{M} \rangle \rightarrow \langle \mathcal{C}', \mathcal{M}' \rangle \rightarrow \tau' \wedge \mathcal{M}' \ominus_\lambda \mathcal{M} = \emptyset \\ \mathcal{M}' \ominus_\lambda \mathcal{M} \rightarrow (\langle \mathcal{C}', \mathcal{M}' \rangle \rightarrow \tau')|_\lambda & \text{if } \tau = \langle \mathcal{C}, \mathcal{M} \rangle \rightarrow \langle \mathcal{C}', \mathcal{M}' \rangle \rightarrow \tau' \wedge \mathcal{M}' \ominus_\lambda \mathcal{M} \neq \emptyset \end{cases}$$

Thus, principal p sees those sequences $\tau|_\lambda$ of observations where $\mathcal{R}(\lambda) \sqsubseteq_R \mathcal{R}(p)$ holds.

Piecewise Noninterference. In the absence of reclassifications, an *illicit λ -flow* is present if executing a command in states having different values of a variable with RIF label λ' satisfying $\lambda' \not\sqsubseteq_\Lambda \lambda$ will result in differences in updates to any variable with RIF label λ or less restrictive. Termination-insensitive noninterference [48] prohibits illicit λ -flows for any join semilattice. It is the starting point for noninterference with RIF labels.

To formalize an extension to termination-insensitive noninterference for accommodating reclassification under our threat model, we require some notation. Given τ a non-empty finite trace or subtrace,

$$\langle \mathcal{C}_1, \mathcal{M}_1 \rangle \rightarrow \langle \mathcal{C}_2, \mathcal{M}_2 \rangle \rightarrow \cdots \rightarrow \langle \mathcal{C}_N, \mathcal{M}_N \rangle$$

and indices i and j satisfying $1 \leq i \leq j \leq N$, define

$$\begin{aligned} \tau[i..] &\triangleq \langle \mathcal{C}_i, \mathcal{M}_i \rangle \rightarrow \cdots \rightarrow \langle \mathcal{C}_N, \mathcal{M}_N \rangle \\ \tau[i..j] &\triangleq \langle \mathcal{C}_i, \mathcal{M}_i \rangle \rightarrow \cdots \rightarrow \langle \mathcal{C}_j, \mathcal{M}_j \rangle \\ \tau[i] &\triangleq \langle \mathcal{C}_i, \mathcal{M}_i \rangle \\ \tau[i].\mathcal{C} &\triangleq \mathcal{C}_i \\ \tau[i].\mathcal{M} &\triangleq \mathcal{M}_i \\ |\tau| &\triangleq N \end{aligned}$$

In addition, we write $\tau.\mathcal{C}$ as an abbreviation for $\tau[1].\mathcal{C}$ and write $\tau.\mathcal{M}$ for $\tau[1].\mathcal{M}$. It is convenient to define $\tau[i..]$, $\tau[i..j]$, $\tau[i]$, $\tau[i].\mathcal{C}$, and $\tau[i].\mathcal{M}$ as equivalent to ϵ when $1 \leq i \leq j \leq N$ does not hold.

We start by formalizing prohibition of illicit λ -flows for commands $\widehat{\mathcal{C}}$ that do not contain reclassifying expressions. By definition, an illicit λ -flow is not present if executing $\widehat{\mathcal{C}}$ in memories \mathcal{M} and \mathcal{M}' satisfying $\mathcal{M} =_\lambda \mathcal{M}'$ does not result in differences in updates to a variable that has a RIF label λ or less restrictive.

$$(\forall \lambda, \mathcal{M}, \mathcal{M}', \tau, \tau': \tau = \Upsilon(\widehat{\mathcal{C}}, \mathcal{M}) \wedge \tau' = \Upsilon(\widehat{\mathcal{C}}, \mathcal{M}') \Rightarrow NI(\lambda, \tau, \tau')) \quad (11)$$

where

$$NI(\lambda, \tau, \tau') \triangleq \tau.\mathcal{M} =_\lambda \tau'.\mathcal{M} \Rightarrow \tau =_\lambda \tau'.$$

When $NI(\lambda, \tau, \tau')$ is *false* then τ and τ' constitute evidence of the illicit λ -flow.

Reclassification. Each transition $\langle \mathcal{C}_i, \mathcal{M}_i \rangle \rightarrow \langle \mathcal{C}_{i+1}, \mathcal{M}_{i+1} \rangle$ in a trace involves evaluating some (possibly empty) set of expressions. The operational semantics of commands determines what those expressions are and how their values are computed. We posit that the operational semantics for each command \mathcal{C}_i also includes a function $\Delta(\mathcal{C}_i)$ that gives the set of reclassifying expressions that a command \mathcal{C}_i evaluates as part of its first transition.

For example, with a simple imperative programming language and ordinary expressions \mathcal{E} and \mathcal{E}' we might expect to have

$$\Delta(x := \mathcal{E}) = \emptyset \quad (12)$$

$$\Delta(x := [\mathcal{E}]_{\bar{f}}) = \{[\mathcal{E}]_{\bar{f}}\} \quad (13)$$

$$\Delta(\mathbf{if} [\mathcal{E}]_{\bar{f}} \mathbf{then} x := [\mathcal{E}']_{\bar{f}}) = \{[\mathcal{E}]_{\bar{f}}\} \quad (14)$$

if we are assuming that assignments are executed as a single transition and that an **if** statement involves a first transition to evaluate its Boolean guard followed by other transitions for the **then** (or **else**) parts.

Handling Downgrades. A reclassifying expression $[\mathcal{E}]_{\bar{f}}$ is considered to perform a λ -downgrade if $\Gamma([\mathcal{E}]_{\bar{f}}) \sqsubseteq_{\Lambda} \lambda$ and $\Gamma(\mathcal{E}) \not\sqsubseteq_{\Lambda} \lambda$ hold. By definition, assigning the value of $[\mathcal{E}]_{\bar{f}}$ to a variable having RIF label λ does not constitute an illicit λ -flow (whereas assigning the value of \mathcal{E} would). So a λ -downgrade eliminates an illicit λ -flow by fiat.

A programming language that supports RIF labels will provide syntax that allows programmers to specify λ -downgrades by using reclassifying expressions. The set of expressions that satisfy the definition of a λ -downgrade and are evaluated by \mathcal{C}_i to make its transition is given by:

$$\{\mathcal{E} \mid [\mathcal{E}]_{\bar{f}} \in \Delta(\mathcal{C}_i) \wedge \Gamma([\mathcal{E}]_{\bar{f}}) \sqsubseteq_{\Lambda} \lambda \wedge \Gamma(\mathcal{E}) \not\sqsubseteq_{\Lambda} \lambda\}$$

To define our security policy for such programs $\widehat{\mathcal{C}}$, we posit that the language semantics includes a function⁷ $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \mathcal{C}_i)$ that satisfies:

$$\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \mathcal{C}_i) \subseteq \{\mathcal{E} \mid [\mathcal{E}]_{\bar{f}} \in \Delta(\mathcal{C}_i) \wedge \Gamma([\mathcal{E}]_{\bar{f}}) \sqsubseteq_{\Lambda} \lambda \wedge \Gamma(\mathcal{E}) \not\sqsubseteq_{\Lambda} \lambda\} \quad (15)$$

⁷The $\widehat{\mathcal{C}}$ subscript in $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \mathcal{C}_i)$ enables this function to depend on an enclosing program $\widehat{\mathcal{C}}$ and/or the position of \mathcal{C}_i within $\widehat{\mathcal{C}}$. For example, a language designer might elect to omit \mathcal{E} from $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, x := [\mathcal{E}]_{\bar{f}})$ if that assignment appears in the scope of an **if** having a Boolean guard \mathcal{E}' , where $\Gamma(\mathcal{E}') \not\sqsubseteq_{\Lambda} \Gamma(x)$ holds.

Trade-offs associated with the exact definition of $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \mathcal{C}_i)$ are discussed at the end of this section.

A generalization of (11) accommodates λ -downgrades that occur in the first transition of a trace that results from executing command $\widehat{\mathcal{C}}$. $NI(\lambda, \tau, \tau')$ was defined in terms of updates to any variable x satisfying $\Gamma(x) \sqsubseteq_{\Lambda} \lambda$. But differences in updates to x that arise when expressions in $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \widehat{\mathcal{C}})$ have different values are, by definition, not illicit λ -flows. Therefore, initial memory pairs \mathcal{M} and \mathcal{M}' that cause expressions in $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \widehat{\mathcal{C}})$ to have different values should be ignored in checking for indistinguishable observations [47]:

$$(\forall \lambda, \mathcal{M}, \mathcal{M}', \tau, \tau': \tau = \Upsilon(\widehat{\mathcal{C}}, \mathcal{M}) \wedge \tau' = \Upsilon(\widehat{\mathcal{C}}, \mathcal{M}') \Rightarrow dNI(\lambda, \tau, \tau')) \quad (16)$$

where

$$dNI(\lambda, \tau, \tau') \triangleq (\forall \mathcal{E} \in \Delta_{\widehat{\mathcal{C}}}^-(\lambda, \tau.\mathcal{C}): \tau.\mathcal{M}(\mathcal{E}) = \tau'.\mathcal{M}(\mathcal{E})) \Rightarrow NI(\lambda, \tau, \tau')$$

By construction $\tau.\mathcal{C} = \widehat{\mathcal{C}}$ and $\tau'.\mathcal{C} = \widehat{\mathcal{C}}$ hold, therefore $\tau.\mathcal{C} = \tau'.\mathcal{C}$ and

$$\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \tau.\mathcal{C}) = \Delta_{\widehat{\mathcal{C}}}^-(\lambda, \tau'.\mathcal{C}) \quad (17)$$

hold too. Thus, ignoring downgraded expressions in the first transition of trace τ (i.e., elements of $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \tau.\mathcal{C})$), $dNI(\lambda, \tau, \tau')$ is also ignoring downgraded expressions in the first transition of trace τ' (i.e., elements of $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \tau'.\mathcal{C})$).

To handle traces that perform downgrades after the first transition in $\widehat{\mathcal{C}}$, it suffices to note that each illicit λ -flow will be visible in some pair of λ -pieces, where a λ -piece in a trace τ is defined as a maximal length subtrace of τ having all of its λ -downgrades (if there are any) in its first transition. A pair of λ -pieces constitute a witness to an illicit λ -flow if they execute the same command in different memories \mathcal{M} and \mathcal{M}' satisfying $\mathcal{M} =_{\lambda} \mathcal{M}'$, and that execution leads to the pair of λ -pieces having (i) different commands⁸ in their last states and/or (ii) different updates to a variable x where $\Gamma(x) \sqsubseteq_{\Lambda} \lambda$ holds.

Given a trace τ and a RIF label λ , we introduce operators $\overset{\rightarrow\lambda}{\tau}$ and $\overset{\lambda\rightarrow}{\tau}$ to identify subtrace $\overset{\rightarrow\lambda}{\tau}$ that is the initial λ -piece of τ and subtrace $\overset{\lambda\rightarrow}{\tau}$ that is the rest of τ . According to the definition of a λ -piece, the last state of $\overset{\rightarrow\lambda}{\tau}$ will

⁸So the pair of λ -pieces do not necessarily produce different observations. The reason this conjunct is needed is seen below with example program (25).

also be the first state of $\overset{\lambda \rightarrow}{\tau}$ (if $\overset{\lambda \rightarrow}{\tau}$ is non-empty). Thus, $\overset{\rightarrow \lambda}{\tau}$ and $\overset{\lambda \rightarrow}{\tau}$ satisfy the following two properties, for traces and subtraces τ satisfying $2 \leq |\tau|$:

$$\begin{aligned} & (\overset{\rightarrow \lambda}{\tau} = \tau \wedge \overset{\lambda \rightarrow}{\tau} = \epsilon) \\ & \vee (\exists 1 < i < |\tau|: \overset{\rightarrow \lambda}{\tau} = \tau[1..i] \wedge \overset{\lambda \rightarrow}{\tau} = \tau[i..] \wedge \Delta_{\mathcal{C}}^-(\lambda, \tau[i].\mathcal{C}) \neq \emptyset) \end{aligned} \quad (18)$$

$$(\forall 1 < i < |\overset{\rightarrow \lambda}{\tau}|: \Delta_{\mathcal{C}}^-(\lambda, \overset{\rightarrow \lambda}{\tau}[i].\mathcal{C}) = \emptyset) \quad (19)$$

Two λ -pieces $\overset{\rightarrow \lambda}{\tau}$ and $\overset{\rightarrow \lambda}{\tau}'$ could exhibit an illicit λ -flow if the following is not *true*, since that would imply $\overset{\rightarrow \lambda}{\tau}$ has a different⁹ final command than $\overset{\rightarrow \lambda}{\tau}'$ or $\overset{\rightarrow \lambda}{\tau}$ involves different updates to a given variable than $\overset{\rightarrow \lambda}{\tau}'$.

$$\begin{aligned} & (\overset{\rightarrow \lambda}{\tau}.\mathcal{C} = \overset{\rightarrow \lambda}{\tau}'.\mathcal{C} \wedge \overset{\rightarrow \lambda}{\tau}.\mathcal{M} =_{\lambda} \overset{\rightarrow \lambda}{\tau}'.\mathcal{M}) \\ & \wedge (\forall \mathcal{E} \in \Delta_{\mathcal{C}}^-(\lambda, \tau.\mathcal{C}): \tau.\mathcal{M}(\mathcal{E}) = \tau'.\mathcal{M}(\mathcal{E})) \\ \Rightarrow & (\tau[|\overset{\rightarrow \lambda}{\tau}|].\mathcal{C} = \tau'[|\overset{\rightarrow \lambda}{\tau}'|].\mathcal{C} \wedge \overset{\rightarrow \lambda}{\tau} =_{\lambda} \overset{\rightarrow \lambda}{\tau}') \end{aligned} \quad (20)$$

By iterating (recursively) through corresponding λ -pieces in τ and τ' , the following formula uses the approach embodied by (20) to identify traces τ and τ' that evidence an illicit λ -flow.

$$\begin{aligned} dpNI(\lambda, \tau, \tau') & \triangleq (\tau \neq \epsilon \wedge \tau' \neq \epsilon \wedge \overset{\rightarrow \lambda}{\tau}.\mathcal{C} = \overset{\rightarrow \lambda}{\tau}'.\mathcal{C} \wedge \overset{\rightarrow \lambda}{\tau}.\mathcal{M} =_{\lambda} \overset{\rightarrow \lambda}{\tau}'.\mathcal{M}) \\ & \wedge (\forall \mathcal{E} \in \Delta_{\mathcal{C}}^-(\lambda, \tau.\mathcal{C}): \tau.\mathcal{M}(\mathcal{E}) = \tau'.\mathcal{M}(\mathcal{E})) \\ \Rightarrow & (\tau[|\overset{\rightarrow \lambda}{\tau}|].\mathcal{C} = \tau'[|\overset{\rightarrow \lambda}{\tau}'|].\mathcal{C} \wedge \overset{\rightarrow \lambda}{\tau} =_{\lambda} \overset{\rightarrow \lambda}{\tau}') \\ & \wedge dpNI(\lambda, \overset{\lambda \rightarrow}{\tau}, \overset{\lambda \rightarrow}{\tau}') \end{aligned} \quad (21)$$

Notice, for traces τ and τ' that each are a single λ -piece (i.e., $\tau = \overset{\rightarrow \lambda}{\tau}$ and $\tau' = \overset{\rightarrow \lambda}{\tau}'$ for all λ) then $dpNI(\lambda, \tau, \tau')$ is equivalent to (20), since $dpNI(\lambda, \overset{\lambda \rightarrow}{\tau}, \overset{\lambda \rightarrow}{\tau}')$ in the consequent of (21) would be $dpNI(\lambda, \epsilon, \epsilon)$, which is trivially satisfied.

A characterization similar to (16) now handles downgrades that appear anywhere in traces that result from executing a command \mathcal{C} .

$$(\forall \lambda, \mathcal{M}, \mathcal{M}', \tau, \tau': \tau = \Upsilon(\mathcal{C}, \mathcal{M}) \wedge \tau' = \Upsilon(\mathcal{C}, \mathcal{M}') \Rightarrow dpNI(\lambda, \tau, \tau')) \quad (22)$$

⁹Example (25) below illustrates what can go wrong.

Some example programs illustrate nuances of (22). These programs use RIF labels $\Lambda_{\mathbf{L}\mathbf{H}} \triangleq \{\mathbf{L}, \mathbf{H}\}$ with $\mathbf{L} \sqsubset_{\Lambda} \mathbf{H}$ and reclassifiers $\mathcal{F} \triangleq \{\downarrow, \uparrow\}$ satisfying $\mathcal{T}(\mathbf{H}, \downarrow) \triangleq \mathbf{L}$ and $\mathcal{T}(\mathbf{L}, \uparrow) \triangleq \mathbf{H}$. Assume that $\Gamma(\mathit{low}) = \Gamma(\mathit{low}') = \mathbf{L}$ and $\Gamma(\mathit{high}) = \Gamma(\mathit{high}') = \mathbf{H}$.

The first program assigns (in \mathcal{C}_3) a value with RIF label \mathbf{H} to a variable with RIF label \mathbf{L} without use of a reclassifying expression and, thus, would seem to exhibit an illicit \mathbf{L} -flow (and indeed this program has a leak according to the memory-reset approach proposed in [11] for characterizing declassification).

$$\begin{aligned} \mathcal{C}_1: \quad & \mathit{low} := [\mathit{op}(\mathit{high})]_{\downarrow}; \\ \mathcal{C}_2: \quad & \mathit{low}' := [\mathit{op}(\mathit{high}')]_{\downarrow}; \\ \mathcal{C}_3: \quad & \mathit{low} := \mathit{op}(\mathit{high}) \end{aligned} \tag{23}$$

Traces for program (23) comprise two \mathbf{L} -pieces.¹⁰ One \mathbf{L} -piece starts with command \mathcal{C}_1 , and the other \mathbf{L} -piece starts with command \mathcal{C}_2 (and includes \mathcal{C}_3), due to the following.

$$\Delta_{(23)}^-(\mathbf{L}, \mathcal{C}_1) = \{\mathit{op}(\mathit{high})\} \quad \Delta_{(23)}^-(\mathbf{L}, \mathcal{C}_2) = \{\mathit{op}(\mathit{high}')\} \quad \Delta_{(23)}^-(\mathbf{L}, \mathcal{C}_3) = \emptyset.$$

Checking, we find (22) is satisfied despite our earlier premonition about \mathcal{C}_3 . A close look shows why the flow to low in assignment \mathcal{C}_3 actually ought to be allowed, as characterization (22) does: \mathcal{C}_3 is assigning an \mathbf{L} -downgraded value, since the value of $\mathit{op}(\mathit{high})$ in the right hand side of \mathcal{C}_3 was an \mathbf{L} -downgraded value in \mathcal{C}_1 and has not been changed since. Gradual release [3], which does not rely on memory-resetting, also handles program (23) correctly.

A second program changes (in \mathcal{C}_2) the value in high after the \mathbf{L} -downgrade in \mathcal{C}_1 .

$$\begin{aligned} \mathcal{C}_1: \quad & \mathit{low} := [\mathit{op}(\mathit{high})]_{\downarrow}; \\ \mathcal{C}_2: \quad & \mathit{high} := \mathit{high}'; \\ \mathcal{C}_3: \quad & \mathit{low} := \mathit{op}(\mathit{high}) \end{aligned} \tag{24}$$

Traces for (24) comprise a single \mathbf{L} -piece that starts with command \mathcal{C}_1 because:

$$\Delta_{(24)}^-(\mathbf{L}, \mathcal{C}_1) = \{\mathit{op}(\mathit{high})\} \quad \Delta_{(24)}^-(\mathbf{L}, \mathcal{C}_2) = \emptyset \quad \Delta_{(24)}^-(\mathbf{L}, \mathcal{C}_3) = \emptyset$$

¹⁰Each trace is also a single \mathbf{H} -piece. For that case, $dpNI(\mathbf{H}, \tau, \tau')$ holds because $\tau.\mathcal{M} =_{\mathbf{H}} \tau'.\mathcal{M}$ is equivalent to $\tau.\mathcal{M} = \tau'.\mathcal{M}$. The same applies to example programs (24) and (25) to come.

Program (24) does not satisfy (22), since traces τ and τ' exist that generate observations that are not (but should be) indistinguishable. This is because there exist memories \mathcal{M} and \mathcal{M}' satisfying $\mathcal{M} =_{\text{L}} \mathcal{M}'$ and $\mathcal{M}(\text{high}') \neq \mathcal{M}'(\text{high}')$. When alternative executions of \mathcal{C}_1 are started in these two memories, \mathcal{C}_3 generates different updates to *low*. But having (22) not satisfied for this program is what we should desire—the value in *high* when \mathcal{C}_3 executes is not the value that was L-downgraded, so in program (24) a value with RIF label H that has not been L-downgraded is being used to update a variable with RIF label L.

A final program illustrates the role of conjunct $\tau[|\overset{\rightarrow\lambda}{\tau}|].\mathcal{C} = \tau[|\overset{\rightarrow\lambda}{\tau'}|].\mathcal{C}$ in the consequent of $dpNI(\lambda, \tau, \tau')$.

$$\begin{aligned}
\mathcal{C}_1: & \text{ if } \text{high}' > 0 \text{ then } \mathcal{C}_2: \text{ high} := 3 \\
& \qquad \qquad \mathcal{C}_3: \text{ low} := [\text{high} + 2]_{\downarrow} \\
& \text{ else } \mathcal{C}_4: \text{ high} := 4 \\
& \qquad \qquad \mathcal{C}_5: \text{ low} := [\text{high} \bmod 2]_{\downarrow}
\end{aligned} \tag{25}$$

Consider traces $\tau = \Upsilon(\mathcal{C}_1, \mathcal{M})$ and $\tau' = \Upsilon(\mathcal{C}_1, \mathcal{M}')$, where the following hold: $\mathcal{M} =_{\text{L}} \mathcal{M}'$, $\mathcal{M}(\text{high}' > 0) = \text{true}$, and $\mathcal{M}'(\text{high}' > 0) = \text{false}$. Trace τ comprises a first L-piece that starts with \mathcal{C}_1 (and includes \mathcal{C}_2), followed by a second L-piece that starts with \mathcal{C}_3 ; trace τ' has a first L-piece that starts with \mathcal{C}_1 (and includes \mathcal{C}_4), followed by a second L-piece that starts with \mathcal{C}_5 .

Program (25) does not satisfy (22) because both $\tau[|\overset{\rightarrow\lambda}{\tau}|].\mathcal{C} = \mathcal{C}_3$ and $\tau[|\overset{\rightarrow\lambda}{\tau'}|].\mathcal{C} = \mathcal{C}_5$ hold, so conjunct $\tau[|\overset{\rightarrow\lambda}{\tau}|].\mathcal{C} = \tau[|\overset{\rightarrow\lambda}{\tau'}|].\mathcal{C}$ in the consequent of $dpNI(\text{L}, \tau, \tau')$ does not hold. Having (22) not be satisfied is what we should desire, because the choice of \mathcal{C}_3 or \mathcal{C}_5 leaks information about the value of *high'* to *low*.

The test we are using for illicit λ -flows involving the program counter, however, can be too conservative. Consider the program obtained when \mathcal{C}_5 in (25) is replaced by assignment $\text{low} := [\text{high} + 1]_{\downarrow}$, which, by design, causes the same sequence of observations to be produced whether the program executes \mathcal{C}_3 or \mathcal{C}_5 . Conjunct $\tau[|\overset{\rightarrow\lambda}{\tau}|].\mathcal{C} = \tau[|\overset{\rightarrow\lambda}{\tau'}|].\mathcal{C}$ in the consequent of $dpNI(\text{L}, \tau, \tau')$ still does not hold, because \mathcal{C}_3 and \mathcal{C}_5 are different statements in the program. Yet the updates produced with the replaced \mathcal{C}_5 now are the same as \mathcal{C}_3 produces, so there is not actually an illicit L-flow.

Handling Upgrades. A reclassifying expression $[\mathcal{E}]_{\bar{f}}$ is considered to perform a λ -upgrade if $\Gamma(\mathcal{E}) \sqsubseteq_{\Lambda} \lambda$ and $\Gamma([\mathcal{E}]_{\bar{f}}) \not\sqsubseteq_{\Lambda} \lambda$ hold. Characterization (22) does not detect illicit λ -flows caused by λ -upgrades. Consider, for example, a program comprising assignment

$$low := [low']_{\uparrow} \quad (26)$$

that uses RIF labels from Λ_{LH} . Since $\Gamma(low) = \mathbf{L}$ and $\Gamma(low') = \mathbf{L}$ hold, this program satisfies (22). Yet the program exhibits an illicit \mathbf{L} -flow: differences in the initial value of an \mathbf{L} -upgraded expression ($[low']_{\uparrow}$) with RIF label \mathbf{H} result in differences in updates to a variable (low) with RIF label \mathbf{L} where $\Gamma([low']_{\uparrow}) \not\sqsubseteq_{\Lambda} \Gamma(low)$ holds.

In checking for evidence of illicit λ -flows, (22) compares traces τ and τ' that differ in initial values of variables whose RIF labels λ' satisfy $\lambda' \not\sqsubseteq_{\Lambda} \lambda$. That set of comparisons covers some expressions whose RIF labels λ' satisfy $\lambda' \not\sqsubseteq_{\Lambda} \lambda$ but it does not cover all expressions—it ignores expressions that involve λ -upgrades.

A programming language that supports RIF labels will provide syntax that allows programmers to specify λ -upgrades by using reclassifying expressions. We posit that the language semantics includes a function $\Delta_{\hat{\mathcal{C}}}^+(\lambda, \mathcal{C}_i)$ that satisfies:

$$\Delta_{\hat{\mathcal{C}}}^+(\lambda, \mathcal{C}_i) \supseteq \{ \mathcal{E} \mid [\mathcal{E}]_{\bar{f}} \in \Delta(\mathcal{C}_i) \wedge \Gamma([\mathcal{E}]_{\bar{f}}) \not\sqsubseteq_{\Lambda} \lambda \wedge \Gamma(\mathcal{E}) \sqsubseteq_{\Lambda} \lambda \} \quad (27)$$

So $\Delta_{\hat{\mathcal{C}}}^+(\lambda, \mathcal{C}_i)$ contains all expressions that satisfy the definition of a λ -upgrade and are evaluated by \mathcal{C}_i (within $\hat{\mathcal{C}}$) to make its transition.

Since (22) correctly identifies illicit λ -flows from variables x for which $\Gamma(x) \not\sqsubseteq_{\Lambda} \lambda$ holds, we transform λ -upgraded expressions to such variables as a way to identify illicit λ -flows for λ -upgraded expressions. Define translation $\mathbb{T}(\lambda, \hat{\mathcal{C}})$ to be the command that results from, in every subcommand \mathcal{C}_i of $\hat{\mathcal{C}}$, substituting an expression $next(h_{\mathcal{E}})$ for every expression $[\mathcal{E}]_{\bar{f}} \in \Delta(\mathcal{C}_i)$ where $\mathcal{E} \in \Delta_{\hat{\mathcal{C}}}^+(\lambda, \mathcal{C}_i)$ and $\Gamma(next(h_{\mathcal{E}})) = \Gamma([\mathcal{E}]_{\bar{f}})$ hold.¹¹ Here, $h_{\mathcal{E}}$ is a fresh list of values that \mathcal{E} could assume, and successive evaluations of $next(h_{\mathcal{E}})$ return successive elements of that list. If a program $\hat{\mathcal{C}}$ exhibits an illicit λ -flow from a λ -upgraded expression \mathcal{E} then, by construction, $\mathbb{T}(\lambda, \hat{\mathcal{C}})$ exhibits an illicit λ -flow from $next(h_{\mathcal{E}})$. Moreover, because $\mathbb{T}(\lambda, \hat{\mathcal{C}})$ contains no λ -upgraded expressions, it exhibits no illicit λ -flows from upgraded λ -expressions *per se*.

¹¹See (29) and (30) below for an example of the translation.

That means we can use (22) to check for illicit λ -flows in $\widehat{\mathcal{C}}$ by checking $\mathbb{T}(\lambda, \widehat{\mathcal{C}})$ for illicit λ -flows.

The resulting characterization of *piecewise noninterference* (PWNI) for a program $\widehat{\mathcal{C}}$ employs $dpNI(\lambda, \tau, \tau')$ —defined in (21)—by extending (22) to handle λ -flows from λ -upgraded expressions.

$$\begin{aligned} PWNI(\widehat{\mathcal{C}}) &\triangleq \\ (\forall \lambda, \mathcal{M}, \mathcal{M}', \tau, \tau': \quad &\tau = \Upsilon(\mathbb{T}(\lambda, \widehat{\mathcal{C}}), \mathcal{M}) \wedge \tau' = \Upsilon(\mathbb{T}(\lambda, \widehat{\mathcal{C}}), \mathcal{M}') \quad (28) \\ &\Rightarrow dpNI(\lambda, \tau, \tau')) \end{aligned}$$

The election example and the conference-management example of §1 both satisfy this definition when coded in a straightforward manner.

$PWNI(\widehat{\mathcal{C}})$ is conservative. Some of the conservatism results from considering executions of a translation in which L-upgraded expressions \mathcal{E} are replaced by different fresh variables $h_{\mathcal{E}}$. Equivalent expressions within an execution and in different executions of the original program may take different values in executions of the translation. Here is an illustration, where RIF labels are from Λ_{LH} .

$$\mathcal{C}_1: \text{ if } [low = low]_{\uparrow} \text{ then } \mathcal{C}_2: low := 1 \text{ else } \mathcal{C}_3: low := 2 \quad (29)$$

$PWNI(\mathcal{C}_1)$ is defined in terms of translated program $\mathbb{T}(\lambda, \mathcal{C}_1)$,

$$\mathcal{C}_1: \text{ if } next(h_{low=low}) \text{ then } \mathcal{C}_2: low := 1 \text{ else } \mathcal{C}_3: low := 2 \quad (30)$$

where $\Gamma(next(h_{low=low})) = \text{H}$ holds. Boolean guard $[low = low]_{\uparrow}$ in (29) is always *true*, so execution of (29) produces no traces in which \mathcal{C}_3 executes. In translated program (30), Boolean guard $next(h_{low=low})$ takes values *true* and *false*, so there are traces of (30) where \mathcal{C}_3 executes as well as traces where \mathcal{C}_2 executes. Thus, $PWNI(\mathcal{C}_1)$ is not satisfied for (30) due to the illicit L-flow from upgraded expression $[low = low]_{\uparrow}$ to low ; in program (29) the corresponding illicit L-flow from upgraded expression $[low = low]_{\uparrow}$ does not actually exist.

Conservatism in $PWNI(\widehat{\mathcal{C}})$ is also related to flexibility that (15) allows in the definition of $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \mathcal{C}_i)$ and that (27) allows in the definition of $\Delta_{\widehat{\mathcal{C}}}^+(\lambda, \mathcal{C}_i)$.

- $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \mathcal{C}_i)$ is constrained only by (15), so $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \mathcal{C}_i)$ may omit expressions that cause λ -downgrades. With fewer expressions in $\Delta_{\widehat{\mathcal{C}}}^-(\lambda, \mathcal{C}_i)$:

- There are more pairs of memories from which execution results in a pair of λ -pieces whose subtraces $\xrightarrow{\lambda} \tau$ and $\xrightarrow{\lambda} \tau'$ must satisfy (20). So the omissions may cause $PWNI(\widehat{\mathcal{C}})$ to reject programs that do not actually exhibit illicit λ -flows.
- Fewer pieces need to be checked by evaluating $dpNI(\lambda, \tau, \tau')$. So the omissions may cause $PWNI(\widehat{\mathcal{C}})$ to accept more programs that do not exhibit illicit λ -flows.
- $\Delta_{\widehat{\mathcal{C}}}^+(\lambda, \mathcal{C}_i)$ is constrained only by (27), so $\Delta_{\widehat{\mathcal{C}}}^+(\lambda, \mathcal{C}_i)$ may include expressions that do not cause λ -upgrades. The additional expressions \mathcal{E} are replaced in each \mathcal{C}_i by expressions involving list $h_{\mathcal{E}}$. That means there could be more pairs of memories (i.e., those that differ in $h_{\mathcal{E}}$) from which execution results in a pair of λ -pieces whose subtraces $\xrightarrow{\lambda} \tau$ and $\xrightarrow{\lambda} \tau'$ must satisfy (20). So the additions may cause $PWNI(\widehat{\mathcal{C}})$ to reject programs that do not exhibit illicit λ -flows.

Assessing the utility of a definition like PWNI can be a challenge. Absent a widely accepted definition of noninterference for systems that support reclassification—and there does not seem to be one—we are forced into studying examples if we want to build confidence about whether a proposed definition is suitable. The examples above give interesting edge-cases. Program (23) is particularly instructive, because it highlights a difference in how downgrading is treated by PWNI as compared with the proposal in [11]. Under PWNI, a value that is downgraded remains so for the remainder of the trace; with the proposal in [11] a value that is downgraded remains so only to the next reclassification operation.

4 Static Enforcement of PWNI

Type-checking allows static enforcement of $PWNI(\widehat{\mathcal{C}})$ when $\widehat{\mathcal{C}}$ is written in a programming language having a type system based on some checkable class of RIF labels. A simple imperative language provides a vehicle for demonstration.

Language and Semantics. Figure 2 gives a syntax of a simple programming language for defining commands. There, ν ranges over constants, x ranges over program variables, $\mathcal{E}, \mathcal{E}_1, \mathcal{E}_2, \dots$ range over ordinary expressions,

$$\begin{array}{l}
\mathcal{E} ::= \nu \mid x \mid op(\mathcal{E}_1, \dots, \mathcal{E}_n) \\
\mathcal{C} ::= \mathbf{skip} \\
\quad \mid x := [\mathcal{E}]_{\bar{f}} \\
\quad \mid \mathbf{if} [\mathcal{E}]_{\bar{f}} \mathbf{then} \mathcal{C}_t \mathbf{else} \mathcal{C}_e \\
\quad \mid \mathbf{while} [\mathcal{E}]_{\bar{f}} \mathbf{do} \mathcal{C}_t \\
\quad \mid \mathcal{C}_1; \mathcal{C}_2
\end{array}$$

Figure 2: Syntax of simple imperative language

and $\mathcal{C}_t, \mathcal{C}_e, \mathcal{C}_1, \mathcal{C}_2, \dots$ range over commands. Note, allowing only reclassifying expressions (rather than ordinary expressions) in the language syntax for commands is not a limitation—due to (4), identity reclassifier ϵ is handled by every RIF label and reclassifying expression $[\mathcal{E}]_{\epsilon}$ has the same value and RIF label as ordinary expression \mathcal{E} .

Figure 3 gives an operational semantics for the programming language of Figure 2. We write $\mathcal{M}[x \mapsto \nu]$ there to denote a mapping that is identical to \mathcal{M} except $\mathcal{M}(x) = \nu$. The rules in Figure 3 define partial function $\Upsilon(\mathcal{C}, \mathcal{M})$.

The final part of the semantics for this programming language is definitions for $\Delta(\mathcal{C}_i)$, $\Delta_{\bar{c}}^-(\lambda, \mathcal{C}_i)$, and $\Delta_{\bar{c}}^+(\lambda, \mathcal{C}_i)$. These are given in Figure 4 through Figure 6. In defining $\Delta_{\bar{c}}^-(\lambda, \mathcal{C}_i)$, we write $lhs(\mathcal{C}_i)$ to denote the set of target variables in assignments appearing in a command \mathcal{C}_i .

- $\Delta_{\bar{c}}^-(\lambda, \mathcal{C}_i)$ excludes λ -downgrades that cannot influence the value being assigned to a variable x where $\Gamma(x) \sqsubseteq_{\Lambda} \lambda$ holds. No illicit λ -flow is possible without such an assignment.
- $\Delta_{\bar{c}}^+(\lambda, \mathcal{C}_i)$ for **if** and **while** is the smallest set that satisfies (27). To exclude λ -upgrade guards because an assignment does not appear in the body of an **if** or **while** (as done in the definition of $\Delta_{\bar{c}}^+(\lambda, \mathcal{C}_i)$) would invalidate (27) by decreasing the size of $\Delta_{\bar{c}}^+(\lambda, \mathcal{C}_i)$.

Typing Rules. Figure 7 gives rules to associate a type with each expression. The rules for ordinary expressions are standard. EXPR-T instantiates (9); ANNEXPRT is based on (10).

Figure 8 shows the familiar rules to deduce whether a command is type-correct. Judgment $\Gamma, \lambda_{\kappa} \vdash \mathcal{C}$ signifies that a command \mathcal{C} is type correct. Parameter λ_{κ} in these rules is called *pc-label*. It is used in checking for

$$\begin{array}{c}
\text{SKIP:} \frac{}{\langle \mathbf{skip}, \mathcal{M} \rangle \rightarrow \langle \bullet, \mathcal{M} \rangle} \\
\\
\text{ASGN:} \frac{\mathcal{M}([\mathcal{E}]_{\bar{f}}) = \nu}{\langle x := [\mathcal{E}]_{\bar{f}}, \mathcal{M} \rangle \rightarrow \langle \bullet, \mathcal{M}[x \mapsto \nu] \rangle} \\
\\
\text{BRCH1:} \frac{\mathcal{M}([\mathcal{E}]_{\bar{f}}) = \text{true}}{\langle \mathbf{if} [\mathcal{E}]_{\bar{f}} \mathbf{then} \mathcal{C}_t \mathbf{else} \mathcal{C}_e, \mathcal{M} \rangle \rightarrow \langle \mathcal{C}_t, \mathcal{M} \rangle} \\
\\
\text{BRCH2:} \frac{\mathcal{M}([\mathcal{E}]_{\bar{f}}) = \text{false}}{\langle \mathbf{if} [\mathcal{E}]_{\bar{f}} \mathbf{then} \mathcal{C}_t \mathbf{else} \mathcal{C}_e, \mathcal{M} \rangle \rightarrow \langle \mathcal{C}_e, \mathcal{M} \rangle} \\
\\
\text{LOOP1:} \frac{\mathcal{M}([\mathcal{E}]_{\bar{f}}) = \text{true}}{\langle \mathbf{while} [\mathcal{E}]_{\bar{f}} \mathbf{do} \mathcal{C}_t, \mathcal{M} \rangle \rightarrow \langle \mathcal{C}_t; \mathbf{while} [\mathcal{E}]_{\bar{f}} \mathbf{do} \mathcal{C}_t, \mathcal{M} \rangle} \\
\\
\text{LOOP2:} \frac{\mathcal{M}([\mathcal{E}]_{\bar{f}}) = \text{false}}{\langle \mathbf{while} [\mathcal{E}]_{\bar{f}} \mathbf{do} \mathcal{C}_t, \mathcal{M} \rangle \rightarrow \langle \bullet, \mathcal{M} \rangle} \\
\\
\text{SEQ1:} \frac{\langle \mathcal{C}_1, \mathcal{M} \rangle \rightarrow \langle \bullet, \mathcal{M}_1 \rangle}{\langle \mathcal{C}_1; \mathcal{C}_2, \mathcal{M} \rangle \rightarrow \langle \mathcal{C}_2, \mathcal{M}_1 \rangle} \\
\\
\text{SEQ2:} \frac{\langle \mathcal{C}_1, \mathcal{M} \rangle \rightarrow \langle \mathcal{C}, \mathcal{M}_1 \rangle \quad \mathcal{C} \neq \bullet}{\langle \mathcal{C}_1; \mathcal{C}_2, \mathcal{M} \rangle \rightarrow \langle \mathcal{C}; \mathcal{C}_2, \mathcal{M}_1 \rangle}
\end{array}$$

Figure 3: Operational Semantics

$$\begin{aligned}
\Delta(\mathbf{skip}) &\triangleq \emptyset \\
\Delta(x := [\mathcal{E}]_{\bar{f}}) &\triangleq \{[\mathcal{E}]_{\bar{f}}\} \\
\Delta(\mathbf{if} [\mathcal{E}]_{\bar{f}} \mathbf{then} \mathcal{C}_t \mathbf{else} \mathcal{C}_e) &\triangleq \{[\mathcal{E}]_{\bar{f}}\} \\
\Delta(\mathbf{while} [\mathcal{E}]_{\bar{f}} \mathbf{do} \mathcal{C}_w) &\triangleq \{[\mathcal{E}]_{\bar{f}}\} \\
\Delta(\mathcal{C}_1; \mathcal{C}_2) &\triangleq \Delta(\mathcal{C}_1)
\end{aligned}$$

Figure 4: Definition of $\Delta(\mathcal{C}_i)$

$$\begin{aligned}
\Delta_{\hat{\mathcal{C}}}^-(\lambda, \mathbf{skip}) &\triangleq \emptyset \\
\Delta_{\hat{\mathcal{C}}}^-(\lambda, x := [\mathcal{E}]_{\bar{f}}) &\triangleq \begin{cases} \{\mathcal{E}\} & \text{if } \Gamma(\mathcal{E}) \not\sqsubseteq_{\Lambda} \lambda \wedge \Gamma([\mathcal{E}]_{\bar{f}}) \sqsubseteq_{\Lambda} \lambda \wedge \Gamma(x) \sqsubseteq_{\Lambda} \lambda \\ \emptyset & \text{otherwise} \end{cases} \\
\Delta_{\hat{\mathcal{C}}}^-(\lambda, \mathbf{if } [\mathcal{E}]_{\bar{f}} \mathbf{then } \mathcal{C}_t \mathbf{else } \mathcal{C}_e) &\triangleq \begin{cases} \{\mathcal{E}\} & \text{if } \Gamma(\mathcal{E}) \not\sqsubseteq_{\Lambda} \lambda \wedge \Gamma([\mathcal{E}]_{\bar{f}}) \sqsubseteq_{\Lambda} \lambda \\ & \wedge (\exists x \in \text{lhs}(\mathcal{C}_t) \cup \text{lhs}(\mathcal{C}_e): \Gamma(x) \sqsubseteq_{\Lambda} \lambda) \\ \emptyset & \text{otherwise} \end{cases} \\
\Delta_{\hat{\mathcal{C}}}^-(\lambda, \mathbf{while } [\mathcal{E}]_{\bar{f}} \mathbf{do } \mathcal{C}_w) &\triangleq \begin{cases} \{\mathcal{E}\} & \text{if } \Gamma(\mathcal{E}) \not\sqsubseteq_{\Lambda} \lambda \wedge \Gamma([\mathcal{E}]_{\bar{f}}) \sqsubseteq_{\Lambda} \lambda \\ & \wedge (\exists x \in \text{lhs}(\mathcal{C}_w): \Gamma(x) \sqsubseteq_{\Lambda} \lambda) \\ \emptyset & \text{otherwise} \end{cases} \\
\Delta_{\hat{\mathcal{C}}}^-(\lambda, \mathcal{C}_1; \mathcal{C}_2) &\triangleq \Delta_{\hat{\mathcal{C}}}^-(\lambda, \mathcal{C}_1)
\end{aligned}$$

Figure 5: Definition of $\Delta_{\hat{\mathcal{C}}}^-(\lambda, \mathcal{C})$

$$\begin{aligned}
\Delta_{\hat{\mathcal{C}}}^+(\lambda, \mathbf{skip}) &\triangleq \emptyset \\
\Delta_{\hat{\mathcal{C}}}^+(\lambda, x := [\mathcal{E}]_{\bar{f}}) &\triangleq \begin{cases} \{\mathcal{E}\} & \text{if } \Gamma(\mathcal{E}) \sqsubseteq_{\Lambda} \lambda \wedge \Gamma([\mathcal{E}]_{\bar{f}}) \not\sqsubseteq_{\Lambda} \lambda \\ \emptyset & \text{otherwise} \end{cases} \\
\Delta_{\hat{\mathcal{C}}}^+(\lambda, \mathbf{if } [\mathcal{E}]_{\bar{f}} \mathbf{then } \mathcal{C}_t \mathbf{else } \mathcal{C}_e) &\triangleq \begin{cases} \{\mathcal{E}\} & \text{if } \Gamma(\mathcal{E}) \sqsubseteq_{\Lambda} \lambda \wedge \Gamma([\mathcal{E}]_{\bar{f}}) \not\sqsubseteq_{\Lambda} \lambda \\ \emptyset & \text{otherwise} \end{cases} \\
\Delta_{\hat{\mathcal{C}}}^+(\lambda, \mathbf{while } [\mathcal{E}]_{\bar{f}} \mathbf{do } \mathcal{C}_w) &\triangleq \begin{cases} \{\mathcal{E}\} & \text{if } \Gamma(\mathcal{E}) \sqsubseteq_{\Lambda} \lambda \wedge \Gamma([\mathcal{E}]_{\bar{f}}) \not\sqsubseteq_{\Lambda} \lambda \\ \emptyset & \text{otherwise} \end{cases} \\
\Delta_{\hat{\mathcal{C}}}^+(\lambda, \mathcal{C}_1; \mathcal{C}_2) &\triangleq \Delta_{\hat{\mathcal{C}}}^+(\lambda, \mathcal{C}_1)
\end{aligned}$$

Figure 6: Definition of $\Delta_{\hat{\mathcal{C}}}^+(\lambda, \mathcal{C})$

$$\begin{array}{c}
\text{VAL-T:} \frac{}{\Gamma \vdash \nu : \perp} \qquad \text{VAR-T:} \frac{\Gamma(x) = \lambda}{\Gamma \vdash x : \lambda} \\
\text{EXPR-T:} \frac{\Gamma \vdash \mathcal{E}_1 : \lambda_1 \quad \dots \quad \Gamma \vdash \mathcal{E}_n : \lambda_n}{\Gamma \vdash \text{op}(\mathcal{E}_1, \dots, \mathcal{E}_n) : \lambda_1 \sqcup_{\Lambda} \dots \sqcup_{\Lambda} \lambda_n} \\
\text{ANNEXPR-T:} \frac{\Gamma \vdash \mathcal{E}_1 : \lambda_1 \quad \dots \quad \Gamma \vdash \mathcal{E}_n : \lambda_n}{\Gamma \vdash [\text{op}(\mathcal{E}_1, \dots, \mathcal{E}_n)]_{f_1, \dots, f_n} : \mathcal{T}(\lambda_1, f_1) \sqcup_{\Lambda} \dots \sqcup_{\Lambda} \mathcal{T}(\lambda_n, f_n)}
\end{array}$$

Figure 7: Typing rules for expressions

$$\begin{array}{c}
\text{SKIP-T:} \frac{}{\Gamma, \lambda_{\kappa} \vdash \mathbf{skip}} \\
\text{ASGN-T:} \frac{\Gamma \vdash [\mathcal{E}]_{\bar{f}} : \lambda_g \quad \Gamma \vdash x : \lambda_x \quad \lambda_{\kappa} \sqcup_{\Lambda} \lambda_g \sqsubseteq_{\Lambda} \lambda_x}{\Gamma, \lambda_{\kappa} \vdash x := [\mathcal{E}]_{\bar{f}}} \\
\text{BRCH-T:} \frac{\Gamma \vdash [\mathcal{E}]_{\bar{f}} : \lambda_g \quad \Gamma, \lambda_{\kappa} \sqcup_{\Lambda} \lambda_g \vdash \mathcal{C}_t \quad \Gamma, \lambda_{\kappa} \sqcup_{\Lambda} \lambda_g \vdash \mathcal{C}_e}{\Gamma, \lambda_{\kappa} \vdash \mathbf{if} [\mathcal{E}]_{\bar{f}} \mathbf{then} \mathcal{C}_t \mathbf{else} \mathcal{C}_e} \\
\text{LOOP-T:} \frac{\Gamma \vdash [\mathcal{E}]_{\bar{f}} : \lambda_g \quad \Gamma, \lambda_{\kappa} \sqcup_{\Lambda} \lambda_g \vdash \mathcal{C}_t}{\Gamma, \lambda_{\kappa} \vdash \mathbf{while} [\mathcal{E}]_{\bar{f}} \mathbf{do} \mathcal{C}_t} \qquad \text{SEQ-T:} \frac{\Gamma, \lambda_{\kappa} \vdash \mathcal{C}_1 \quad \Gamma, \lambda_{\kappa} \vdash \mathcal{C}_2}{\Gamma, \lambda_{\kappa} \vdash \mathcal{C}_1; \mathcal{C}_2}
\end{array}$$

Figure 8: Typing rules for commands

implicit flows. A pc-label λ_{κ} associates a type with the conjunction of the guards for all of the enclosing conditional commands; λ_{κ} is combined with the type of the right hand side of an assignment statement.

The following theorem states that if a program is type correct, then this program satisfies PWNI.

Theorem 1. *If $\Gamma, \lambda_{\kappa} \vdash \mathcal{C}$, then $PWNI(\mathcal{C})$ holds.*

Proof. See [31, Appendix A.1] □

5 Some Checkable Classes of RIF Labels

5.1 RIF Automata

Finite state automata provide the basis for a checkable class of RIF labels, called *RIF automata*, that have broad practical utility. This class of RIF labels is supported in the JRIF¹² programming language [32]. We built JRIF to gain practical experience with using RIF automata for specifying security policies and to understand the compiler modifications needed when a programming language that supports ordinary security label types is converted to use RIF automata. The *privacy automata* used in the Avance language [8] for specifying use-based privacy are also instances of RIF automata.

Formalization of RIF Automata. A finite state automaton can serve as a RIF label λ_α by: (i) having the set of reclassifiers be the automaton’s input alphabet, and (ii) associating restrictions with each automaton state. Restrictions imposed by λ_α are those associated with the automaton’s current state. Reclassifiers, which by construction change the automaton’s current state, thus cause a (potentially) different set of restrictions to be imposed.

Formally, a set Λ_{RA} comprising RIF automata is defined relative to some join semilattice of restrictions $\langle R, \sqcup_R, \sqsubseteq_R \rangle$. Each RIF automaton $\lambda_\alpha \in \Lambda_{\text{RA}}$ is described by a 5-tuple

$$\lambda_\alpha \triangleq \langle Q_\alpha, \mathcal{F}, \delta_\alpha, q_\alpha, \rho_\alpha \rangle \quad (31)$$

where:

Q_α is a finite set of automaton states

\mathcal{F} is the finite set of reclassifiers

$\delta_\alpha: Q_\alpha \times \mathcal{F} \rightarrow Q_\alpha$ is a (deterministic) next-state transition function¹³

$q_\alpha \in Q_\alpha$ is the current state of the RIF automaton

$\rho_\alpha: Q_\alpha \rightarrow R$ gives the restrictions associated with each automaton state

¹²JRIF is a variant of the Jif security-typed programming language, which extends Java with support for information flow control.

¹³Closure $\delta_\alpha^*: Q_\alpha \times \mathcal{F}^* \rightarrow Q_\alpha$ is derived from δ_α in the usual way: $\delta_\alpha^*(q, \epsilon) \triangleq q$ and $\delta_\alpha^*(q, Ff) \triangleq \delta_\alpha(\delta_\alpha^*(q, F), f)$

We require transition function δ_α to be total, so any sequence of reclassifiers from \mathcal{F} causes a sequence of transitions. And for $F \in \mathcal{F}^*$ we write $\lambda_\alpha(F)$ to denote the automaton that results when λ_α performs the transitions indicated by a sequence F of reclassifiers. Automaton $\lambda_\alpha(F)$ thus replaces current state q_α of λ_α with $\delta_\alpha^*(q_\alpha, F)$; we define $\lambda_\alpha(F)$ formally by:

$$\begin{aligned} \lambda_\alpha(F) &\triangleq \lambda_\alpha[q_\alpha \mapsto \delta_\alpha^*(q_\alpha, F)] \\ \text{where for } q_i \in Q_\alpha: \lambda_\alpha[q_\alpha \mapsto q_i] &\triangleq \langle Q_\alpha, \mathcal{F}, \delta_\alpha, q_i, \rho_\alpha \rangle \end{aligned}$$

\mathcal{R}_{RA} and \mathcal{T}_{RA} are defined as expected.

$$\mathcal{R}_{\text{RA}}(\lambda_\alpha) \triangleq \rho_\alpha(q_\alpha) \quad \mathcal{T}_{\text{RA}}(\lambda_\alpha, f) \triangleq \lambda_\alpha(f)$$

Instantiating definition (7) of \sqsubseteq_Δ using these definitions gets:

$$\lambda_\alpha \sqsubseteq_{\text{RA}} \lambda_{\alpha'} \triangleq (\forall F \in \mathcal{F}^*: \rho_\alpha(\delta_\alpha^*(q_\alpha, F)) \sqsubseteq_R \rho_{\alpha'}(\delta_{\alpha'}^*(q_{\alpha'}, F))) \quad (32)$$

A computable test for deciding $\lambda_\alpha \sqsubseteq_{\text{RA}} \lambda_{\alpha'}$ comes if we construct a *restrictiveness product* automaton $\lambda_\alpha \otimes \lambda_{\alpha'}$ that is the product of RIF automata λ_α and $\lambda_{\alpha'}$ but with unreachable automaton states eliminated.

$$\begin{aligned} \lambda_\alpha \otimes \lambda_{\alpha'} &\triangleq \langle Q_{\alpha \otimes \alpha'}, \mathcal{F}, \delta_{\alpha \otimes \alpha'}, \langle q_\alpha, q_{\alpha'} \rangle, \rho_{\alpha \otimes \alpha'} \rangle \\ \text{where } Q_{\alpha \otimes \alpha'} &\subseteq Q_\alpha \times Q_{\alpha'} \\ \delta_{\alpha \otimes \alpha'}(\langle q, q' \rangle, f) &\triangleq \langle \delta_\alpha(q, f), \delta_{\alpha'}(q', f) \rangle \\ \rho_{\alpha \otimes \alpha'}(\langle q, q' \rangle) &\triangleq \langle \rho_\alpha(q), \rho_{\alpha'}(q') \rangle \\ (\forall q \in Q_{\alpha \otimes \alpha'}: (\exists F_q \in \mathcal{F}^*: \delta_{\alpha \otimes \alpha'}^*(\langle q_\alpha, q_{\alpha'} \rangle, F_q) = q)) &\quad (33) \end{aligned}$$

Note that the signature for $\rho_{\alpha \otimes \alpha'}$ in $\lambda_\alpha \otimes \lambda_{\alpha'}$

$$\rho_{\alpha \otimes \alpha'}: Q_{\alpha \otimes \alpha'} \rightarrow R \times R,$$

implies that restrictiveness product automata are not RIF automata; a RIF automaton must associate a single element of R (rather than a pair of elements) with each automaton state.

Condition (33), which stipulates that all states in $Q_{\alpha \otimes \alpha'}$ are reachable, is straightforward to check with a linear-time depth-first search of the state-transition graph for $\delta_{\alpha \otimes \alpha'}$. And condition (33) implies that for any predicate $\mathcal{P}(\langle q, q' \rangle)$ on states $\langle q, q' \rangle \in Q_{\alpha \otimes \alpha'}$

$$(\forall \langle q, q' \rangle \in Q_{\alpha \otimes \alpha'}: \mathcal{P}(\langle q, q' \rangle)) = (\forall F \in \mathcal{F}^*: \mathcal{P}(\delta_{\alpha \otimes \alpha'}^*(\langle q_\alpha, q_{\alpha'} \rangle, F))) \quad (34)$$

holds. It means that a decision procedure for $\lambda_\alpha \sqsubseteq_{\text{RA}} \lambda_{\alpha'}$ need only check a predicate on the pair of restrictions associated with each state $\langle q, q' \rangle \in Q_{\alpha \otimes \alpha'}$ —namely that $r \sqsubseteq_R r'$ holds if $\rho_{\alpha \otimes \alpha'}(\langle q, q' \rangle) = \langle r, r' \rangle$.¹⁴

$$\lambda_\alpha \sqsubseteq_{\text{RA}} \lambda_{\alpha'} = (\forall \langle q, q' \rangle \in Q_{\alpha \otimes \alpha'}: \rho_\alpha(q) \sqsubseteq_R \rho_{\alpha'}(q')) \quad (35)$$

The properties of \sqcup_{RA} for RIF automata are satisfied by using a form of product construction that does yield a RIF automaton.

$$\begin{aligned} \lambda_\alpha \sqcup_{\text{RA}} \lambda_{\alpha'} &\triangleq \langle Q_{\alpha \sqcup \alpha'}, \mathcal{F}, \delta_{\alpha \sqcup \alpha'}, \langle q_\alpha, q_{\alpha'} \rangle, \rho_{\alpha \sqcup \alpha'} \rangle \\ \text{where } Q_{\alpha \sqcup \alpha'} &= Q_\alpha \times Q_{\alpha'} \\ \delta_{\alpha \sqcup \alpha'}(\langle q, q' \rangle, f) &\triangleq \langle \delta_\alpha(q, f), \delta_{\alpha'}(q', f) \rangle \\ \rho_{\alpha \sqcup \alpha'}(\langle q, q' \rangle) &\triangleq \rho_\alpha(q) \sqcup_R \rho_{\alpha'}(q') \end{aligned}$$

Note that (6) relating \sqcup_R and \sqcup_{RA} is satisfied by this definition.

Finally, we prove that $\langle \Lambda_{\text{RA}}, \sqsubseteq_{\text{RA}}, \sqcup_{\text{RA}} \rangle$ is a join semilattice.¹⁵

¹⁴To prove that (35) is equivalent to definition (32) of $\lambda_\alpha \sqsubseteq_{\text{RA}} \lambda_{\alpha'}$, observe that

$$\begin{aligned} &(\forall \langle q, q' \rangle \in Q_{\alpha \otimes \alpha'}: \rho_\alpha(q) \sqsubseteq_R \rho_{\alpha'}(q')) \\ &= \text{due to (34), with } \mathcal{P}(\langle q, q' \rangle) \triangleq \rho_\alpha(q) \sqsubseteq_R \rho_{\alpha'}(q') \\ &(\forall F \in \mathcal{F}^*: \rho_\alpha(\delta_\alpha^*(q_\alpha, F)) \sqsubseteq_R \rho_{\alpha'}(\delta_{\alpha'}^*(q_{\alpha'}, F))) \end{aligned}$$

¹⁵To prove that $\langle \Lambda_{\text{RA}}, \sqsubseteq_{\text{RA}}, \sqcup_{\text{RA}} \rangle$ is a join semilattice requires showing (i) $\lambda_\alpha \sqsubseteq_{\text{RA}} \lambda_\alpha \sqcup_{\text{RA}} \lambda_{\alpha'}$, (ii) $\lambda_{\alpha'} \sqsubseteq_{\text{RA}} \lambda_\alpha \sqcup_{\text{RA}} \lambda_{\alpha'}$, and (iii) that $\lambda_\alpha \sqcup_{\text{RA}} \lambda_{\alpha'}$ is the least upper bound.

To prove (i), it suffices to use decision procedure (35) for \sqsubseteq_{RA} , and prove:

$$(\forall \langle q, \langle q', q'' \rangle \rangle \in Q_{\alpha \otimes (\alpha \sqcup \alpha')}: \rho_\alpha(q) \sqsubseteq_R \rho_{\alpha \sqcup \alpha'}(\langle q', q'' \rangle)), \quad (36)$$

because reachability condition (33) for $\lambda_\alpha \otimes (\lambda_\alpha \sqcup_{\text{RA}} \lambda_{\alpha \sqcup \alpha'})$ implies that if

$$\langle q, \langle q', q'' \rangle \rangle \in Q_{\alpha \otimes (\alpha \sqcup \alpha')}$$

holds then $q = q'$. Completing the proof of (36) is just a matter of expanding the definition of $\rho_{\alpha \sqcup \alpha'}(\langle q, q'' \rangle)$.

The proof of (ii) is similar.

The proof of (iii) involves showing that $\lambda_\alpha \sqcup_{\text{RA}} \lambda_{\alpha'}$ is the least RIF label satisfying

$$\begin{aligned} \lambda_\alpha &\sqsubseteq_{\text{RA}} \lambda_\alpha \sqcup_{\text{RA}} \lambda_{\alpha'} \\ \lambda_{\alpha'} &\sqsubseteq_{\text{RA}} \lambda_\alpha \sqcup_{\text{RA}} \lambda_{\alpha'} \end{aligned}$$




Figure 9: RIF automaton $\lambda_{voter(A)}$ for secret ballots

Examples of RIF Automata. A security policy that we might associate with the ballot that each participant A casts in an election is: (i) only A may read the ballot's value and (ii) anyone may read the majority value derived from all of the ballots cast. We formalize this security policy as a RIF automaton $\lambda_{voter(A)}$, where $Prins$ is the set of principals eligible to learn the election outcome, the join semilattice of underlying restrictions is $\langle 2^{Prins}, \cap, \supseteq \rangle$, and set \mathcal{F} of reclassifiers includes **tally**, which will be associated with calculating the election outcome.

Figure 9 gives a graphic depiction¹⁶ of $\lambda_{voter(A)}$; the formal definition is:

$$\lambda_{voter(A)} \triangleq \langle \{q_1, q_2\}, \mathcal{F}, \delta_{voter}, q_1, \rho_{voter(A)} \rangle$$

where

$$\delta_{voter}(q, f) \triangleq \begin{cases} q_2 & \text{if } q = q_1 \wedge f = \text{tally} \\ q & \text{otherwise} \end{cases}$$

Our proof by contradiction derives *false* from the assumption that a RIF label $\lambda_{\alpha''}$ satisfying $\lambda_{\alpha} \sqsubseteq_{RA} \lambda_{\alpha''}$, $\lambda_{\alpha'} \sqsubseteq_{RA} \lambda_{\alpha''}$, and $\lambda_{\alpha''} \sqsubset_{RA} \lambda_{\alpha} \sqcup_{RA} \lambda_{\alpha'}$ exists.

From $\lambda_{\alpha} \sqsubseteq_{RA} \lambda_{\alpha''}$ we have $\mathcal{R}_{RA}(\mathcal{T}_{RA}(\lambda_{\alpha}, F)) \sqsubseteq_R \mathcal{R}_{RA}(\mathcal{T}_{RA}(\lambda_{\alpha''}, F))$ for any F ; from $\lambda_{\alpha'} \sqsubseteq_{RA} \lambda_{\alpha''}$ we have $\mathcal{R}_{RA}(\mathcal{T}_{RA}(\lambda_{\alpha'}, F)) \sqsubseteq_R \mathcal{R}_{RA}(\mathcal{T}_{RA}(\lambda_{\alpha''}, F))$ for any F . Since $\langle R, \sqcup_R, \sqsubseteq_R \rangle$ defines a lattice, we conclude

$$\mathcal{R}_{RA}(\mathcal{T}_{RA}(\lambda_{\alpha}, F)) \sqcup_R \mathcal{R}_{RA}(\mathcal{T}_{RA}(\lambda_{\alpha'}, F)) \sqsubseteq_R \mathcal{R}_{RA}(\mathcal{T}_{RA}(\lambda_{\alpha''}, F)). \quad (37)$$

From assumption $\lambda_{\alpha''} \sqsubset_{RA} \lambda_{\alpha} \sqcup_{RA} \lambda_{\alpha'}$ we conclude (i) $\lambda_{\alpha''} \sqsubseteq_{RA} \lambda_{\alpha} \sqcup_{RA} \lambda_{\alpha'}$ and (ii) $\lambda_{\alpha''} \neq \lambda_{\alpha} \sqcup_{RA} \lambda_{\alpha'}$. But (37) conjoined with (i) contradict (ii).

¹⁶A conventional graphical representation for finite-state automata is used. Circles denote states of the automaton. Arrows between states are labeled with sets of reclassifiers and define allowed transitions, where $*$ abbreviates the list of all reclassifiers and \neg denotes set complement. The label inside each state q indicates associated restrictions $\rho_{\alpha}(q)$, and the grey-filled state indicates the current automaton state.




Figure 10: RIF automaton of document excerpting.

$$\rho_{voter(A)}(q) \triangleq \begin{cases} \{A\} & \text{if } q = q_1 \\ Prins & \text{if } q = q_2 \end{cases}$$

Restriction $\rho_{voter(A)}$ defines which principals can read the value being labeled by $\lambda_{voter(A)}$:

- the value is secret to A if $\lambda_{voter(A)}$ is in state q_1 , because $\rho_{voter(A)}(q_1) = \{A\}$, and
- any value derived by using a **tally** operation becomes public because the current state transitions to q_2 and $\rho_{voter(A)}(q_2) = Prins$ hold.

Thus, according to the terminology of Figure 1, **tally** causes declassification.

A programmer would use reclassifying expression $[maj(v_A, v_B, \dots, v_Z)]_{\text{tally}}$ to assert that computing the majority of votes in v_A, v_B, \dots, v_Z implements the intended effect of a **tally** operation. That derived value can be stored in a variable whose RIF label imposes no restriction on readers. Assignment

$$winner := [maj(v_A, v_B, \dots, v_Z)]_{\text{tally}} \quad (38)$$

has exactly that effect if the RIF automaton associated with each variable v_A is $\lambda_{voter(A)}$ in state q_1 , and the RIF automaton associated with variable $winner$ is at least as restrictive as

$$\lambda_{voter(A)}(\text{tally}) \sqcup_{\text{RA}} \lambda_{voter(B)}(\text{tally}) \sqcup_{\text{RA}} \dots \sqcup_{\text{RA}} \lambda_{voter(Z)}(\text{tally})$$

which happens to be equivalent to $\lambda_{voter(x)}[q_1 \mapsto q_2]$ for any $x \in \{A, \dots, Z\}$. So assignment (38) typechecks according to the rules of §4 if $winner$ has type $\lambda_{voter(x)}[q_1 \mapsto q_2]$ for some $x \in \{A, \dots, Z\}$.

A second example sketches RIF automata that enforce integrity policies for a document management system.¹⁷ A set of *original* documents is given;

¹⁷This example is inspired by TruDocs [50].

these are trusted by all principals for all purposes. Operation $ext(D, parms)$ derives a new document by excerpting from document D according to $parms$. Because “creative” excerpting can be used to generate a document that has different meaning from the original, cautious principals should hesitate to use such *derived* documents for certain purposes. Using the terminology of Figure 1, excerpting causes deprecation.

One RIF automaton λ_D for supporting such a policy might employ a join semilattice of underlying restrictions $\langle \{all, limited\}, \sqcup_R, \sqsubseteq_R \rangle$, where $limited \sqsubseteq_R all$ holds; these restrictions indicate whether a document is trusted for *all* purposes or for *limited* purposes. λ_D would employ a set $\mathcal{F}_{\text{Docs}}$ of reclassifiers that includes ext , which will correspond to excerpting operations. And λ_D would have two automaton states, where:

$$\rho_{\text{Docs}}(q) \triangleq \begin{cases} all & \text{if } q = q_1 \\ limited & \text{if } q = q_2 \end{cases}$$

Figure 10 gives a graphic depiction for a RIF automaton λ_D associated with an original document D ; the formal definition of λ_D is:

$$\lambda_D \triangleq \langle \{q_1, q_2\}, \mathcal{F}_{\text{Docs}}, \delta_{\text{Docs}}, q_1, \rho_{\text{Docs}} \rangle$$

where

$$\delta_{\text{Docs}}(q, f) \triangleq \begin{cases} q_2 & \text{if } q = q_1 \wedge f = \text{ext} \\ q & \text{otherwise} \end{cases}$$

RIF automaton $\lambda_D[q_1 \mapsto q_2]$ would be associated with any document produced by excerpting from D .

Some applications might require a more refined basis for deciding whether a document should be trusted for a specific purpose. One obvious basis for making such trust assessments is the set of principals participating in the document’s derivation. A RIF automaton can specify such policies. There would be an automaton state q_S for each set S of principals corresponding to a subset of $Prins$. And restrictions being associated with an automaton state q_S would depend on members of S . So the join semilattice of underlying restrictions is $\langle 2^{Prins}, \cup, \subseteq \rangle$. Transitions are facilitated by having a set \mathcal{F}_{Doc} of reclassifiers contain an element ext_p for each principal p that invokes an excerpting operation. Here is the formal definition of an automaton λ_D that is associated with a document D that some principal $W \in Prins$ has written:

$$\lambda_D \triangleq \langle Q_{Prins}, \mathcal{F}_{\text{Doc}}, \delta_{\text{Doc}}, q_{\{W\}}, \rho_{\text{Doc}} \rangle$$

where

$$\begin{aligned} Q_{Prins} &\triangleq \{q_S \mid S \in 2^{Prins}\} \\ \delta_{Doc}(q_S, \mathbf{exc}_p) &\triangleq q_{S \cup \{p\}} \\ \rho_{Doc}(q_S) &\triangleq S \end{aligned}$$

5.2 κ -Labels

Another class of RIF labels—called κ -labels—suffice to establish that a value v flows to v' only if principals authorized to read v' are (i) authorized to read v or (ii) unable to reverse some cryptographic operation used in generating v' from v . Our modelling of cryptographic operations follows Dolev and Yao [21].

The starting point for defining allowed flows is mapping KN .

- $KN(k)$ is the set of principals allowed to know the value of cryptographic key k .
- $KN(x)$ is the set of principals allowed to know the initial value of variable x .

The reclassifiers for κ -labels are associated with cryptographic operations, and rewrite rules for sequences of reclassifiers characterize when cryptographic operations are inverses. Different cryptosystems and cryptographic operations give rise to different reclassifiers along with different sets of rewrite rules.

κ -Atoms. A κ -atom κ_i for a value v_i to which v has flowed is a pair $\langle F_i, \beta_i \rangle$, where $F_i \in \mathcal{F}_\kappa^*$ gives the sequence of cryptographic operations involved in deriving v_i from v , and β_i is a set of principals authorized to read v . Let \mathcal{A}_κ denote the set of κ -atoms. Reclassifier $\theta(a_1, \dots, a_n) \in \mathcal{F}_\kappa$ is associated with a cryptographic operation θ whose invocation has a_1, \dots, a_n in its list of arguments, which may be cryptographic keys or other expressions.

A *rewrite rule* will have form

$$\theta(a) \theta'(a') \rightsquigarrow \epsilon$$

where $\theta(a)$ and $\theta'(a')$ are reclassifiers. This rewrite rule defines $\theta'(a')$ to be the *complement* of $\theta(a)$ and allows appearances of “ $\theta(a) \theta'(a')$ ” to be

deleted in a sequence of reclassifiers. Notation $\theta(a)^c$ will be used to denote the complement of reclassifier $\theta(a)$. Reclassifiers are not required to have complements, but if complements $\theta(a)^c$ and $(\theta(a)^c)^c$ both do exist then we require that $(\theta(a)^c)^c = \theta(a)$ holds.

To illustrate, we formulate κ -atoms for a symmetric-key cryptosystem having two cryptographic operations—encryption and decryption. $Enc(p, k)$ produces a *ciphertext* from *plaintext* p and key k . Its reclassifiers are defined by regarding appearances of $Enc(p, k)$ to be instances of reclassifying expression

$$[Enc(p, k)]_{\theta_{Enc}^1(k), \theta_{Enc}^2(p)}$$

which posits flows occur from p (associated with reclassifier $\theta_{Enc}^1(k)$) and from k (associated with reclassifier $\theta_{Enc}^2(p)$) to the value that $Enc(p, k)$ produces. $Dec(c, k)$ recovers the plaintext iff ciphertext c was previously encrypted using k . It has reclassifiers for the flows from c and from k to the value that $Dec(c, K)$ produces:

$$[Dec(c, k)]_{\theta_{Dec}^1(k), \theta_{Dec}^2(c)}$$

And we assume Enc and Dec satisfy certain properties.

$$Dec(Enc(v, k), k) = v \text{ holds for all values } v \text{ and keys } k. \quad (39)$$

$$Dec(c, k) \text{ is the only way to recover } p \text{ from } c = Enc(p, k). \quad (40)$$

$$\text{Key } k \text{ cannot be recovered from ciphertexts created using } k. \quad (41)$$

Property (39) stipulates that the effects of $Enc(\cdot, k)$ on plaintext can be reversed by using $Dec(\cdot, k)$; it is that basis for rewrite rule

$$\forall k: \quad \theta_{Enc}^1(k) \theta_{Dec}^1(k) \rightsquigarrow \epsilon. \quad (42)$$

Property (40) further implies that $\theta_{Dec}^1(k)$ is the only complement of $\theta_{Enc}^1(k)$. Thus, (40) stipulates an absence of certain rewrite rules. Finally, (41) prohibits rewrite rules that would define complements for reclassifiers associated with flows from k : $\theta_{Enc}^2(\cdot)$ and $\theta_{Dec}^2(\cdot)$. This prohibition corresponds to forbidding cryptographic functions where a key used as input to that function can be recovered from the output. It also rules out cryptographic systems that admit plaintext attacks.¹⁸

¹⁸A *plaintext attack* uses plaintext p and corresponding ciphertext c to recover the encryption key k satisfying $c = Enc(p, k)$. The attack can be viewed as having a crypto-

By definition, no reclassifier is followed by its complement in a *reduced sequence* $\langle F \rangle$ of reclassifiers, where $\langle F \rangle$ denotes the sequence obtained by repeatedly applying some given rewrite rules to F until no longer possible. If the full set of rewrite rules defines at most one complement for each reclassifier then (i) reduced sequence $\langle F \rangle$ is unique and (ii) the order in which rewrite rules are applied does not matter, so $\langle F \rangle$ is associative; see [31, Appendix A.3] for proofs.

All reclassifiers in rewrite rule (42) have a single argument; it is a key. We define set $\mathcal{X}(\theta(k))$ of principals that can recover a value produced by a cryptographic operation associated with a reclassifier $\theta(k)$ as follows.

$$\mathcal{X}(\theta(k)) \triangleq \begin{cases} \emptyset & \text{if complement } \theta(k)^c \text{ does not exist} \\ KN(k') & \text{if } \theta(k)^c = \theta'(k') \end{cases}$$

Set complement $\overline{\mathcal{X}(f)}$ thus contains those principals that cannot recover a value produced by the cryptographic operation identified by f .

If a value v' is derived from v by performing a sequence of cryptographic operations described with sequence F of reclassifiers then v can be recovered from v' only by those principals able to perform inverses of all operations mentioned in $\langle F \rangle$. So the set $\overline{\mathcal{X}(F)}$ of principals that cannot recover a value produced using a sequence F of cryptographic operations are those principals that cannot invert at least one of the operations in F . We characterize that set formally as follows, writing $f \in \langle F \rangle$ to indicate that f ranges over the reclassifiers appearing in reduced sequence $\langle F \rangle$.

$$\overline{\mathcal{X}(F)} \triangleq \bigcup_{f \in \langle F \rangle} \overline{\mathcal{X}(f)}$$

κ -atoms concern confidentiality for principals in a set $Prins$; subsets are more restrictive. Therefore, we use $\langle 2^{Prins}, \cap, \supseteq \rangle$ as the join semilattice

graphic function $[PA(p, c)]_{\theta_{PA}^1(c), \theta_{PA}^2(p)}$ that satisfies:

$$PA(p, c) = k \quad \text{if } c = Enc(p, k)$$

This semantics for $PA(p, c)$ would lead to rewrite rules—precluded by (41) stipulating that plaintext attacks are infeasible—for the flows from k in $Enc(p, k)$:

$$\theta_{Enc(p)}^2 \theta_{PA}^2(p) \mapsto \epsilon$$

$\langle R, \sqcup_R, \sqsubseteq_R \rangle$ of underlying restrictions. $\mathcal{R}_A(\langle F, \beta \rangle)$ is defined to be the set of principals to which the value associated with $\langle F, \beta \rangle$ may flow—the set of principals in β along with principals $\overline{\mathcal{X}}(F)$ that cannot recover an input to the sequence F of cryptographic operations:

$$\mathcal{R}_A(\langle F, \beta \rangle) \triangleq \beta \cup \overline{\mathcal{X}}(F)$$

And $\mathcal{T}_A(\langle F, \beta \rangle, f)$ specifies the flows allowed by $\langle F, \beta \rangle$ when its associated value is transformed by some cryptographic operation being identified with reclassifier f .

$$\mathcal{T}_A(\langle F, \beta \rangle, f) \triangleq \langle Ff, \beta \rangle$$

For example, an initial value v that can flow to principals in β_v would be associated with κ -atom $\langle \epsilon, \beta_v \rangle$. And κ -atom $\mathcal{T}_A(\langle F, \beta \rangle, f)$ associated with a transformed value potentially changes the set of principals where that transformed value might flow.¹⁹

The definition of restrictiveness relation \sqsubseteq_A on κ -atoms is obtained by substituting \mathcal{R}_A and \mathcal{T}_A for \mathcal{R} and \mathcal{T} in \sqsubseteq_Δ definition (7). Notice that for all $F \in \mathcal{F}_\kappa^*$

$$\langle F, \beta \rangle = \langle \langle F \rangle, \beta \rangle$$

and, therefore, an implementation of κ -atoms need only store reduced sequences.

A computable test to decide \sqsubseteq_A for κ -atoms can be obtained by extending complement sequences from Dolev-Yao [21], as follows. For any finite sequence F of reclassifiers, define *maximal complement sequence* F^c to be the sequence of reclassifiers that minimizes the length of $\langle F F^c \rangle$. F^c can be computed by taking the complement of each element in F , starting at the end.

$$F^c \triangleq \begin{cases} \epsilon & \text{if } F = \epsilon \\ f^c F_1^c & \text{if } F = F_1 f \text{ and } f^c \text{ exists} \\ \epsilon & \text{otherwise} \end{cases}$$

A test for \sqsubseteq_A is then given by:

$$\langle F_1, \beta_1 \rangle \sqsubseteq_A \langle F_2, \beta_2 \rangle = \mathcal{R}_A(\langle F_1 F_1^c, \beta_1 \rangle) \supseteq \mathcal{R}_A(\langle F_2 F_1^c, \beta_2 \rangle) \quad (43)$$

¹⁹The new set of principals is a superset when $\langle Ff \rangle$ extends $\langle F \rangle$ and $\overline{\mathcal{X}}(f) \not\subseteq \overline{\mathcal{X}}(F)$ holds. It is a subset when $\langle Ff \rangle$ is a proper prefix of $\langle F \rangle$, $\overline{\mathcal{X}}(f) \not\subseteq \overline{\mathcal{X}}(Ff)$ and $\overline{\mathcal{X}}(f^c) \not\subseteq \overline{\mathcal{X}}(Ff)$ hold. (If $\langle Ff \rangle$ is a proper prefix of $\langle F \rangle$ then f is the complement of the final reclassifier in $\langle F \rangle$, so $\langle Ff \rangle$ contains neither the final reclassifier in $\langle F \rangle$ nor the f being added.)

Soundness and completeness proofs for this test are given in [31, Appendix A.3].

κ -Labels. A κ -label \mathcal{K} is a finite set of κ -atoms. Given the following definitions for $\mathcal{R}_{\mathcal{L}}(\mathcal{K})$, and $\mathcal{T}_{\mathcal{L}}(\mathcal{K}, \mathfrak{f})$,

$$\mathcal{R}_{\mathcal{L}}(\mathcal{K}) \triangleq \sqcup_R \kappa \in \mathcal{K}: \mathcal{R}_A(\kappa) \quad (44)$$

$$\mathcal{T}_{\mathcal{L}}(\mathcal{K}, \mathfrak{f}) \triangleq \{\mathcal{T}_A(\kappa, \mathfrak{f}) \mid \kappa \in \mathcal{K}\} \quad (45)$$

κ -labels form a join semilattice $\langle 2^{\mathcal{A}^\kappa}, \sqcup_{\mathcal{L}}, \sqsubseteq_{\mathcal{L}} \rangle$ where $\sqcup_{\mathcal{L}}$ is defined by

$$\mathcal{K} \sqcup_{\mathcal{L}} \mathcal{K}' \triangleq \mathcal{K} \cup \mathcal{K}' \quad (46)$$

and restrictiveness relation $\sqsubseteq_{\mathcal{L}}$ is defined by substituting $\mathcal{R}_{\mathcal{L}}$ and $\mathcal{T}_{\mathcal{L}}$ into \sqsubseteq_{Λ} definition (7). The resulting class of RIF labels is then characterized by:

$$\langle \langle R, \sqcup_R, \sqsubseteq_R \rangle, \langle 2^{\mathcal{A}^\kappa}, \sqcup_{\mathcal{L}}, \sqsubseteq_{\mathcal{L}} \rangle, \mathcal{F}_\kappa, \mathcal{R}_{\mathcal{L}}, \mathcal{T}_{\mathcal{L}} \rangle \quad (47)$$

κ -labels in action. To illustrate, we typecheck a simple command that invokes operations of the symmetric cryptosystem formalized above. Suppose x is allowed to flow to some subset P_x of set *Prins* of all principals, principals in P_x are authorized to know key k , y is allowed to flow to some subset P_y of *Prins*, and principals in P_y are authorized to know key k' :

$$KN(x) = P_x \quad KN(k) = P_x \quad KN(y) = P_y \quad KN(k') = P_y$$

Here is a command that combines the symmetric-key cryptographic operations discussed above.

$$\begin{aligned} w &:= [Enc(x, k)]_{\theta_{Enc}^1(k), \theta_{Enc}^2(x)}; \\ y &:= [Dec(w, k')]_{\theta_{Dec}^1(k'), \theta_{Dec}^2(w)} \end{aligned} \quad (48)$$

A first assignment encrypts x using key k and then a second assignment attempts to decrypt that ciphertext using a key k' .

Given the assumptions above about x and k , we posit the following type declarations for the variables named in the first assignment of (48). Notice, the type for w corresponds to a value that has been encrypted under key k .

$$\Gamma(x) \triangleq \{\langle \epsilon, P_x \rangle\} \quad \Gamma(k) \triangleq \{\langle \epsilon, P_x \rangle\} \quad \Gamma(w) \triangleq \{\langle \theta_{Enc}^1(k), P_x \rangle\} \quad (49)$$

Using these types and typing rule ANNEXPRT in Figure 7, we obtain a type for the RHS of the first assignment of (48).

$$\mathcal{T}_{\mathcal{L}}(\Gamma(x), \theta_{Enc}^1(k)) \sqcup_{\mathcal{L}} \mathcal{T}_{\mathcal{L}}(\Gamma(k), \theta_{Enc}^2(x))$$

So, according to typing rule ASGN-T in Figure 8, that assignment is type correct provided the following holds.

$$(\mathcal{T}_{\mathcal{L}}(\Gamma(x), \theta_{Enc}^1(k)) \sqcup_{\mathcal{L}} \mathcal{T}_{\mathcal{L}}(\Gamma(k), \theta_{Enc}^2(x))) \sqsubseteq_{\mathcal{L}} \Gamma(w) \quad (50)$$

Substituting according to type declarations (49), obligation (50) for type correctness simplifies as follows:

$$\begin{aligned} & \{\langle \theta_{Enc}^1(k), P_x \rangle\} \cup \{\langle \theta_{Enc}^2(x), P_x \rangle\} \sqsubseteq_{\mathcal{L}} \{\langle \theta_{Enc}^1(k), P_x \rangle\} \\ & = \{\langle \theta_{Enc}^1(k), P_x \rangle, \langle \theta_{Enc}^2(x), P_x \rangle\} \sqsubseteq_{\mathcal{L}} \{\langle \theta_{Enc}^1(k), P_x \rangle\} \\ & = (\forall F: \mathcal{R}_{\mathcal{L}}(\{\langle \theta_{Enc}^1(k)F, P_x \rangle, \langle \theta_{Enc}^2(x)F, P_x \rangle\}) \sqsubseteq_{\mathcal{R}} \mathcal{R}_{\mathcal{L}}(\{\langle \theta_{Enc}^1(k)F, P_x \rangle\})) \\ & = (\forall F: \mathcal{R}_{\mathcal{A}}(\langle \theta_{Enc}^1(k)F, P_x \rangle) \sqcup_R \mathcal{R}_{\mathcal{A}}(\langle \theta_{Enc}^2(x)F, P_x \rangle) \sqsubseteq_{\mathcal{R}} \mathcal{R}_{\mathcal{A}}(\langle \theta_{Enc}^1(k)F, P_x \rangle)) \\ & = (\forall F: (P_x \cup \overline{\mathcal{X}}(\theta_{Enc}^1(k)F)) \cap (P_x \cup \overline{\mathcal{X}}(\theta_{Enc}^2(x)F)) \supseteq P_x \cup \overline{\mathcal{X}}(\theta_{Enc}^1(k)F)) \\ & = (\forall F: (P_x \cup \overline{\mathcal{X}}(\theta_{Enc}^1(k)F)) \cap (P_x \cup \text{Prins}) \supseteq P_x \cup \overline{\mathcal{X}}(\theta_{Enc}^1(k)F)) \\ & = (\forall F: (P_x \cup \overline{\mathcal{X}}(\theta_{Enc}^1(k)F)) \supseteq P_x \cup \overline{\mathcal{X}}(\theta_{Enc}^1(k)F)) \end{aligned}$$

The final formula is trivially equivalent to *true*.

Analogous reasoning establishes that the second assignment is type correct provided the following holds.

$$(\mathcal{T}_{\mathcal{L}}(\Gamma(w), \theta_{Dec}^1(k')) \sqcup_{\mathcal{L}} \mathcal{T}_{\mathcal{L}}(\Gamma(k'), \theta_{Dec}^2(w))) \sqsubseteq_{\mathcal{L}} \Gamma(y) \quad (51)$$

The validity of (51) depends on whether $k = k'$ holds and on the choice of P_y and $P_{k'}$ in the following proposed types for $\Gamma(y)$ and $\Gamma(k')$:

$$\Gamma(y) \triangleq \{\langle \epsilon, P_y \rangle\} \quad \Gamma(k') \triangleq \{\langle \epsilon, P_{k'} \rangle\}$$

- Case $k = k'$: From $\Gamma(k') = \Gamma(k)$, we conclude from $\Gamma(k)$ and $\Gamma(k')$ that $P_{k'} = P_x$ holds. Expanding (51) gets

$$\begin{aligned} & (\forall F: (P_x \cup \overline{\mathcal{X}}(\theta_{Enc}^1(k) \theta_{Dec}^1(k) F)) \cap (P_x \cup \overline{\mathcal{X}}(\theta_{Dec}^2(w) F)) \supseteq (P_y \cup \overline{\mathcal{X}}(F))) \\ & = \text{due to (41), } \overline{\mathcal{X}}(\theta_{Dec}^2(w) F) = \text{Prins} \\ & (\forall F: P_x \cup \overline{\mathcal{X}}(\theta_{Enc}^1(k) \theta_{Dec}^1(k) F) \supseteq (P_y \cup \overline{\mathcal{X}}(F))) \end{aligned}$$

$$\begin{aligned}
&= \text{since } \overline{\mathcal{X}}(\theta_{Enc}^1(k) \theta_{Dec}^1(k) F) \text{ simplifies to } \overline{\mathcal{X}}(F) \\
&(\forall F: P_x \cup \overline{\mathcal{X}}(F) \supseteq (P_y \cup \overline{\mathcal{X}}(F)))
\end{aligned}$$

This final formula is valid provided $P_x \supseteq P_y$ holds. So we have established that if $k = k'$ holds then the assignment is type correct provided $P_x \supseteq P_y$ holds. This conclusion should not be surprising, since if $P_x \not\supseteq P_y$ holds then once w has been decrypted, the value of x (which would now be stored in y) could flow to principals not in P_x (because they are in $P_y - P_x$).

- Case $k \neq k'$: Expanding (51) yields::

$$(\forall F: (P_x \cup \overline{\mathcal{X}}(\theta_{Enc}^1(k) \theta_{Dec}^1(k') F)) \cap (P_{k'} \cup \overline{\mathcal{X}}(\theta_{Dec}^2(w) F)) \supseteq (P_y \cup \overline{\mathcal{X}}(F)))$$

This formula is valid for all values of P_x and P_y since, by definition of $\overline{\mathcal{X}}(\cdot)$,

$$\overline{\mathcal{X}}(\theta_{Enc}^1(k) \theta_{Dec}^1(k') F) = Prins \text{ and } \overline{\mathcal{X}}(\theta_{Dec}^2(w) F) = Prins$$

because complements $\theta_{Dec}^1(k')^c$ and $\theta_{Dec}^2(w)^c$ do not exist. Thus, we conclude the assignment is type correct when $k \neq k'$, for any choices of P_y and $P_{k'}$. Again, this outcome should not be surprising. With no way to recover x from the value stored in y by the second assignment (since we have assumed that cryptographic operation $Dec(w, k')$ has no inverse), the second assignment cannot cause a flow violation.

The connection between type correctness and flows proved for this example command is an instance of a more-general guarantee about type correctness for this instance of κ -labels. To formalize this guarantee, note that any value a principal \mathfrak{p} could read must be tagged with a κ -label \mathcal{K}_L satisfying $\mathcal{K}_L \sqsubseteq_{\mathcal{L}} \mathcal{K}_{\mathfrak{p}}$, where

$$\mathcal{K}_{\mathfrak{p}} \triangleq \{ \langle F, \beta \rangle \in \mathcal{A}_{\kappa} \mid \mathfrak{p} \in \mathcal{R}_{\mathcal{A}}(\langle F, \beta \rangle) \}.$$

We consider a κ -label \mathcal{K} to be \mathfrak{p} -low iff $\mathcal{K} \sqsubseteq_{\mathcal{L}} \mathcal{K}_{\mathfrak{p}}$ holds and to be \mathfrak{p} -high iff $\mathcal{K} \not\sqsubseteq_{\mathcal{L}} \mathcal{K}_{\mathfrak{p}}$ holds. Thus the type-correctness guarantee provided by (any) κ -labels can be formulated as limiting what \mathfrak{p} can learn from values that flow from \mathfrak{p} -high to \mathfrak{p} -low—in particular, type correctness ensures that \mathfrak{p} cannot recover a \mathfrak{p} -high value from \mathfrak{p} -low values.

To show why this guarantee persists when using the κ -labels we defined above for the symmetric-key cryptosystem, consider a command \mathcal{C} that is type correct. Theorem 1 implies $PWNI(\mathcal{C})$ will hold and, therefore, if a value flows from \mathfrak{p} -high to \mathfrak{p} -low, then it does so through some operation that performs a $\mathcal{K}_{\mathfrak{p}}$ -downgrade. Since cryptographic operations $Enc(\cdot, \cdot)$ and $Dec(\cdot, \cdot)$ are the sole reclassifying expressions for our κ -labels, we conclude that all $\mathcal{K}_{\mathfrak{p}}$ -downgrades are caused by these operations. Without loss of generality, let such an operation be described by

$$[\theta(\mathcal{E}, k)]_{\theta^1(k), \theta^2(\mathcal{E})} \quad (52)$$

which has a flow from \mathcal{E} (associated with reclassifier $\theta^1(k)$) and a flow from k (associated with reclassifier $\theta^2(\mathcal{E})$) to the value that $\theta(\mathcal{E}, k)$ produces.

By definition, for (52) to exhibit a $\mathcal{K}_{\mathfrak{p}}$ -downgrade then

$$\Gamma([\theta(\mathcal{E}, k)]_{\theta^1(k), \theta^2(\mathcal{E})}) \sqsubseteq_{\mathcal{L}} \mathcal{K}_{\mathfrak{p}} \quad \wedge \quad \Gamma(\theta(\mathcal{E}, k)) \not\sqsubseteq_{\mathcal{L}} \mathcal{K}_{\mathfrak{p}}$$

must hold. From the first conjunct we get:

$$\mathfrak{p} \in \mathcal{R}_{\mathcal{L}}(\Gamma([\theta(\mathcal{E}, k)]_{\theta^1(k), \theta^2(\mathcal{E})})) \quad (53)$$

To validate the type-correctness guarantee, it suffices to prove that \mathfrak{p} cannot recover either \mathcal{E} or k from $\theta(\mathcal{E}, k)$ or, equivalently that the following hold.

$$\mathfrak{p} \in \overline{\mathcal{X}(\theta^1(k))} \quad \mathfrak{p} \in \overline{\mathcal{X}(\theta^2(\mathcal{E}))}$$

- Establishing $\mathfrak{p} \in \overline{\mathcal{X}(\theta^1(k))}$.

$$\begin{aligned} & \mathfrak{p} \in \mathcal{R}_{\mathcal{L}}(\Gamma([\theta(\mathcal{E}, k)]_{\theta^1(k), \theta^2(\mathcal{E})})) \\ &= \text{ANNEXPRT} \\ & \mathfrak{p} \in \mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(\Gamma(\mathcal{E}), \theta^1(k)) \sqcup_{\mathcal{L}} \mathcal{T}_{\mathcal{L}}(\Gamma(k), \theta^2(\mathcal{E}))) \\ &= \text{definition of } \mathcal{R}_{\mathcal{L}} \text{ where } \sqcup_R \text{ is } \cap \\ & \mathfrak{p} \in \mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(\Gamma(\mathcal{E}), \theta^1(k))) \cap \mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(\Gamma(k), \theta^2(\mathcal{E}))) \\ &= (\theta^2(\mathcal{E}))^c \text{ not defined so } \mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(\Gamma(k), \theta^2(\mathcal{E}))) = \text{Prins} \\ & \mathfrak{p} \in \mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(\Gamma(\mathcal{E}), \theta^1(k))) \\ &= \text{let } \Gamma(\mathcal{E}) = \{\langle F_1, \beta_1 \rangle, \dots, \langle F_n, \beta_n \rangle\} \\ & \mathfrak{p} \in (\mathcal{R}_{\mathcal{A}}(\langle F_1 \theta^1(k), \beta_1 \rangle) \sqcup_R \dots \sqcup_R \mathcal{R}_{\mathcal{A}}(\langle F_n \theta^1(k), \beta_n \rangle)) \\ &= \text{definition of } \mathcal{R}_{\mathcal{A}}; \sqcup_R \text{ is } \cap \end{aligned}$$

$$\begin{aligned}
& \mathbf{p} \in ((\beta_1 \cup \overline{\mathcal{X}}(F_1\theta^1(k))) \cap \dots \cap (\beta_n \cup \overline{\mathcal{X}}(F_n\theta^1(k)))) \\
& = \Gamma(\theta(\mathcal{E}, k)) \not\sqsubseteq_{\mathcal{L}} \mathcal{K}_{\mathbf{p}} \Rightarrow \mathbf{p} \notin ((\beta_1 \cup \overline{\mathcal{X}}(F_1)) \cap \dots \cap (\beta_n \cup \overline{\mathcal{X}}(F_n))) \\
& \quad \mathbf{p} \in \overline{\mathcal{X}(\theta^1(k))}
\end{aligned}$$

- Establishing $\mathbf{p} \in \overline{\mathcal{X}(\theta^2(\mathcal{E}))}$. According to (41), $\theta^2(\mathcal{E})$ has no complement. Thus, by definition, $\mathcal{X}(\theta^2(\mathcal{E})) = \text{Prins}$.

5.3 RIF Labels Based on RIF-Atoms

The construction of κ -labels from κ -atoms suggests a general approach: Have the RIF label for v be the set containing a *RIF-atom* associated with each v_i that flows to v . Generalizing from κ -atoms, a RIF-atom a_i would be specified by a pair $\langle F_i, \beta_i \rangle$, where F_i is a sequence of reclassifiers that transforms restrictions β_i describes. We then define RIF labels that are subsets of a set \mathcal{A} of such RIF-atoms, by defining functions

$$\mathcal{R}_{\mathcal{A}}: \mathcal{A} \rightarrow R \qquad \mathcal{T}_{\mathcal{A}}: \mathcal{A} \times \mathcal{F} \rightarrow \mathcal{A}$$

where $\mathcal{T}_{\mathcal{A}}(\langle F_i, \beta_i \rangle, \mathbf{f})$ appends \mathbf{f} to F_i and adjusts β_i . The result is a class \mathcal{L} of RIF labels

$$\langle \langle R, \sqcup_R, \sqsubseteq_R \rangle, \langle 2^{\mathcal{A}}, \sqcup_{\mathcal{L}}, \sqsubseteq_{\mathcal{L}} \rangle, \mathcal{F}_{\mathcal{L}}, \mathcal{R}_{\mathcal{L}}, \mathcal{T}_{\mathcal{L}} \rangle \quad (54)$$

where the labels are subsets of \mathcal{A} by defining

$$\mathcal{R}_{\mathcal{L}}: 2^{\mathcal{A}} \rightarrow R \qquad \mathcal{T}_{\mathcal{L}}: 2^{\mathcal{A}} \times \mathcal{F}_{\mathcal{L}} \rightarrow 2^{\mathcal{A}}$$

and defining, for RIF labels $A \subseteq \mathcal{A}$ and $A' \subseteq \mathcal{A}$,

$$A \sqcup_{\mathcal{L}} A' \triangleq A \cup A' \quad (55)$$

$$\mathcal{R}_{\mathcal{L}}(A) \triangleq \sqcup_R a \in A: \mathcal{R}_{\mathcal{A}}(a) \quad (56)$$

$$\mathcal{T}_{\mathcal{L}}(A, \mathbf{f}) \triangleq \{ \mathcal{T}_{\mathcal{A}}(a, \mathbf{f}) \mid a \in A \} \quad (57)$$

It is not hard to prove that $\langle 2^{\mathcal{A}}, \sqcup_{\mathcal{L}}, \sqsubseteq_{\mathcal{L}} \rangle$ is a join semilattice for $A \sqcup_{\mathcal{L}} A'$ defined by (55) and restrictiveness relation $\sqsubseteq_{\mathcal{L}}$ defined by substituting $\mathcal{R}_{\mathcal{L}}$ and $\mathcal{T}_{\mathcal{L}}$ into \sqsubseteq_{Λ} definition (7). Notice also that (6), relating \sqcup_R and $\sqcup_{\mathcal{L}}$, is straightforward to establish by expanding definitions (55) and (56).

Deciding $\sqsubseteq_{\mathcal{L}}$. Instances of class \mathcal{L} admit various computable—albeit, conservative—tests for deciding $\sqsubseteq_{\mathcal{L}}$. By being conservative, such a test may offer an affirmation that $A \sqsubseteq_{\mathcal{L}} A'$ holds for a pair of RIF labels $A \subseteq \mathcal{A}$ and $A' \subseteq \mathcal{A}$, but nothing can be concluded when the test does not offer an affirmation.

One test for $A \sqsubseteq_{\mathcal{L}} A'$ involves checking whether $A \subseteq A'$ holds:

$$A \subseteq A' \Rightarrow A \sqsubseteq_{\mathcal{L}} A' \quad (58)$$

By using (58), a property ($A \sqsubseteq_{\mathcal{L}} A'$) that involves quantification over all finite sequences of reclassifiers is being affirmed simply by checking a subset relation on bounded-size sets—a check that is feasible to discharge.

Here is the justification for (58). By definition, $A \sqsubseteq_{\mathcal{L}} A'$ holds if

$$\mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(A, F)) \sqsubseteq_R \mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(A', F))$$

holds for each $F \in \mathcal{F}_{\mathcal{L}}^*$. So it suffices to prove that

$$\mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(A, F)) \sqsubseteq_R \mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(A', F))$$

is equivalent to *true* if $A \subseteq A'$ holds:

$$\begin{aligned} & \mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(A, F)) \sqsubseteq_R \mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(A', F)) \\ &= \text{definition (56) of } \mathcal{R}_{\mathcal{L}}(A) \\ & (\sqcup_R a \in \mathcal{T}_{\mathcal{L}}(A, F): \mathcal{R}_{\mathcal{A}}(a)) \sqsubseteq_R (\sqcup_R a \in \mathcal{T}_{\mathcal{L}}(A', F): \mathcal{R}_{\mathcal{A}}(a)) \\ &= A \subseteq A' \text{ is equivalent to } A' = (A \cup (A' - A)) \text{ and (55)} \\ & (\sqcup_R a \in \mathcal{T}_{\mathcal{L}}(A, F): \mathcal{R}_{\mathcal{A}}(a)) \sqsubseteq_R \\ & (\sqcup_R a \in \mathcal{T}_{\mathcal{L}}(A, F): \mathcal{R}_{\mathcal{A}}(a)) \sqcup_R (\sqcup_R a \in \mathcal{T}_{\mathcal{L}}(A - A', F): \mathcal{R}_{\mathcal{A}}(a)) \\ &= \text{since } x \sqsubseteq_R (x \sqcup y) \text{ for all } x \in R, y \in R \\ & \textit{true} \end{aligned}$$

A second test for $A \sqsubseteq_{\mathcal{L}} A'$ is sometimes possible by defining restrictiveness relation $a \sqsubseteq_{\mathcal{A}} a'$ on the RIF-atoms that constitute RIF labels A and A' . If there is a feasible test for $a \sqsubseteq_{\mathcal{A}} a'$, then we can leverage the monotonicity of \sqcup_R with respect to \sqsubseteq_R to obtain a test for $A \sqsubseteq_{\mathcal{L}} A'$. That test is described by

$$(\forall a \in A, a' \in A': a \sqsubseteq_{\mathcal{A}} a') \Rightarrow A \sqsubseteq_{\mathcal{L}} A' \quad (59)$$

where the definition of $\sqsubseteq_{\mathcal{A}}$ is obtained by substituting $\mathcal{R}_{\mathcal{A}}$ and $\mathcal{T}_{\mathcal{A}}$ for \mathcal{R} and \mathcal{T} in \sqsubseteq_{Λ} definition (7). Here is a proof that (59) is valid.

$$\begin{aligned}
& (\forall a \in A, a' \in A': a \sqsubseteq_{\mathcal{A}} a') \\
& \Rightarrow \text{definition of } \sqsubseteq_{\mathcal{A}} \text{ and monotonicity of } \sqcup_R \text{ with respect to } \sqsubseteq_R \\
& (\forall F \in \mathcal{F}_{\mathcal{L}}^*: (\sqcup_R a \in A: \mathcal{R}_{\mathcal{A}}(\mathcal{T}_{\mathcal{A}}(a, F))) \sqsubseteq_{\mathcal{R}} (\sqcup_R a' \in A': \mathcal{R}_{\mathcal{A}}(\mathcal{T}_{\mathcal{A}}(a', F)))) \\
& = \text{rearrange bound variables} \\
& (\forall F \in \mathcal{F}_{\mathcal{L}}^*: (\sqcup_R a \in \mathcal{T}_{\mathcal{A}}(A, F): \mathcal{R}_{\mathcal{A}}(a)) \sqsubseteq_{\mathcal{R}} (\sqcup_R a' \in \mathcal{T}_{\mathcal{A}}(A', F): \mathcal{R}_{\mathcal{A}}(a'))) \\
& = \text{definition of } \mathcal{R}_{\mathcal{L}}, \mathcal{T}_{\mathcal{L}} \\
& (\forall F \in \mathcal{F}_{\mathcal{L}}^*: \mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(A, F)) \sqsubseteq_{\mathcal{R}} \mathcal{R}_{\mathcal{L}}(\mathcal{T}_{\mathcal{L}}(A', F))) \\
& = \text{definition of } \sqsubseteq_{\mathcal{L}} \\
& A \sqsubseteq_{\mathcal{L}} A'
\end{aligned}$$

We use RIF automata set Λ_{RA} to illustrate this second test, as follows. Let \mathcal{A} be the set of RIF-atoms $\langle F_i, \lambda_i \rangle$, where $F_i \in \mathcal{F}^*$ and $\lambda_i \in \Lambda_{\text{RA}}$, and define:

$$\begin{aligned}
\mathcal{R}_{\mathcal{A}}(\langle F, \lambda_i \rangle) & \triangleq \mathcal{R}_{\text{RA}}(\lambda_i(F)) \\
\mathcal{T}_{\mathcal{A}}(\langle F, \lambda_i \rangle, f) & \triangleq \langle Ff, \lambda_i \rangle
\end{aligned}$$

Expanding definitions, we get:

$$\langle F, \lambda \rangle \sqsubseteq_{\mathcal{A}} \langle F', \lambda' \rangle \Leftrightarrow \langle \epsilon, \lambda(F) \rangle \sqsubseteq_{\mathcal{A}} \langle \epsilon, \lambda'(F') \rangle \quad (60)$$

The definition of $\langle F, \lambda \rangle \sqsubseteq_{\mathcal{A}} \langle F', \lambda' \rangle$ (in the hypothesis of (59)) involves a quantification over the infinite set of sequences $F'' \in \mathcal{F}^*$:

$$(\forall F'' \in \mathcal{F}^*: \mathcal{R}_{\text{RA}}(\lambda(FF'')) \sqsubseteq_{\mathcal{R}} \mathcal{R}_{\text{RA}}(\lambda'(F'F''))) \quad (61)$$

However, for a RIF automaton λ_{α} in current state q_{α} with set of states Q_{α} , we have

$$\{\lambda_{\alpha}(F) \mid F \in \mathcal{F}^*\} \subseteq \{\lambda_{\alpha}[q_{\alpha} \mapsto q_i] \mid q_i \in Q_{\alpha}\}$$

Because Q_{α} is finite, we conclude that $\{\lambda_{\alpha}(F) \mid F \in \mathcal{F}^*\}$ must be finite. Thus, (61) can be discharged by checking the finite subset of sequences $F'' \in \mathcal{F}^*$ —specifically those sequences $F'' \in \mathcal{F}^*$ satisfying $|F''| \leq |Q_{\alpha}| \times |Q_{\alpha'}|$, where Q_{α} is the set of states for RIF automaton λ and $Q_{\alpha'}$ is the set of states for RIF automaton λ' .

Another approach to obtaining feasible test for $\langle F, \lambda \rangle \sqsubseteq_{\mathcal{A}} \langle F', \lambda' \rangle$ in the hypothesis of (59) is to give rewrite rules that transform one sequence of reclassifiers into another. The goal is a set of rewrite rules where checking a property of all sequences F is equivalent to checking that property for an easily computed finite set of reduced sequences ($\downarrow F$). The set of rewrite rules are determined by the semantics for the operations associated with each given reclassifier. κ -labels, discussed above, are an example of this approach.

6 Related Work

RIF labels specify restrictions that depend on the history of applied operations. And our static type system ensures that programs using RIF labels will satisfy PWNI, which extends noninterference [26] to accommodate reclassification. Our work thus extends an approach initiated by Volpano *et al.* [55]—a type system based on Denning’s lattice model [20, 18, 19] that enforces noninterference.

We are not the first to explore types that depend on a history of applied operations. Strom and Yemini [53] propose types that incorporate *typestate* to summarize the history of operations previously applied to an object and to govern the set of operations that may next be invoked. But Hartson and Hsiao [28] seem to be the first to use history in access control. Contemporary uses of history for defining authorization include *stack inspection* [57] and *history-based access control* [1].

Reclassification: Specification and Enforcement

State Based Reclassification. A good deal of prior work ties reclassification to changes in state predicates. Chong *et al.* [15] is closest to RIF labels. There, information flow policies for confidentiality are defined in terms of state predicates that specify when a value may be reclassified. A static type system enforces these declassification and erasure policies. And when a reclassification occurs in the approach of Chong *et al.* [15], the value and all derived values are reclassified—in contrast, with RIF labels, reclassifications apply only to the single value.

With ClickRelease, Micinski *et al.* [40] extend the approach in Chong *et al.* [15] by allowing a more-expressive language for the formulas that specify when values may be declassified. In particular, ClickRelease uses temporal

logic formulas over events, in contrast to the state predicates used by Chong *et al.* [15]

Paralocks [12] formulates a security policy as $\Sigma \Rightarrow \alpha$, where Σ is a set of state predicates and α is a principal. A value associated with such a policy is allowed to flow to α only if all predicates in Σ are true. Changes to state predicates in Σ alter the allowed flows. Paralocks policies are enforced using a static type system.

A policy in Thoth [23] and its successor system SHAI [22] specifies confidentiality and integrity restrictions for data containers called *conduits*. The policies comprise two layers. The *access control layer* specifies which principals may read and update the associated conduit and under what conditions. The *declassification layer* specifies conditions under which policies for data derived from the associated conduit can change. The conditions employed by a Thoth policy are predicates on conduits’ state, which include both data and metadata (including a type) of conduits. Thoth uses dynamic analysis for enforcement, SHAI uses static analysis too.

Jeeves [5] employs *faceted values* [4] to specify declassification between two levels of confidentiality (i.e., **public** and **secret**). A faceted value is a pair comprising a real and a dummy value, guarded by a state predicate. If that state predicate (which itself can be a faceted value) holds, then the real, possibly secret, value is allowed to flow to low outputs. Otherwise, the dummy value flows to low outputs.

All of this prior work thus ties reclassification to changes in state predicates. In contrast, RIF labels tie reclassification to expression evaluation. A comparison of RIF labels with reclassification based on state predicates thus depends first on what state predicates are available and second on whether those state predicates could be incorporated into RIF transition function \mathcal{T} .

Operation Based Reclassification. Explicit expressions for declassification (for confidentiality) and endorsement (for integrity) have been proposed [41, 43]. However, the approach can be unsatisfying because output restrictions are not connected to input restrictions or to the operation that transforms inputs to outputs. For example, [41, 43] allows an arbitrary label to be assigned to the result of evaluating any expression—type-casting has the same flavor.

RIF labels tie reclassification to expression evaluation, but it is not the first work to connect restrictions on outputs to an operation. With FJifP [29], a principal declares *trusted methods* to declassify input values.²⁰ A FJifP

²⁰Trusted methods are similar to *trusted subjects*, first introduced by Bell and La-

trusted method always performs the same declassification of an input whereas, depending on the RIF label associated with an input value, a specific re-classifier in a reclassifying expression could trigger different changes to the restrictions. So reclassifying expressions (in conjunction with RIF labels) are more expressive than trusted methods.

Rocha *et al.* [44] propose *policy graphs* to specify declassification of information from **secret** to **public** (as compared to RIF labels, which handle arbitrary reclassification); similar techniques to those presented in [44] are later employed by Hammer *et al.* [27] and Johnson *et al.* [30]. In a policy graph, nodes represent variables and edges represent operation identifiers (similar to our reclassifiers). The tail node of an edge is an input of the corresponding operation, and the head node of that edge is an output of that operation. Some of the nodes in a policy graph are defined *final*. Values in variables of non-final nodes are considered **secret**, whereas values in variables that correspond to final nodes are declassified and considered **public**. Data and control-flow analysis is used to check whether some given program satisfies a specific policy graph. Subsequently, Rocha *et al.* [45] introduce a specification language for defining policies that might also depend on how many times a function is applied to a given value. This specification language seems more expressive than Rocha *et al.* [44], although the properties being enforced have not been formally characterized. We believe restrictions defined using any of these policy graphs could be described using a set of RIF automata.

Li and Zdancewic [37] suggest that downgrades between two security levels be specified as lambda terms. Apply one of these lambda terms to a value and the result, by definition, is given a downgraded label. (A type system is given to enforce these policies.) This approach to characterizing downgrades is attractive because it is independent from the program code. Enforcement, however, involves a conservative approach to deciding equivalence of functions because that problem is undecidable in general. Also, the approach is not well suited for handling reclassifications based on how a value has been derived, whereas RIF labels do handle that.

Padula [7] to handle declassification.

Extending Noninterference for Reclassification

Declassification violates classical noninterference [26], which has prompted researchers to develop alternatives. One example is *conditional noninterference* [26, 25]. It proscribes **secret** information flows to **public** information—unless a given predicate is satisfied by (i) the sequence of operations involved in this flow and (ii) the principals invoking those operations. For declassification, conditional noninterference is more expressive than PWNI, because PWNI ignores identities of the principals that invoke operations. PWNI, however, handles reclassification in full generality (i.e., arbitrary lattices of labels and declassification as well as classification), in contrast to conditional noninterference, which only deals with declassification in a 2-level lattice.

Gradual release (GR) [3] dictates that declassification of expressions, encrypted exceptions, and released cryptographic keys are the only execution points where an attacker’s knowledge about initial secret values may increase. Other work (*delimited release* [47] and *relaxed noninterference* [37]) specifies what expressions of **secret** values in the initial state could be declassified, with no restriction about where in the program such expressions are declassified. With *conditional gradual release* [6], as with our approach, any expression in a program may be declassified. In doing so, however, it is not always possible to determine what initial **secret** values become available to an observer at a given declassification.²¹ Moreover, conditional gradual release allows declassifications to depend on **secret** guards, causing a declassification to disclose more information than might have been intended. (PWNI does not suffer this defect, because of the way λ -pieces are defined.)

Non-disclosure policy [39] is a variation of noninterference for handling local declassification constructs. These declassification constructs augment allowed flows during execution of some code M , restoring the previously allowed flows once execution leaves M . To satisfy the non-disclosure policy, noninterference must hold for flow relations that are allowed at every execution step. This policy, then, embodies a different design choice from PWNI about the scope of a declassification—with PWNI the declassification persists to the end of the execution, whereas with the non-disclosure policy it ends when the declassification construct has completed. The same design choice is adopted in [11] for handling flow-locks [10].

²¹The restriction on programs imposed by [6] and [47]—namely that declassified variables may not be updated prior to declassification— would allow determination of what **secret** values become available to an observer.

Other alternatives to noninterference associate declassification with special commands, such as match queries²² [56], one-way functions [54], or various other cryptographic operations. With *computational noninterference* (CNI) [34], disclosure of **secret** values is permitted only when those values are encrypted; CNI is enforced by type systems introduced in Laud *et al.* [36] for *passive* and Fournet *et al.* [24] for *active* adversaries. Smith *et al.* [51] propose a variant of noninterference that handles both encryption and decryption, and it is enforced using a type system.

κ -labels demonstrate how RIF labels can handle cryptographic operations by treating these operations as special and translating them into reclassifying expressions. The literature on type checking of cryptographic operations has explored two general approaches [16]: computational and symbolic. κ -labels embrace the symbolic approach, using an analysis approach derived from Dolev-Yao [21]. *Cryptographically masked flows* [2] also employs a symbolic analysis, and it too is enforced by type systems. Laud [35] shows that type correctness according to [2], together with some additional conditions, imply CNI, thereby establishing a connection between cryptographically masked flows (which is based on symbolic analysis) and CNI (which is based on computational analysis). Cortier *et al.* [16] generalizes this connection by showing that programs secure according to a symbolic analysis are also secure according to a computational analysis.

Semantic properties have been proposed to handle erasure, too. Del Tedesco *et al.* [17] use knowledge semantics to express a hierarchy of erasure policies. These erasure policies are categorized based on (i) whether they specify erasure of all or part of the information, and (ii) whether erasure depends on program state (either high or low). Erasure is a form of classification, and thus, can be specified by RIF labels. However, with the RIF labels in this paper, erasure only affects a value being derived and cannot be formulated to depend on program state.

Noninterference and its variations (including PWNI) characterize allowed flows of information; they do not handle required flows. Chong [14] gives a semantic definition for *required information release* policies and presents a static type system to enforce these policies. A required information release policy specifies what information should be released and how this information can be learned by the authorized observers. The semantics is based on

²²A *match query* checks whether two objects are equal. For example, a match query is used to check whether a certain string is the password of a given user.

algorithmic knowledge.

Views of the Reclassification Landscape

The survey by Sabelfeld *et al.* [49] introduces a four-dimensional categorization for declassification policies (though the categorization seems applicable reclassification, too): *what* information is declassified, *who* declassifies information, *where* in the system information is declassified, and *when* information can be declassified. RIF labels and our reclassifying expressions specify *what*, *where*, and *when*, but not *who*. The *what* is the value produced by the expression; *where* is the position of the reclassifying expression in the program text; *when* is determined by the program’s control flow.

Nothing prevents the semantics of our reclassifiers from incorporating information about *who* is evaluating a given reclassifying expression. The *decentralized label model* (DLM) [42, 41] is an obvious starting point for such an extension. According to DLM, a value may be declassified only if the declassification command is executed on behalf of the value’s *owner* or on behalf of a principal that *acts-for* that owner. To adopt this approach, we would add an additional input argument to \mathcal{T} —the identity of the principal undertaking the reclassification. The semantics of \mathcal{T} would then be extended so that a reclassifier triggers a transition only for certain principals.

Broberg *et al.* [13] offers an orthogonal view for information flow policies and declassification. It is based on a three-level *hierarchy of control*. *Level 0 control* is a set of possible *flow relations* between information *sources* (e.g., input variables) and *sinks* (e.g., output channels). A flow relation indicates that information from the source is allowed to flow to the sink. *Level 1 controls* select which flow relations are allowed. *Level 2 controls* constitute a meta policy for controlling the way in which the current flow relations (Level 1 controls) may be changed. RIF label function \mathcal{T} incorporates aspects of *Level 1* and *Level 2* controls.

Acknowledgments

We appreciate the comments we received on early discussions of this work from Eleanor Birrell, Michael Clarkson, Josée Desharnais, Michael George, José Meseguer, Brad Martin, Andrew Myers, David Naumann, Karn Seth, Geoffrey Smith, and Nadia Tawbi. Steve Chong and Deepak Garg, in addition, provided detailed and very helpful feedback on an earlier draft of this

paper.

References

- [1] M. Abadi and C. Fournet. Access control based on execution history. In *Proceedings of the 10th Annual Network and Distributed System Security Symposium*, pages 107–121, 2003.
- [2] A. Askarov, D. Hedin, and A. Sabelfeld. Cryptographically-masked flows. *Theor. Comput. Sci.*, 402(2-3):82–101, July 2008.
- [3] A. Askarov and A. Sabelfeld. Gradual release: Unifying declassification, encryption and key release policies. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 207–221, 2007.
- [4] T. H. Austin and C. Flanagan. Multiple facets for dynamic information flow. In *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '12, pages 165–178, New York, NY, USA, 2012. ACM.
- [5] T. H. Austin, J. Yang, C. Flanagan, and A. Solar-Lezama. Faceted execution of policy-agnostic programs. In *Proceedings of the 8th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security*, PLAS '13, pages 15–26, New York, NY, USA, 2013. ACM.
- [6] A. Banerjee, D. Naumann, and S. Rosenberg. Expressive declassification policies and modular static enforcement. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 339–353, 2008.
- [7] D. Bell and L. La Padula. Secure computer systems: Unified exposition and MULTICS interpretation. Technical Report ESD-TR-75306, Bedford, MA, 1976.
- [8] E. Birrell and F. B. Schneider. A reactive approach to use-based privacy. Technical Report 54843, Cornell University, Computing and Information Science, November 2017.
- [9] N. Broberg, B. Delft, and D. Sands. Paragon for practical programming with information-flow control. In *Proceedings of the 11th Asian Sym-*

- posium on Programming Languages and Systems*, volume 8301, pages 217–232, New York, NY, USA, 2013. Springer-Verlag New York, Inc.
- [10] N. Broberg and D. Sands. Flow locks: Towards a core calculus for dynamic flow policies. In *Proceedings of the 15th European Conference on Programming Languages and Systems*, ESOP’06, pages 180–196, Berlin, Heidelberg, 2006. Springer-Verlag.
 - [11] N. Broberg and D. Sands. Flow locks: Towards a core calculus for dynamic flow policies. Technical report, Chalmers University of Technology and Göteborgs University, May 2006.
 - [12] N. Broberg and D. Sands. Paralocks: Role-based information flow control and beyond. In *Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL ’10, pages 431–444, New York, NY, USA, 2010. ACM.
 - [13] N. Broberg, B. van Delft, and D. Sands. The anatomy and facets of dynamic policies. In *Proceedings of the 28th IEEE Computer Security Foundations Symposium*, pages 122–136, July 2015.
 - [14] S. Chong. Required information release. In *Proceedings of the 23rd IEEE Computer Security Foundations Symposium*, pages 215–227, Piscataway, NJ, USA, July 2010. IEEE Press.
 - [15] S. Chong and A. C. Myers. End-to-end enforcement of erasure and declassification. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium*, pages 98–111, June 2008.
 - [16] V. Cortier, S. Kremer, and B. Warinschi. A survey of symbolic methods in computational analysis of cryptographic systems. *J. Autom. Reason.*, 46(3-4):225–259, Apr. 2011.
 - [17] F. Del Tedesco, S. Hunt, and D. Sands. A semantic hierarchy for erasure policies. In *Proceedings of the 7th International Conference on Information Systems Security*, ICISS’11, pages 352–369, Berlin, Heidelberg, 2011. Springer-Verlag.
 - [18] D. E. Denning. A lattice model of secure information flow. *Commun. ACM*, 19(5):236–243, May 1976.

- [19] D. E. Denning and P. J. Denning. Certification of programs for secure information flow. *Commun. ACM*, 20(7):504–513, July 1977.
- [20] D. E. R. Denning. *Secure information flow in computer systems*. PhD thesis, Purdue University, West Lafayette, IN, USA, 1975.
- [21] D. Dolev and A. C. Yao. On the security of public key protocols. *Information Theory, IEEE Transactions on*, 29(2):198–208, Mar 1983.
- [22] E. Elnikety, D. Garg, and P. Druschel. SHAI: Enforcing Data-Specific Policies with Near-Zero Runtime Overhead. Technical report, Max Planck Institute for Software Systems, Saarland Informatics Campus, Germany, January 2018.
- [23] E. Elnikety, A. Mehta, A. Vahldiek-Oberwagner, D. Garg, and P. Druschel. Thoth: Comprehensive policy compliance in data retrieval systems. In *Proceedings of the 25th USENIX Conference on Security Symposium, SEC'16*, pages 637–654. USENIX Association, 2016.
- [24] C. Fournet and T. Rezk. Cryptographically sound implementations for typed information-flow security. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '08*, pages 323–335, New York, NY, USA, 2008. ACM.
- [25] J. A. Goguen and J. Meseguer. Unwinding and inference control. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 75–87, 1984.
- [26] J. A. Goguen and J. Meseguer. Security policies and security models. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 11–20, 1982.
- [27] C. Hammer and G. Snelting. Flow-sensitive, context-sensitive, and object-sensitive information flow control based on program dependence graphs. *International Journal of Information Security*, 8(6):399–422, 2009.
- [28] H. R. Hartson and D. K. Hsiao. Full protection specifications in the semantic model for database protection languages. In *Proceedings of the 1976 Annual Conference, ACM '76*, pages 90–95, New York, NY, USA, 1976. ACM.

- [29] B. Hicks, D. King, P. McDaniel, and M. Hicks. Trusted declassification: High-level policy for a security-typed language. In *Proceedings of the Workshop on Programming Languages and Analysis for Security*, PLAS '06, pages 65–74, New York, NY, USA, 2006. ACM.
- [30] A. Johnson, L. Wayne, S. Moore, and S. Chong. Exploring and enforcing security guarantees via program dependence graphs. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '15, pages 291–302, New York, NY, USA, 2015. ACM.
- [31] E. Kozyri. *Enhancing Expressiveness of Information Flow Labels: Reclassification and Permissiveness*. PhD thesis, Cornell University, Ithaca, New York, USA, 2018. <https://search.proquest.com/docview/2167492985?accountid=10267>.
- [32] E. Kozyri, O. Arden, A. C. Myers, and F. B. Schneider. JRIF: Reactive Information Flow Control for Java. Technical report, Cornell University, February 2016.
- [33] M. Krohn, A. Yip, M. Brodsky, N. Cliffer, M. F. Kaashoek, E. Kohler, and R. Morris. Information flow control for standard OS abstractions. In *Proceedings of the 21st ACM SIGOPS Symposium on Operating Systems Principles*, SOSP '07, pages 321–334, New York, NY, USA, 2007. ACM.
- [34] P. Laud. Semantics and program analysis of computationally secure information flow. In *Proceedings of the 10th European Symposium on Programming Languages and Systems*, ESOP '01, pages 77–91, London, UK, 2001. Springer-Verlag.
- [35] P. Laud. On the computational soundness of cryptographically masked flows. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '08, pages 337–348, New York, NY, USA, 2008. ACM.
- [36] P. Laud and V. Vene. A type system for computationally secure information flow. In *Proceedings of the 15th International Conference on Fundamentals of Computation Theory*, FCT'05, pages 365–377, Berlin, Heidelberg, 2005. Springer-Verlag.

- [37] P. Li and S. Zdancewic. Downgrading policies and relaxed noninterference. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '05, pages 158–170, New York, NY, USA, 2005. ACM.
- [38] P. Li and S. Zdancewic. Practical information-flow control in web-based information systems. In *Proceedings of the 18th IEEE Workshop on Computer Security Foundations*, CSFW '05, pages 2–15, Washington, DC, USA, 2005. IEEE Computer Society.
- [39] A. A. Matos and G. Boudol. On declassification and the non-disclosure policy. In *Proceedings of the 18th IEEE Computer Security Foundations Workshop (CSFW'05)*, pages 226–240, June 2005.
- [40] K. Micinski, J. Fetter-Degges, J. Jeon, J. S. Foster, and M. R. Clarkson. Checking interaction-based declassification policies for Android using symbolic execution. In *Proceedings of the European Symposium on Research in Computer Security*, ESORICS 2015, pages 520–538, Cham, 2015. Springer International Publishing.
- [41] A. C. Myers. Jflow: Practical mostly-static information flow control. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '99, pages 228–241, New York, NY, USA, 1999. ACM.
- [42] A. C. Myers and B. Liskov. A decentralized model for information flow control. In *Proceedings of the 16th ACM Symposium on Operating Systems Principles*, SOSP '97, pages 129–142, New York, NY, USA, 1997. ACM.
- [43] A. C. Myers, L. Zheng, S. Zdancewic, S. Chong, and N. Nystrom. Jif 3.0: Java information flow. Software release. <http://www.cs.cornell.edu/jif>, July 2006.
- [44] B. Rocha, S. Bandhakavi, J. den Hartog, W. Winsborough, and S. Etalle. Towards static flow-based declassification for legacy and untrusted programs. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 93–108, 2010.

- [45] B. Rocha, M. Conti, S. Etalle, and B. Crispo. Hybrid static-runtime information flow and declassification enforcement. *Information Forensics and Security, IEEE Transactions on*, 8(8):1294–1305, 2013.
- [46] I. Roy, D. E. Porter, M. D. Bond, K. S. McKinley, and E. Witchel. Laminar: Practical fine-grained decentralized information flow control. In *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '09*, pages 63–74, New York, NY, USA, 2009. ACM.
- [47] A. Sabelfeld and A. C. Myers. A model for delimited information release. In *Proceedings of the International Symposium on Software Security (ISSS03), volume 3233 of LNCS*, pages 174–191. Springer-Verlag, 2004.
- [48] A. Sabelfeld and D. Sands. A per model of secure information flow in sequential programs. In *Proceedings of the 8th European Symposium on Programming Languages and Systems, ESOP '99*, pages 40–58, Berlin, Heidelberg, 1999. Springer-Verlag.
- [49] A. Sabelfeld and D. Sands. Declassification: Dimensions and principles. *J. Comput. Secur.*, 17(5):517–548, Oct. 2009.
- [50] F. B. Schneider, K. Walsh, and E. G. Sirer. Nexus Authorization Logic (NAL): Design rationale and applications. *ACM Trans. Inf. Syst. Secur.*, 14(1):8:1–8:28, June 2011.
- [51] G. Smith and R. Alpizar. Secure information flow with random assignment and encryption. In *Proceedings of the 4th ACM Workshop on Formal Methods in Security, FMSE '06*, pages 33–44, New York, NY, USA, 2006. ACM.
- [52] D. Stefan, A. Russo, J. C. Mitchell, and D. Mazières. Flexible dynamic information flow control in Haskell. In *Proceedings of the 4th ACM Symposium on Haskell, Haskell '11*, pages 95–106, New York, NY, USA, 2011. ACM.
- [53] R. E. Strom and S. Yemini. Timestep: A programming language concept for enhancing software reliability. *IEEE Trans. Softw. Eng.*, 12(1):157–171, Jan. 1986.

- [54] D. Volpano. Secure introduction of one-way functions. In *Proceedings of the 13th IEEE Computer Security Foundations Workshop*, pages 246–254, 2000.
- [55] D. Volpano, C. Irvine, and G. Smith. A sound type system for secure flow analysis. *J. Comput. Secur.*, 4(2-3):167–187, Jan. 1996.
- [56] D. Volpano and G. Smith. Verifying secrets and relative secrecy. In *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '00, pages 268–276, New York, NY, USA, 2000. ACM.
- [57] D. S. Wallach, J. A. Roskind, and E. W. Felten. Flexible, extensible Java security using digital signatures. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 38:59–74, Dec. 1996.