

Foundations and Trends® in Programming

Languages

From Fine- to Coarse-Grained Dynamic Information Flow Control and Back

A Tutorial on Dynamic Information Flow

Suggested Citation: Marco Vassena, Alejandro Russo, Deepak Garg, Vineet Rajani and Deian Stefan (2023), “From Fine- to Coarse-Grained Dynamic Information Flow Control and Back”, Foundations and Trends® in Programming Languages: Vol. 8, No. 1, pp 1–117. DOI: 10.1561/25000000046.

Marco Vassena

Utrecht University
m.vassena@uu.nl

Alejandro Russo

Chalmers University of Technology
russo@chalmers.se

Deepak Garg

Max Planck Institute for Software Systems
dg@mpi-sws.org

Vineet Rajani

University of Kent
v.rajani@kent.ac.uk

Deian Stefan

University of California, San Diego
deian@cs.ucsd.edu

This article may be used only for the purpose of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval.

now

the essence of knowledge

Boston — Delft

Contents

1	Introduction	3
2	Fine-Grained IFC Calculus	8
2.1	Dynamics	10
2.2	Security	16
2.3	Flow-Sensitive References	23
3	Coarse-Grained IFC Calculus	42
3.1	Dynamics	44
3.2	Security	50
3.3	Flow-Sensitive References	56
4	Verified Artifacts	67
4.1	Artifact Analysis	68
5	Fine- to Coarse-Grained Program Translation	72
5.1	Types and Values	72
5.2	Expressions	73
5.3	References	78
5.4	Correctness	79
5.5	Recovery of Non-Interference	80

6	Coarse- to Fine-Grained Program Translation	84
6.1	Types and Values	85
6.2	Expressions and Thunks	87
6.3	References	89
6.4	Cross-Language Equivalence Relation	90
6.5	Correctness	94
6.6	Recovery of Non-Interference	96
7	Related work	100
7.1	Relative Expressiveness of IFC Systems	100
7.2	Coarse-Grained Dynamic IFC	102
7.3	Fine-Grained Dynamic IFC	103
7.4	Label Introspection and Flow-Sensitive References	104
7.5	Proof Techniques	104
8	Conclusion	107
	References	108

From Fine- to Coarse-Grained Dynamic Information Flow Control and Back

Marco Vassena¹, Alejandro Russo², Deepak Garg³, Vineet Rajani⁴ and Deian Stefan⁵

¹*Utrecht University, The Netherlands; m.vassena@uu.nl*

²*Chalmers University of Technology, Sweden; russo@chalmers.se*

³*Max Planck Institute for Software Systems, Germany; dg@mpi-sws.org*

⁴*University of Kent, UK; v.rajani@kent.ac.uk*

⁵*University of California, San Diego, USA; deian@cs.ucsd.edu*

ABSTRACT

This tutorial provides a complete and homogeneous account of the latest advances in fine- and coarse-grained dynamic information-flow control (IFC) security. Since the 1970s, the programming language and the operating system communities proposed different IFC approaches. IFC operating systems track information flows in a coarse-grained fashion, at the granularity of a process. In contrast, traditional language-based approaches to IFC are fine-grained: they track information flows at the granularity of program variables. For decades, researchers believed coarse-grained IFC to be strictly less permissive than fine-grained IFC—coarse grained IFC systems seem inherently less precise because they track less information—and so granularity appeared to be a fundamental feature of IFC systems.

Marco Vassena, Alejandro Russo, Deepak Garg, Vineet Rajani and Deian Stefan (2023), “From Fine- to Coarse-Grained Dynamic Information Flow Control and Back”, Foundations and Trends® in Programming Languages: Vol. 8, No. 1, pp 1–117. DOI: 10.1561/25000000046.

©2023 M. Vassena *et al.*

We show that the granularity of the tracking system does *not* fundamentally restrict how precise or permissive dynamic IFC systems can be. To this end, we mechanize two mostly standard languages, one with a fine-grained dynamic IFC system and the other with a coarse-grained dynamic IFC system, and prove a semantics-preserving translation from each language to the other. In addition, we derive the standard security property of non-interference of each language from that of the other, via our verified translation.

These translations stand to have important implications on the usability of IFC approaches. The coarse- to fine-grained direction can be used to remove the label annotation burden that fine-grained systems impose on developers, while the fine- to coarse-grained translation shows that coarse-grained systems—which are easier to design and implement—can track information as precisely as fine-grained systems and provides an algorithm for automatically retrofitting legacy applications to run on existing coarse-grained systems.

1

Introduction

Dynamic *information-flow control* (IFC) is a principled approach to protecting the confidentiality and integrity of data in software systems. Conceptually, dynamic IFC systems are very simple—they associate *security* levels or *labels* with every bit of data in the system to subsequently track and restrict the flow of labeled data throughout the system, e.g., to enforce a security property such as *non-interference* (Goguen and Meseguer, 1982). In practice, dynamic IFC implementations are considerably more complex—the *granularity* of the tracking system alone has important implications for the usage of IFC technology. Indeed, until somewhat recently (Roy *et al.*, 2009; Stefan *et al.*, 2017), granularity was the main distinguishing factor between dynamic IFC operating systems and programming languages. Most IFC operating systems (e.g., Efstathopoulos *et al.*, 2005; Zeldovich *et al.*, 2006; Krohn *et al.*, 2007) are *coarse-grained*, i.e., they track and enforce information flow at the granularity of a process or thread. Conversely, most programming languages with dynamic IFC (e.g., Austin and Flanagan, 2009; Zdancewic, 2002; Hedin *et al.*, 2014; Hritcu *et al.*, 2013; Yang *et al.*, 2012) track the flow of information in a more *fine-grained* fashion, e.g., at the granularity of program variables and references.

Dynamic coarse-grained IFC systems in the style of LIO (Stefan *et al.*, 2017; Stefan *et al.*, 2011; Stefan *et al.*, 2012; Heule *et al.*, 2015; Buiras *et al.*, 2015; Vassena *et al.*, 2017) have several advantages over dynamic fine-grained IFC systems. Such coarse-grained systems are often easier to design and implement—they inherently track less information. For example, LIO protects against control-flow-based *implicit flows* by tracking information at a coarse-grained level—to branch on secrets, LIO programs must first taint the context where secrets are going to be observed. Finally, coarse-grained systems often require considerably fewer programmer annotations—unlike fine-grained ones. More specifically, developers often only need a single label-annotation to protect everything in the scope of a thread or process responsible to handle sensitive data.

Unfortunately, these advantages of coarse-grained systems give up on the many benefits of fine-grained ones. For instance, one main drawback of coarse-grained systems is that it requires developers to compartmentalize their application in order to avoid both false alarms and the *label creep* problem, i.e., wherein the program gets too “tainted” to do anything useful. To this end, coarse-grained systems often create special abstractions (e.g., event processes (Efstathopoulos *et al.*, 2005), gates (Zeldovich *et al.*, 2006), and security regions (Roy *et al.*, 2009)) that compensate for the conservative approximations of the coarse-grained tracking approach. Furthermore, fine-grained systems do not impose the burden of focusing on avoiding the label creep problem on developers. By tracking information at fine granularity, such systems are seemingly more flexible and do not suffer from false alarms and label creep issues (Austin and Flanagan, 2009) as coarse-grained systems do. Indeed, fine-grained systems such as JSFlow (Hedin *et al.*, 2014) can often be used to secure existing, legacy applications; they only require developers to properly annotate the application.

This tutorial removes the division between fine- and coarse-grained dynamic IFC systems and the belief that they are fundamentally different. In particular, we show that *dynamic* fine-grained and coarse-grained IFC are equally expressive. Our work is inspired by the recent work of Rajani *et al.* (2017) and Rajani and Garg (2018), who prove similar results for *static* fine-grained and coarse-grained IFC systems. Specifi-

cally, they establish a semantics- and type-preserving translation from a coarse-grained IFC type system to a fine-grained one and vice-versa. We complete the picture by showing a similar result for dynamic IFC systems that additionally allow *introspection on labels* at run-time. While label introspection is meaningless in a static IFC system, in a dynamic IFC system this feature is key to both writing practical applications and mitigating the label creep problem (Stefan *et al.*, 2017).

Using the Agda proof assistant (Norell, 2009; Bove *et al.*, 2009), we formalize a traditional fine-grained system (in the style of Austin and Flanagan, 2009) extended with label introspection primitives, as well as a coarse-grained system (in the style of Stefan *et al.*, 2017). We then define and formalize modular semantics-preserving translations between them. Our translations are macro-expressible in the sense of Felleisen (1991), i.e., they can be expressed as a pure source program rewriting.

We show that a translation from fine- to coarse-grained is possible when the coarse-grained system is equipped with a primitive that limits the scope of tainting (e.g., when reading sensitive data). In practice, this is not an imposing requirement since most coarse-grained systems rely on such primitives for compartmentalization. For example, Stefan *et al.* (2017) and Stefan *et al.* (2012), provide **toLabeled**(\cdot) blocks and threads for precisely this purpose. Dually, we show that the translation from coarse- to fine-grained is possible when the fine-grained system has a primitive **taint**(\cdot) that relaxes precision to keep the *program counter label* synchronized when translating a program to the coarse-grained language. While this primitive is largely necessary for us to establish the coarse- to fine-grained translation, extending existing fine-grained systems with it is both secure and trivial.

The implications of our results are multi-fold. The fine- to coarse-grained translation formally confirms an old OS-community hypothesis that it is possible to restructure a system into smaller compartments to address the label creep problem—indeed our translation is a (naive) algorithm for doing so. This translation also allows running legacy fine-grained IFC compatible applications atop coarse-grained systems like LIO. Dually, the coarse- to fine-grained translation allows developers building new applications in a fine-grained system to avoid the annotation burden of the fine-grained system by writing some of the

code in the coarse-grained system and compiling it automatically to the fine-grained system with our translation. The technical contributions of this monograph are:

- A pair of semantics-preserving translations between traditional dynamic fine-grained and coarse-grained IFC systems equipped with label introspection and flow-insensitive references (Theorems 5 and 7).
- Two different proofs of *termination-insensitive* non-interference (TINI) for each calculus: one is derived directly in the usual way (Theorems 1 and 3), while the other is recovered via our verified translation (Theorems 6 and 8).
- Mechanized Agda proofs of our results (~4,000 LOC).

This monograph is based on our conference paper (Vassena *et al.*, 2019) and extended with:

- A tutorial-style introduction to fine- and coarse-grained dynamic IFC, which (i) illustrates their specific features and (apparent) differences through examples, and (ii) supplements our proof artifacts with general explanations of the proof techniques used.
- *Flow-sensitive* references, a key feature for boosting the permissiveness of dynamic IFC systems (Austin and Flanagan, 2009). We extend both fine- and coarse-grained language with flow-sensitive references (Sections 2.3 and 3.3), adapt their security proofs (Theorems 2 and 4), and the verified translations to each other.
- A discussion and analysis of our extended proof artifact (~6,900 LOC)¹. Our analysis finds that the security proofs for fine-grained languages are between 43% and 74% longer than for coarse-grained languages. These empirical results suggests that it is indeed easier to reason about coarse-grained IFC languages than fine-grained languages.

¹The extended artifact is available at <https://hub.docker.com/r/marcovassena/granularity-ftpl> and supersedes the artifact archived with the conference paper.

This tutorial is organized as follows. Our dynamic fine- and coarse-grained IFC calculi are introduced in Sections 2 and 3, and then extended with flow-sensitive references in Sections 2.3 and 3.3, respectively. We prove the soundness guarantees (i.e., termination-insensitive non-interference) of the original languages (Sections 2.2 and 3.2), and of the extended languages (Sections 2.3.3 and 3.3.3). In Section 4, we discuss our mechanized proof artifacts and compare the security proofs of the two calculi, before and after the extension. In Section 5, we present the fine- to coarse-grained translation and a proof of non-interference for the fine-grained calculus recovered from non-interference of the other calculus through our verified translation. Section 6 presents similar results in the other direction. Related work is described in Section 7 and Section 8 concludes the tutorial.

2

Fine-Grained IFC Calculus

In order to compare in a rigorous way fine- and coarse-grained dynamic IFC techniques, we formally define the operational semantics of two λ -calculi that respectively perform fine- and coarse-grained IFC dynamically. Figure 2.1 shows the syntax of the dynamic fine-grained IFC calculus λ^{dFG} , which is inspired by Austin and Flanagan (2009) and extended with a standard (security unaware) type system $\Gamma \vdash e : \tau$ (omitted), sum and product data types and security labels $\ell \in \mathcal{L}$ that form a lattice $(\mathcal{L}, \sqsubseteq)$.¹ In order to capture flows of information precisely at run-time, the λ^{dFG} -calculus features *intrinsically labeled* values, written r^ℓ , meaning that raw value r has security level ℓ . Compound values, e.g., pairs and sums, carry labels to tag the security level of each component, for example a pair containing a secret and a public boolean would be written $(\mathbf{true}^H, \mathbf{false}^L)$.² Functional values are closures $(x.e, \theta)$, where x is the variable that binds the argument in the body of the function e and all other free variables are mapped to some labeled value in the environment θ . The λ^{dFG} -calculus features

¹The lattice is arbitrary and fixed. In examples we will often use the two point lattice $\{L, H\}$, which only disallows secret to public flow of information, i.e., $H \not\sqsubseteq L$.

²We define the boolean type $\mathbf{bool} = \mathbf{unit} + \mathbf{unit}$, boolean values as raw values, i.e., $\mathbf{true} = \mathbf{inl}(()^L)$, $\mathbf{false} = \mathbf{inr}(()^L)$ and $\mathbf{if } e \mathbf{ then } e_1 \mathbf{ else } e_2 = \mathbf{case } e \mathbf{ _}.e_1 \mathbf{ _}.e_2$.

Types: $\tau ::= \mathbf{unit} \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 + \tau_2 \mid \tau_1 \times \tau_2 \mid \mathcal{L} \mid \mathbf{Ref} \tau$
 Labels: $\ell, pc \in \mathcal{L}$
 Addresses: $n \in \mathbb{N}$
 Environments: $\theta \in \mathit{Var} \rightarrow \mathit{Value}$
 Raw Values: $r ::= () \mid (x.e, \theta) \mid \mathbf{inl}(v) \mid \mathbf{inr}(v) \mid (v_1, v_2) \mid \ell \mid n_\ell$
 Values: $v ::= r^\ell$
 Expression: $e ::= x \mid \lambda x.e \mid e_1 e_2 \mid () \mid \ell \mid (e_1, e_2) \mid \mathbf{fst}(e) \mid \mathbf{snd}(e)$
 $\mid \mathbf{inl}(e) \mid \mathbf{inr}(e) \mid \mathbf{case}(e, x.e_1, x.e_2)$
 $\mid \mathbf{getLabel} \mid \mathbf{labelOf}(e) \mid e_1 \sqsubseteq^? e_2 \mid \mathbf{taint}(e_1, e_2)$
 $\mid \mathbf{new}(e) \mid !e \mid e_1 := e_2 \mid \mathbf{labelOfRef}(e)$
 Type System: $\Gamma \vdash e : \tau$
 Configurations: $c ::= \langle \Sigma, e \rangle$
 Stores: $\Sigma \in (\ell : \mathit{Label}) \rightarrow \mathit{Memory} \ell$
 Memory ℓ : $M ::= [] \mid r : M$

Figure 2.1: Syntax of λ^{dFG} .

a labeled partitioned store, i.e., $\Sigma \in (\ell : \mathcal{L}) \rightarrow \mathit{Memory} \ell$, where $\mathit{Memory} \ell$ is the memory that contains values at security level ℓ . Each reference carries an additional label annotation that records the label of the memory it refers to—reference n_ℓ points to the n -th cell of the ℓ -labeled memory, i.e., $\Sigma(\ell)$. Notice that this label has nothing to do with the *intrinsic* label that decorates the reference itself. For example, a reference $(n_H)^L$ represents a secret reference in a public context, whereas $(n_L)^H$ represents a public reference in a secret context. Notice that there is no order invariant between those labels—in the latter case, the IFC runtime monitor prevents writing data to the reference to avoid *implicit flows*. A program can create, read and write a labeled reference via constructs $\mathbf{new}(e)$, $!e$ and $e_1 := e_2$ and inspect its subscripted label with the primitive $\mathbf{labelOfRef}(\cdot)$.

2.1 Dynamics

The operational semantics of λ^{dFG} includes a security monitor that propagates the label annotations of input values during program execution and assigns security labels to the result accordingly. The monitor prevents information leakage by stopping the execution of potentially leaky programs, which is reflected in the semantics by not providing reduction rules for the cases that may cause insecure information flow.³ In Figure 2.2, the relation $\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', v \rangle$ denotes the evaluation of program e with initial store Σ that terminates with labeled value v and final store Σ' . The environment θ stores the input values of the program and is extended with intermediate results during function application and case analysis. The subscript pc is the *program counter* label (Sabelfeld and Myers, 2006)—it is a label that represents the security level of the context in which the expression is evaluated. The semantics employs the program counter label to (i) propagate and assign labels to values computed by a program, and (ii) prevent implicit flow leaks that exploit the control flow and the store (explained below).

In particular, when a program produces a value, the monitor tags the raw value with the program counter label in order to record the security level of the context in which it was computed. For this reason all the introduction rules for ground and compound types ($[\text{UNIT}, \text{LABEL}, \text{FUN}, \text{INL}, \text{INR}, \text{PAIR}]$) assign security level pc to the result. Other than that, these rules are fairly standard—we simply note that rule $[\text{FUN}]$ creates a closure by capturing the current environment θ .

When the control flow of a program *depends* on some intermediate value, the program counter label is joined with the value’s label so that the label of the final result will be tainted with the result of the intermediate value. For instance, consider case analysis, i.e., **case** $e x.e_1 x.e_2$. Rules $[\text{CASE}_1]$ and $[\text{CASE}_2]$ evaluate the scrutinee e to a value (either $\mathbf{inl}(v)^\ell$ or $\mathbf{inr}(v)^\ell$), add the value to the environment, i.e., $\theta[x \mapsto v]$, and then evaluate the appropriate branch with a program counter label tainted with v ’s security label, i.e., $pc \sqcup \ell$. As a result, the monitor tracks data dependencies across control flow constructs through the

³In this work, we ignore leaks that exploit program termination and prove *termination insensitive* non-interference for λ^{dFG} (Theorem 1).

$$\begin{array}{c}
\text{(VAR)} \qquad \langle \Sigma, x \rangle \Downarrow_{pc}^{\theta} \langle \Sigma, \theta(x) \sqcup pc \rangle \qquad \text{(UNIT)} \qquad \langle \Sigma, () \rangle \Downarrow_{pc}^{\theta} \langle \Sigma, ()^{pc} \rangle \qquad \text{(LABEL)} \qquad \langle \Sigma, \ell \rangle \Downarrow_{pc}^{\theta} \langle \Sigma, \ell^{pc} \rangle \\
\\
\text{(FUN)} \qquad \langle \Sigma, \lambda x. e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma, (x. e, \theta)^{pc} \rangle \\
\\
\text{(APP)} \qquad \frac{\langle \Sigma, e_1 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', (x. e, \theta')^{\ell} \rangle \quad \langle \Sigma', e_2 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', v_2 \rangle \quad \langle \Sigma'', e \rangle \Downarrow_{pc \sqcup \ell}^{\theta' [x \mapsto v_2]} \langle \Sigma''', v \rangle}{\langle \Sigma, e_1 \ e_2 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma''', v \rangle} \\
\\
\text{(INL)} \qquad \frac{\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', v \rangle}{\langle \Sigma, \mathbf{inl}(e) \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', \mathbf{inl}(v)^{pc} \rangle} \qquad \text{(INR)} \qquad \frac{\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', v \rangle}{\langle \Sigma, \mathbf{inr}(e) \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', \mathbf{inr}(v)^{pc} \rangle} \\
\\
\text{(CASE}_1\text{)} \qquad \frac{\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', \mathbf{inl}(v_1)^{\ell} \rangle \quad \langle \Sigma', e_1 \rangle \Downarrow_{pc \sqcup \ell}^{\theta [x \mapsto v_1]} \langle \Sigma'', v \rangle}{\langle \Sigma, \mathbf{case}(e, x. e_1, x. e_2) \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', v \rangle} \\
\\
\text{(CASE}_2\text{)} \qquad \frac{\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', \mathbf{inr}(v_2)^{\ell} \rangle \quad \langle \Sigma', e_2 \rangle \Downarrow_{pc \sqcup \ell}^{\theta [x \mapsto v_2]} \langle \Sigma'', v \rangle}{\langle \Sigma, \mathbf{case}(e, x. e_1, x. e_2) \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', v \rangle} \\
\\
\text{(PAIR)} \qquad \frac{\langle \Sigma, e_1 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', v_1 \rangle \quad \langle \Sigma', e_2 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', v_2 \rangle}{\langle \Sigma, (e_1, e_2) \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', (v_1, v_2)^{pc} \rangle} \\
\\
\text{(FST)} \qquad \frac{\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', (v_1, v_2)^{\ell} \rangle}{\langle \Sigma, \mathbf{fst}(e) \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', v_1 \sqcup \ell \rangle} \qquad \text{(SND)} \qquad \frac{\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', (v_1, v_2)^{\ell} \rangle}{\langle \Sigma, \mathbf{snd}(e) \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', v_2 \sqcup \ell \rangle}
\end{array}$$

Figure 2.2: Big-step semantics for λ^{dFG} (part I).

$$\begin{array}{c}
\text{(LABELOF)} \\
\frac{\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', r^{\ell} \rangle}{\langle \Sigma, \mathbf{labelOf}(e) \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', \ell^{\ell} \rangle}
\end{array}
\qquad
\begin{array}{c}
\text{(GETLABEL)} \\
\langle \Sigma, \mathbf{getLabel} \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', pc^{pc} \rangle
\end{array}$$

$$\begin{array}{c}
\text{(\(\sqsubseteq\)-T)} \\
\frac{\langle \Sigma, e_1 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', \ell_1^{\ell_1} \rangle \quad \langle \Sigma', e_2 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', \ell_2^{\ell_2} \rangle \quad \ell_1 \sqsubseteq \ell_2}{\langle \Sigma, e_1 \sqsubseteq^? e_2 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', \mathbf{inl}((\)^{pc})^{\ell_1 \sqcup \ell_2} \rangle}
\end{array}$$

$$\begin{array}{c}
\text{(\(\sqsubseteq\)-F)} \\
\frac{\langle \Sigma, e_1 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', \ell_1^{\ell_1} \rangle \quad \langle \Sigma', e_2 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', \ell_2^{\ell_2} \rangle \quad \ell_1 \not\sqsubseteq \ell_2}{\langle \Sigma, e_1 \sqsubseteq^? e_2 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', \mathbf{inr}((\)^{pc})^{\ell_1 \sqcup \ell_2} \rangle}
\end{array}$$

$$\begin{array}{c}
\text{(TAINT)} \\
\frac{\langle \Sigma, e_1 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', \ell^{\ell'} \rangle \quad \ell' \sqsubseteq \ell \quad \langle \Sigma', e_2 \rangle \Downarrow_{\ell}^{\theta} \langle \Sigma'', v \rangle}{\langle \Sigma, \mathbf{taint}(e_1, e_2) \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', v \rangle}
\end{array}$$

Figure 2.3: Big-step semantics for λ^{dFG} (part II).

label of the result. Function application follows the same principle. In rule [APP], since the first premise evaluates the function to some closure $(x.e, \theta')$ at security level ℓ , the third premise evaluates the body with program counter label raised to $pc \sqcup \ell$. The evaluation strategy is call-by-value: it evaluates the argument before the body in the second premise and binds the corresponding variable to its value in the environment of the closure, i.e., $\theta'[x \mapsto v_2]$. Notice that the security level of the argument is irrelevant at this stage and that this is beneficial to not over-tainting the result: if the function never uses its argument then the label of the result depends exclusively on the program counter label, e.g., $(\lambda x.()) y \Downarrow_L^{[y \mapsto \mathbf{true}^H]} ()^L$. The elimination rules for variables and pairs taint the label of the corresponding value with the program counter label for security reasons. In rules [VAR, FST, SND] the notation, $v \sqcup \ell'$ upgrades the label of v with ℓ' —it is a shorthand for $r^{\ell \sqcup \ell'}$ with $v = r^{\ell}$. Intuitively, public values must be considered secret when the program counter is secret, for example $x \Downarrow_H^{[x \mapsto ()^L]} ()^H$.

2.1.1 Label Introspection

The λ^{dFG} -calculus features primitives for label introspection, namely **getLabel**, **labelOf**(\cdot) and $\sqsubseteq^?$ —see Figure 2.3. These operations allow to respectively retrieve the current program counter label, obtain the label annotations of values, and compare two labels (inspecting labels at run-time is useful for controlling and mitigating the label creep problem).

Enabling label introspection raises the question of what label should be assigned to the label itself (in λ^{dFG} every value, including all label values, must be annotated with a label). As a matter of fact, labels can be used to encode secret information and thus careless label introspection may open the doors to information leakage (Stefan *et al.*, 2017). Notice that in λ^{dFG} , the label annotation on the result is computed by the semantics together with the result and thus it is as sensitive as the result itself (the label annotation on a value depends on the sensitivity of all values affecting the *control-flow* of the program up to the point where the result is computed). This motivates the design choice to protect each projected label with the label itself, i.e., ℓ^ℓ and pc^{pc} in rules [GETLABEL] and [LABELOF] in Figure 2.3. We remark that this choice is consistent with previous work on coarse-grained IFC languages (Buiras *et al.*, 2014; Stefan *et al.*, 2017), but novel in the context of fine grained IFC.

Finally, primitive **taint**(e_1, e_2) temporarily raises the program counter label to the label given by the first argument in order to evaluate the second argument. The fine-to-coarse translation in Section 5 uses **taint**(\cdot) to loosen the precision of λ^{dFG} in a controlled way and match the *coarse* approximation of our coarse-grained IFC calculus (λ^{dCG}) by upgrading the labels of intermediate values systematically. In rule [TAINT], the constraint $\ell' \sqsubseteq \ell$ ensures that the label of the nested context ℓ is at least as sensitive as the program counter label pc . In particular, this constraint ensures that the operational semantics have Property 1 (“the label of the result of any λ^{dFG} program is always at least as sensitive as the program counter label”) even with rule [TAINT].

Property 1. If $\langle \Sigma, e \rangle \Downarrow_{pc}^\theta \langle \Sigma', r^\ell \rangle$ then $pc \sqsubseteq \ell$.

$$\begin{array}{c}
\text{(NEW)} \\
\frac{\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', r^{\ell} \rangle \quad n = |\Sigma'(\ell)|}{\langle \Sigma, \mathbf{new}(e) \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'[\ell \mapsto \Sigma'(\ell)[n \mapsto r]], (n_{\ell})^{pc} \rangle} \\
\\
\text{(READ)} \\
\frac{\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', n_{\ell}^{\ell'} \rangle \quad \Sigma'(\ell)[n] = r}{\langle \Sigma, !e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', r^{\ell \sqcup \ell'} \rangle} \\
\\
\text{(WRITE)} \\
\frac{\langle \Sigma, e_1 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', n_{\ell}^{\ell_1} \rangle \quad \ell_1 \sqsubseteq \ell \quad \langle \Sigma', e_2 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', r^{\ell_2} \rangle \quad \ell_2 \sqsubseteq \ell}{\langle \Sigma, e_1 := e_2 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma''[\ell \mapsto \Sigma''(\ell)[n \mapsto r]],^{pc} \rangle} \\
\\
\text{(LABELOFREF)} \\
\frac{\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', n_{\ell}^{\ell'} \rangle}{\langle \Sigma, \mathbf{labelOfRef}(e) \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', \ell^{\ell \sqcup \ell'} \rangle}
\end{array}$$

Figure 2.4: Big-step semantics for λ^{dFG} (references).

Proof. By induction on the given evaluation derivation. \square

2.1.2 References

We now extend the semantics presented earlier with primitives that inspect, access and modify the labeled store via labeled references. See Figure 2.4. Rule [NEW] creates a reference n_{ℓ} , labeled with the security level of the initial content, i.e., label ℓ , in the ℓ -labeled memory $\Sigma(\ell)$ and updates the memory store accordingly.⁴ Since the security level of the reference is as sensitive as the content, which is at least as sensitive as the program counter label by Property 1 ($pc \sqsubseteq \ell$) this operation does not leak information via *implicit flows*. When reading the content of reference n_{ℓ} at security level ℓ' , rule [READ] retrieves the corresponding raw value from the n -th cell of the ℓ -labeled memory, i.e., $\Sigma'(\ell)[n] = r$ and upgrades its label to $\ell \sqcup \ell'$ since the decision to read from that

⁴ $|M|$ denotes the length of memory M —memory indices start at 0.

particular reference depends on information at security level ℓ' . When writing to a reference the monitor performs security checks to avoid leaks via explicit or implicit flows. Rule [WRITE] achieves this by evaluating the reference, i.e., $(n_\ell)^{\ell_1}$ and replacing its content with the value of the second argument, i.e., r^{ℓ_2} , under the conditions that the decision of “which” reference to update does not depend on data more sensitive than the reference itself, i.e., $\ell_1 \sqsubseteq \ell$ (not checking this would leak via an *implicit flow*),⁵ and that the new content is no more sensitive than the reference itself, i.e., $\ell_2 \sqsubseteq \ell$ (not checking this would leak sensitive information to a less sensitive reference via an *explicit flow*). Lastly, rule [LABELOFREF] retrieves the label of the reference and protects it with the label itself (as explained before) and taints it with the security level of the reference, i.e., $\ell^\ell \sqcup \ell'$ to avoid leaks. Intuitively, the label of the reference, i.e., ℓ , depends also on data at security level ℓ' as seen in the premise.

Other Extensions. We consider λ^{dFG} equipped with references as sufficient foundation to study the relationship between fine-grained and coarse-grained IFC. We remark that extending it with other side-effects such as file operations, or other IO-operations would not change our claims in Section 5 and 6. The main reason for this is that, typically, handling such effects would be done at the same granularity in both IFC enforcements. For instance, when adding file operations, both fine- (e.g., Broberg *et al.*, 2013) and coarse-grained (e.g., Russo *et al.*, 2009; Stefan *et al.*, 2011; Efstathopoulos *et al.*, 2005; Krohn *et al.*, 2007) enforcements are likely to assign a single *flow-insensitive* (i.e., immutable) label to each file in order to denote the sensitivity of its content. Then, those features could be handled *flow-insensitively* in both systems (e.g., Pottier and Simonet, 2003; Myers *et al.*, 2006; Stefan *et al.*, 2011; Vassena and Russo, 2016), in a manner similar to what we have just shown for references in λ^{dFG} .

Importantly, fine- and coarse-grained IFC are equally expressive also when extended with *flow-sensitive* (i.e., mutable) labels. Unlike

⁵Notice that $pc \sqsubseteq \ell_1$ by Property 1, thus $pc \sqsubseteq \ell_1 \sqsubseteq \ell$ by transitivity. An *implicit flow* would occur if a reference is updated in a *high branch*, i.e., depending on the secret, e.g., `let x = new(0) in if secret then x := 1 else ()`.

flow-insensitive labels, these labels can *change* during program execution to reflect the current sensitivity of the content of various resources (e.g., references, files etc.). In order to show that fine- and coarse-grained IFC equally support flow-sensitive features, we follow the same approach described above. In particular, we add *flow-sensitive* references (Austin and Flanagan, 2009) to λ^{dFG} (Section 2.3) and λ^{dCG} (Section 3.3) and then complete our pair of verified semantics- and security-preserving translations from one language to the other (Sections 5 and 6).

2.2 Security

We now prove that λ^{dFG} is secure, i.e., it satisfies *termination insensitive non-interference* (TINI) (Goguen and Meseguer, 1982; Volpano and Smith, 1997). Intuitively, the security condition says that no terminating λ^{dFG} program leaks information, i.e., changing secret inputs does not produce any publicly visible effect. The proof technique is standard and based on the notion of L -equivalence, written $v_1 \approx_L v_2$, which relates values (and similarly raw values, environments, stores and configurations) that are indistinguishable for an attacker at security level L . For clarity we use the 2-points lattice, assume that secret data is labeled with H and that the attacker can only observe data at security level L . Our mechanized proofs are parametric in the lattice and in the security level of the attacker.

2.2.1 L -Equivalence

L -equivalence for values and raw-values is defined formally by mutual induction in Figure 2.5. Rule [VALUE $_L$] relates observable values, i.e., raw values labeled below the security level of the attacker. These values have the *same* observable label ($\ell \sqsubseteq L$) and related raw values, i.e., $r_1 \approx_L r_2$. Rule [VALUE $_H$] relates non-observable values, which may have different labels not below the attacker level, i.e., $\ell_1 \not\sqsubseteq L$ and $\ell_2 \not\sqsubseteq L$. In this case, the raw values can be arbitrary. Raw values are L -equivalent when they consist of the same ground value (i.e., rules [UNIT] and [LABEL]), or are homomorphically related for compound values. For example, for the sum type the relation requires that both values are

either a left or a right injection through rules [INL] and [INR]. Closures are related by rule [CLOSURE], if they contain the *same* function (up to α -renaming)⁶ and L -equivalent environments, i.e., the environments are L -equivalent pointwise. Formally, $\theta_1 \approx_L \theta_2$ iff $\text{dom}(\theta_1) \equiv \text{dom}(\theta_2)$ and $\forall x. \theta_1(x) \approx_L \theta_2(x)$.

$$\begin{array}{c}
\text{(VALUE}_L\text{)} \\
\frac{\ell \sqsubseteq L \quad r_1 \approx_L r_2}{r_1^\ell \approx_L r_2^\ell}
\end{array}
\quad
\begin{array}{c}
\text{(VALUE}_H\text{)} \\
\frac{\ell_1 \not\sqsubseteq L \quad \ell_2 \not\sqsubseteq L}{r_1^{\ell_1} \approx_L r_2^{\ell_2}}
\end{array}
\quad
\begin{array}{c}
\text{(UNIT)} \\
() \approx_L ()
\end{array}
\quad
\begin{array}{c}
\text{(LABEL)} \\
\ell \approx_L \ell
\end{array}$$

$$\begin{array}{c}
\text{(CLOSURE)} \\
\frac{e_1 \equiv_\alpha e_2 \quad \theta_1 \approx_L \theta_2}{(e_1, \theta_1) \approx_L (e_2, \theta_2)}
\end{array}
\quad
\begin{array}{c}
\text{(INL)} \\
\frac{v_1 \approx_L v_2}{\mathbf{inl}(v_1) \approx_L \mathbf{inl}(v_2)}
\end{array}
\quad
\begin{array}{c}
\text{(INR)} \\
\frac{v_1 \approx_L v_2}{\mathbf{inr}(v_1) \approx_L \mathbf{inr}(v_2)}
\end{array}$$

$$\begin{array}{c}
\text{(PAIR)} \\
\frac{v_1 \approx_L v'_1 \quad v_2 \approx_L v'_2}{(v_1, v_2) \approx_L (v'_1, v'_2)}
\end{array}
\quad
\begin{array}{c}
\text{(REF}_L\text{)} \\
\frac{\ell \sqsubseteq L}{n_\ell \approx_L n_\ell}
\end{array}
\quad
\begin{array}{c}
\text{(REF}_H\text{)} \\
\frac{\ell_1 \not\sqsubseteq L \quad \ell_2 \not\sqsubseteq L}{n_{1\ell_1} \approx_L n_{2\ell_2}}
\end{array}$$

Figure 2.5: L -equivalence for λ^{dFG} values and raw values.

We define L -equivalence for stores pointwise, i.e., $\Sigma_1 \approx_L \Sigma_2$ iff for all labels $\ell \in \mathcal{L}$, $\Sigma_1(\ell) \approx_L \Sigma_2(\ell)$. Memory L -equivalence relates arbitrary ℓ -labeled memories if $\ell \not\sqsubseteq L$, and pointwise otherwise, i.e., $M_1 \approx_L M_2$ iff M_1 and M_2 are memories labeled with $\ell \sqsubseteq L$, $|M_1| = |M_2|$ and for all $n \in \{0 \dots |M_1| - 1\}$, $M_1[n] \approx_L M_2[n]$. Similarly, L -equivalence relates any two secret references through rule [REF_H], but requires the same label and address for public references in rule [REF_L]. We naturally lift L -equivalence to initial configurations, i.e., $c_1 \approx_L c_2$ iff $c_1 = \langle \Sigma_1, e_1 \rangle$, $c_2 = \langle \Sigma_2, e_2 \rangle$, $\Sigma_1 \approx_L \Sigma_2$ and $e_1 \equiv_\alpha e_2$, and final configurations, i.e., $c'_1 \approx_L c'_2$ iff $c'_1 = \langle \Sigma'_1, v_1 \rangle$, $c'_2 = \langle \Sigma'_2, v_2 \rangle$ and $\Sigma'_1 \approx_L \Sigma'_2$ and $v_1 \approx_L v_2$.

The L -equivalence relation defined above is *reflexive*, *symmetric*, and *transitive*.

⁶Symbol \equiv_α denotes α -equivalence. In our mechanized proofs we use De Bruijn indexes and syntactic equivalence.

$$\begin{array}{ccc}
 \Sigma_1 & \xrightarrow{\approx_L} & \Sigma'_1 \\
 \approx_L \downarrow & & \downarrow \approx_L \\
 \Sigma_2 & \xrightarrow{\approx_L} & \Sigma'_2
 \end{array}$$

Figure 2.6: Square commutative diagram for stores (Lemma 2.1).

Property 2. Let x, y, z range over labeled values, raw values, environments, labeled memories, stores, and configurations:

1. **Reflexivity.** For all x , $x \approx_L x$.
2. **Symmetry.** For all x and y , if $x \approx_L y$, then $y \approx_L x$.
3. **Transitivity.** For all x, y, z , if $x \approx_L y$ and $y \approx_L z$, then $x \approx_L z$.

These properties simplify the security analysis of λ^{dFG} : they let us reason about L -equivalent terms using *commutative diagrams*. For example, consider the *Square Commutative Diagram for Stores* outlined in Figure 2.6. In the diagram, the arrows connect L -equivalent stores (e.g., the arrow from Σ_1 to Σ'_1 indicates that $\Sigma_1 \approx_L \Sigma'_1$). The diagram provides a visual representation of Lemma 2.1: solid arrows represent the assumptions of the lemma and the dashed arrow represents the conclusion. To prove the lemma, we have to show that the diagram *commutes*, i.e., we need to construct a path from Σ'_1 to Σ'_2 using the solid arrows. Thanks to Property 2, we can derive additional arrows to construct this path. For example, we can reverse arrows using *symmetry* (e.g., $\Sigma'_1 \approx_L \Sigma_1$ from $\Sigma_1 \approx_L \Sigma'_1$) and *transitivity* let us compose consecutive arrows (e.g., $\Sigma_1 \approx_L \Sigma'_2$ from $\Sigma_1 \approx_L \Sigma_2$ and $\Sigma_2 \approx_L \Sigma'_2$).

Lemma 2.1 (Square Commutative Diagram for Stores). If $\Sigma_1 \approx_L \Sigma'_1$, $\Sigma_1 \approx_L \Sigma_2$, $\Sigma_2 \approx_L \Sigma'_2$, then $\Sigma'_1 \approx_L \Sigma'_2$.

Proof. We show that the square diagram commutes using *symmetry* (Property 2.2) and *transitivity* (Property 2.3) to draw the red arrows in Figure 2.7.

□




Figure 2.7: Proof of Lemma 2.1. The red arrows represent L -equivalent relations derived via symmetricity and transitivity.

2.2.2 Termination-Insensitive Non-Interference

The security monitor of λ^{dFG} enforces termination-insensitive non-interference. Intuitively, this property guarantees that *terminating* programs do not leak secret data into public values and memories of the store. More formally, a program satisfies non-interference if, given indistinguishable inputs (initial stores and environments), then it produces outputs (final stores and values) that are also indistinguishable to the attacker. We prove this result through two key lemmas: *store confinement* and *L -equivalence preservation*. In the following, we give a high-level overview of these lemmas and their proof, focusing on the general proof technique and illustrative cases. We refer to our mechanized proof scripts for complete proofs.

Store Confinement. At a high-level, store confinement ensures that programs cannot leak secret data *implicitly* through observable side-effects in the labeled store. Intuitively, the side-effects of programs running in secret contexts must be *confined* to secret memories in the labeled store to enforce security—programs that do otherwise may leak and are therefore conservatively aborted by the security monitor. This lemma holds for λ^{dFG} precisely because the constraints in rules [NEW] and [WRITE] only allow programs to write memories labeled above the program counter label. In particular, these constraints prevent programs running in secret contexts from writing public memories, which remain unchanged and thus indistinguishable to the attacker.

Lemma 2.2 (Store Confinement). For all configurations $c = \langle \Sigma, e \rangle$, $c' = \langle \Sigma', v \rangle$, program counter labels $pc \not\sqsubseteq L$, if $c \Downarrow_{pc}^{\theta} c'$, then $\Sigma \approx_L \Sigma'$.

Proof. The proof is by induction on the big-step reduction, using *reflexivity* (Property 2.1) in the base cases and *transitivity* (Property 2.3) in the inductive cases. In the inductive cases (e.g., [CASE₁]), we observe that the program counter label of the nested computations is always *at least as sensitive* as the initial program counter pc , and therefore above the attacker’s label L , i.e., if $pc \not\sqsubseteq L$, then $pc \sqcup \ell \not\sqsubseteq L$ for any label ℓ . The interesting cases are those that change the store, i.e., cases [NEW] and [WRITE], where we use Property 1 to show that these rules can modify only secret memories. For example, in case [NEW], the program creates a reference in a secret context ($pc \not\sqsubseteq L$). First, the program computes a value labeled ℓ above the attacker’s label L , i.e., $\ell \not\sqsubseteq L$ by Property 1, and then allocates it in the corresponding secret memory also labeled ℓ . Since $\ell \not\sqsubseteq L$, the original and the extended memory are indistinguishable by the attacker and so are the initial and final stores. \square

L -equivalence Preservation. Termination-insensitive non-interference ensures that programs that receive L -equivalent inputs produce outputs that are also L -equivalent, i.e., terminating programs *preserve* L -equivalence. Importantly, programs must preserve L -equivalence regardless of the sensitivity of the context in which they are executed. Therefore, we consider L -equivalence preservation in public and secret contexts *separately*. Then, we combine these individual results to prove L -equivalence preservation in arbitrary contexts, i.e., termination-insensitive non-interference. More precisely, we prove two preservation lemmas: the first relates executions in secret contexts ($pc \not\sqsubseteq L$) involving *arbitrary expressions*, while the other relates executions of the *same expression* in public contexts ($pc \sqsubseteq L$).

The first lemma ensures that programs cannot leak secret data implicitly through the *program control-flow*. For example, consider the program **if** s **then** e_1 **else** e_2 , which branches on a secret boolean s . Depending on the value of s , this program evaluates either expression e_1 or e_2 , which may reveal the value of the secret through secret-dependent store updates (e.g., **if** s **then** $p := \mathbf{true}$ **else** $()$) for some public reference p) or results (e.g., **if** s **then** **true** **else** **false**). This lemma ensures that even these programs cannot leak secrets through

the final result or observable changes to the stores. Formally, we have to show that programs preserve L -equivalence in secret contexts even if they evaluate different expressions, i.e., if $\Sigma_1 \approx_L \Sigma_2$, $\langle \Sigma_1, e_1 \rangle \Downarrow_H^{\theta_1} c_1$ and $\langle \Sigma_2, e_2 \rangle \Downarrow_H^{\theta_2} c_2$, then $c_1 \approx_L c_2$.

How should we prove this lemma? At first, we might try to prove it *directly* by induction on the reduction steps. However, this approach is not practical: since these executions involve *arbitrary* expressions, we would have to reason about a large number of completely unrelated reduction steps! Instead, we observe that these programs are executed in secret contexts ($pc = H$), so they are restricted by the security monitor to avoid leaks. From these restrictions, we establish two program *invariants* to show that the final stores and values are L -equivalent: in secret contexts, programs can (i) modify *only* secret memories (*Store Confinement*), and (ii) produce results labeled secret (Property 1). Notice that these invariants hold for individual executions: we still need to combine them to relate the final configurations of the two executions. To do that, we reason separately about the final stores and values. The final stores are related via a *Square Commutative Diagram* (Figure 2.6), while the secret results are trivially L -equivalent by rule [VALUE $_H$] (Figure 2.5).

Lemma 2.3 (L -Equivalence Preservation in Secret Contexts). For all program counter labels $pc_1 \not\sqsubseteq L$ and $pc_2 \not\sqsubseteq L$, and *arbitrary* expressions e_1 and e_2 , if $\Sigma_1 \approx_L \Sigma_2$, $\langle \Sigma_1, e_1 \rangle \Downarrow_{pc_1}^{\theta_1} c_1$, and $\langle \Sigma_2, e_2 \rangle \Downarrow_{pc_2}^{\theta_2} c_2$, then $c_1 \approx_L c_2$.

Proof. Assume $pc_1 \not\sqsubseteq L$, $pc_2 \not\sqsubseteq L$, $\Sigma_1 \approx_L \Sigma_2$, and let the final configurations be $c_1 = \langle \Sigma'_1, v_1 \rangle$ and $c_2 = \langle \Sigma'_2, v_2 \rangle$. First, we apply *Store Confinement* (Lemma 2.2) to $\langle \Sigma_1, e_1 \rangle \Downarrow_{pc_1}^{\theta_1} \langle \Sigma'_1, v_1 \rangle$ and $\langle \Sigma_2, e_2 \rangle \Downarrow_{pc_2}^{\theta_2} \langle \Sigma'_2, v_2 \rangle$ and obtain $\Sigma_1 \approx_L \Sigma'_1$ and $\Sigma_2 \approx_L \Sigma'_2$, respectively. Then, we construct the *Square Commutative Diagram for Stores* (Lemma 2.1) using $\Sigma_1 \approx_L \Sigma_2$, $\Sigma'_1 \approx_L \Sigma_1$, $\Sigma_2 \approx_L \Sigma'_2$, and obtain $\Sigma'_1 \approx_L \Sigma'_2$. To show that the values $v_1 = r_1^{\ell_1}$ and $v_2 = r_2^{\ell_2}$ are L -equivalent, it suffices to show that they are labeled *secret*. Since $pc_1 \not\sqsubseteq L$ and $pc_2 \not\sqsubseteq L$ by assumption, $pc_1 \sqsubseteq \ell_1$ and $pc_2 \sqsubseteq \ell_2$ by Property 1, we have $\ell_1 \not\sqsubseteq L$ and $\ell_2 \not\sqsubseteq L$, and thus $v_1 = r_1^{\ell_1} \approx_L r_2^{\ell_2} = v_2$ by rule [VALUE $_H$]. Since $\Sigma'_1 \approx_L \Sigma'_2$ and $v_1 \approx_L v_2$, we have $c_1 = \langle \Sigma'_1, v_1 \rangle \approx_L \langle \Sigma'_2, v_2 \rangle = c_2$, as desired. \square

We now turn our attention to L -equivalence preservation in *public contexts*. Unlike L -equivalence preservation in secret contexts, we must inspect the program executions and examine their specific reduction steps to prove this lemma. This is because these executions occur in a public context ($pc \sqsubseteq L$), therefore their side-effects and final results may be observable by the attacker and so we have to verify that no leaks occur in each case. Luckily, this task is simplified by the fact that the initial configurations are L -equivalent and thus contain α -equivalent expressions. Intuitively, this means that the two programs are *synchronized* and proceed in *lock-step*. In other words, their reductions are almost identical, i.e., the programs perform the same operation on L -equivalent inputs and so produce L -equivalent outputs. This makes the proof of this lemma straightforward in most cases. However, reductions that involve control-flow expressions (e.g., $\mathbf{case}(e, x.e_1, x.e_2)$) are more complicated. In general, these programs may follow different paths and evaluate different expressions: how can they preserve L -equivalence? Intuitively, since the programs considered in this lemma are initially synchronized, they can start following different paths *only* after branching on secret data. As a result, their program counter label gets tainted and the programs enter a *secret context*, in which public side-effects are not allowed and results are guaranteed to be labeled secret, as proved in the previous lemma.

Lemma 2.4 (*L -Equivalence Preservation in Public Contexts*). For all program counter labels $pc \sqsubseteq L$, if $c_1 \approx_L c_2$, $\theta_1 \approx_L \theta_2$, $c_1 \Downarrow_{pc}^{\theta_1} c'_1$, and $c_2 \Downarrow_{pc}^{\theta_2} c'_2$, then $c'_1 \approx_L c'_2$.

Proof. Let $c_1 = \langle \Sigma_1, e_1 \rangle$ and $c_2 = \langle \Sigma_2, e_2 \rangle$ and proceed by induction on the big-step reductions. Since $c_1 \approx_L c_2$, we know that the initial stores are L -equivalent, i.e., $\Sigma_1 \approx_L \Sigma_2$, and the big-step reductions evaluate α -equivalent expressions, i.e., $e_1 \equiv_\alpha e_2$. As a result, the expressions step following the same rule in most cases and thus produce L -equivalent values and stores.

The interesting cases are those that influence the control-flow of the program, where the two executions may deviate from each other, i.e., case [CASE₁], [CASE₂], and [APP]. In these cases, the security label ℓ of the scrutinee determines whether the two executions stay in a public

context ($\ell \sqsubseteq L$) and remain *synchronized* on the same path, or not ($\ell \not\sqsubseteq L$). For example, consider expression $\mathbf{case}(e, x.e_1, x.e_2)$, which steps through rule [CASE₁] in the first reduction and either through rule [CASE₁] or [CASE₂] in the second reduction (the opposite cases are symmetric). If both rules step through rule [CASE₁], then the scrutinees are both left injections, i.e., $\mathbf{inl}(v_1)^{\ell_1} \approx_L \mathbf{inl}(v_2)^{\ell_2}$ by induction hypothesis. Then, we perform case analysis on the L -equivalence judgment and have two sub-cases: [VALUE_L] and [VALUE_H]. In the first case, both scrutinees are public, i.e., $\ell_1 = \ell_2 \sqsubseteq L$ and $v_1 \approx_L v_2$ by rule [VALUE_L], and the proof follows by induction. In the other case, both scrutinees are secret, i.e., $\ell_1 \not\sqsubseteq L$ and $\ell_2 \not\sqsubseteq L$ by rule [VALUE_H], and the program enters a secret context and we apply *L-equivalence Preservation in Secret Contexts* (Lemma 2.3). Finally, if $\mathbf{case}(e, x.e_1, x.e_2)$ steps through different rules (e.g., [CASE₁] and [CASE₂]), then the scrutinees *must* be secret and thus we also apply *L-equivalence Preservation in Secret Contexts*. In particular, it is impossible for L -equivalent public scrutinees to have different injection. To see that, assume $\mathbf{inl}(v_1)^{\ell_1} \approx_L \mathbf{inr}(v_2)^{\ell_2}$ and $\ell_1 = \ell_2 \sqsubseteq L$ by rule [VALUE_L]. Then, we also have a proof that the *raw values* are L -equivalent, i.e., $\mathbf{inl}(v_1) \approx_L \mathbf{inr}(v_2)$. But this is impossible: raw values of sum type can be related only if they are the same injection (i.e., rules [INL] and [INR] in Figure 2.5). \square

Finally, we combine the L -equivalence preservation lemmas from above and prove *termination-insensitive non-interference* (TINI) for λ^{dFG} .

Theorem 1 (λ^{dFG} -TINI). If $c_1 \Downarrow_{pc}^{\theta_1} c'_1$, $c_2 \Downarrow_{pc}^{\theta_2} c'_2$, $\theta_1 \approx_L \theta_2$ and $c_1 \approx_L c_2$ then $c'_1 \approx_L c'_2$.

Proof. By case analysis over the program counter label pc . If $pc \sqsubseteq L$, we apply *L-equivalence Preservation in Public Contexts* (Lemma 2.4). If $pc \not\sqsubseteq L$, we apply *L-equivalence Preservation in Secret Contexts* (Lemma 2.3). \square

2.3 Flow-Sensitive References

This section extends λ^{dFG} with *flow-sensitive* references (Austin and Flanagan, 2009), an important feature to boost the permissiveness of

IFC systems. These references differ slightly from the labeled references presented in Section 2.1.2, which are instead *flow-insensitive*. The key difference between them lies in the way the IFC system treats their label. In particular, the label of flow-insensitive references is *immutable*, i.e., when a program creates a reference, the IFC monitor assigns it a label, which remains fixed throughout the execution of the program. In contrast, the label of flow-sensitive references is *mutable*. Intuitively, the label of these references can change to reflect the sensitivity of the data that they currently store. This change is completely transparent to the program: when the program writes some data to a flow-sensitive reference, the IFC monitor simply replaces the label of the reference with the label of the new content. However, if added naively, mutable labels open a new channel for implicitly leaking data, therefore, the IFC monitor changes the labels of flow-sensitive references only as long as this operation does not leak information.

In this section, we formally discuss the differences of flow-sensitive references and their subtleties for security in the context of λ^{dFG} . First, we show that flow-sensitive references are more permissive than flow-insensitive references with an example.

Example 2.1. Consider the following program, which creates a new reference, initially containing some public data p , overwrites it with secret data s , and finally reads it back from the reference.

$$\begin{aligned} &\mathbf{let} \ r = \mathbf{new}(p) \ \mathbf{in} \\ &\quad r := s; \\ &\quad !r \end{aligned}$$

The execution of the λ^{dFG} program above with program counter label public, i.e., $pc = L$, environment $\theta = [p \mapsto \mathbf{true}^L, s \mapsto \mathbf{false}^H]$, and empty store $\Sigma = \lambda\ell.[\]$ is aborted by the IFC monitor with flow-insensitive references from Section 2.1.2. Rule [NEW] (Figure 2.4) assigns the *fixed* public label L to the reference, which then causes rule [WRITE] to fail, as the program tries to write secret data (i.e., \mathbf{false}^H) into a public reference (i.e., 0_L). In particular, the last premise of rule [WRITE] does not hold ($H \not\sqsubseteq L$) and thus the program gets stuck. Had the

monitor not aborted the execution, the program would have leaked the secret value of s . Specifically, the program would have written the *raw* value of s , i.e., **false**, into the public L -labeled memory, from where it would be extracted by the last instruction through rule [READ] and labeled with L , i.e., **false** ^{L} , causing the leak. In contrast, the IFC monitor presented in this section and extended with flow-sensitive references does *not* abort the program. The extended monitor allows the program to write secret data into the public reference, but after doing so, it *upgrades* the label of the reference from L to H to indicate that it now contains secret data. As a result, when the program reads back the reference in the last instruction, the result gets tainted with H , i.e., **false** ^{H} , thus eliminating the leak.

In the following, we extend λ^{dFG} with flow-sensitive references (Section 2.3.1), define their operational semantics (Section 2.3.2), and finally reestablish non-interference (Section 2.3.3).

2.3.1 Syntax

We introduce the syntax of λ^{dFG} extended with flow-sensitive references in Figure 2.8a, where we omit the constructs and the syntactic categories that remain unchanged. First, we annotate the type of references with flow-sensitivity tags, e.g., **Ref** $s \tau$, where the tag s allows programs to distinguish between flow-insensitive ($s = I$) and flow-sensitive ($s = S$) references. We do not introduce new constructs to create, write, and read flow-sensitive references. Instead, we reuse the same constructs introduced for flow-sensitive references (i.e., **new**(e), $e_1 := e_2$, **!** e , and **labelOfRef**(e) from Figure 2.1)⁷ and rely on the type-level sensitivity tag to determine which kind of reference is used.⁸ Raw values for flow-sensitive references are plain addresses $n \in \mathbb{N}$. Since the label of flow-sensitive references is *mutable* (as explained above), these addresses are *not* annotated with a label like flow-insensitive references (e.g., n_ℓ) and hence do not point into a labeled memory in the partitioned

⁷We do not bother to introduce separate constructs because the translations of these constructs given in Sections 5 and 6 are the same for both kinds of references.

⁸In this presentation we assume that terms are implicitly well-typed. In our mechanized proofs, terms are explicitly and intrinsically well-typed.

Flow-sensitivity tags: $s ::= I \mid S$
 Types: $\tau ::= \dots \mid \mathbf{Ref} \ s \ \tau$
 Raw Values: $r ::= \dots \mid n$
 Heaps: $\mu ::= [] \mid r^\ell : \mu$
 Configurations: $c ::= \langle \Sigma, \mu, e \rangle$

(a) Syntax.

(NEW-FS)

$$\frac{\langle \Sigma, \mu, e \rangle \Downarrow_{pc}^\theta \langle \Sigma', \mu', v \rangle \quad n = |\mu'| \quad \mu'' = \mu'[n \mapsto v]}{\langle \Sigma, \mu, \mathbf{new}(e) \rangle \Downarrow_{pc}^\theta \langle \Sigma', \mu'', n^{pc} \rangle}$$

(READ-FS)

$$\frac{\langle \Sigma, \mu, e \rangle \Downarrow_{pc}^\theta \langle \Sigma', \mu', n^\ell \rangle \quad \mu'[n] = r^{\ell'}}{\langle \Sigma, \mu, !e \rangle \Downarrow_{pc}^\theta \langle \Sigma', \mu', r^{\ell \sqcup \ell'} \rangle}$$

(LABELOFREF-FS)

$$\frac{\langle \Sigma, \mu, e \rangle \Downarrow_{pc}^\theta \langle \Sigma', \mu', n^{\ell_1} \rangle \quad \mu'[n] = r^{\ell_2}}{\langle \Sigma, \mu, \mathbf{labelOfRef}(e) \rangle \Downarrow_{pc}^\theta \langle \Sigma', \mu', \ell_2^{\ell_1 \sqcup \ell_2} \rangle}$$

(WRITE-FS)

$$\frac{\langle \Sigma, \mu, e_1 \rangle \Downarrow_{pc}^\theta \langle \Sigma', \mu', n^\ell \rangle \quad \langle \Sigma', \mu', e_2 \rangle \Downarrow_{pc}^\theta \langle \Sigma'', \mu'', r_2^{\ell_2} \rangle \quad \mu''[n] = r_1^{\ell_1} \quad \ell \sqsubseteq \ell_1 \quad \mu''' = \mu''[n \mapsto r_2^{\ell_2 \sqcup \ell}]}{\langle \Sigma, \mu, e_1 := e_2 \rangle \Downarrow_{pc}^\theta \langle \Sigma'', \mu''', ()^{pc} \rangle}$$

(b) Dynamics. The shaded constraint $\ell \sqsubseteq \ell_1$ is the *no-sensitive upgrade* security check.

Figure 2.8: λ^{dFG} extended with flow-sensitive references.

store. Instead, we store both the content and the label of flow-sensitive references in the new, linear heap μ , which can contain data at different security levels. Specifically, a flow-sensitive reference n of type **Ref** S τ points to the n -th cell of the heap μ , i.e., $\mu[n] = r^\ell$, for some raw value r of type τ at security level ℓ , which also represents the label of the reference. Finally, we extend program configurations with a heap μ , e.g., $c = \langle \Sigma, \mu, e \rangle$.

2.3.2 Dynamics

Figure 2.8b gives the semantics rules for the λ^{dFG} constructs that operate on flow-sensitive references.⁹ Rule [NEW-FS] allocates a new flow-sensitive reference in the heap at fresh address $n = |\mu'|$, where value v gets stored, i.e., $\mu'[n \mapsto v]$. Similarly to rule [NEW] (Figure 2.4), this rule does not leak data through implicit flows because value v has security level at least equal to the program counter label.¹⁰ Rule [READ-FS] is also similar to rule [READ]. The rule reads reference n at security level ℓ from the heap μ' , i.e., $\mu'[n] = r^{\ell'}$, and upgrades the label of r to $\ell \sqcup \ell'$ to reflect the fact that reading this particular reference depends on information at security level ℓ . Rule [LABELOFREF-FS] retrieves the label of reference n by reading the corresponding cell in the heap, i.e., $\mu'[n] = r^{\ell_2}$, and protects it with the label itself and taints it with the security level of the reference, i.e., $\ell_2^{\ell_1 \sqcup \ell_2}$. Although the label of values in the heap represent also the label of the reference, we cannot obtain exactly the label of the reference by reading the reference and then projecting the label of the value, i.e., in general $\mathbf{labelOfRef}(e) \not\equiv \mathbf{labelOf}(!e)$ for a flow-sensitive reference e . To see this, suppose e evaluates to reference n^{ℓ_1} such that $\mu[n] = r^{\ell_2}$ for some raw value r . Then, expression $\mathbf{labelOfRef}(e)$ results in $\ell_2^{\ell_1 \sqcup \ell_2}$ through rule [LABELOFREF-FS], which is in general different from $(\ell_1 \sqcup \ell_2)^{\ell_1 \sqcup \ell_2}$ obtained from $\mathbf{labelOf}(!e)$ through rule [LABELOF] (Figure 2.3) applied to rule [READ-FS]. Hence, we need to define $\mathbf{labelOfRef}(\cdot)$ as a primitive construct of the calculus.

⁹The rules from Figure 2.2-2.4 are adapted by simply threading the heap through the intermediate evaluations.

¹⁰In particular, Property 1 holds also for λ^{dFG} extended with flow-sensitive references.

Rule [WRITE-FS] updates flow-sensitive references but, in contrast to rule [WRITE], it also updates the label of the (flow-sensitive) reference. In particular, the rule updates reference n at security level ℓ by writing value $r_2^{\ell_2}$ tainted with ℓ in the n -th cell of the heap, i.e., $\mu''[n \mapsto r_2^{\ell_2} \sqcup \ell]$. Notice that the label of the updated reference, i.e., $\ell_2 \sqcup \ell$, depends *only* on the security level of the data written into the reference and of the reference that gets updated—this label can be completely different from the label of the reference before the update, i.e., ℓ_1 from $\mu''[n] = r_1^{\ell_1}$. For example, the rule allows turning a public reference into a secret one (i.e., changing the label of the reference from L to H) by writing secret data into it (like in Example 2.1). Perhaps more surprisingly, the rule also allows changing the label in the opposite direction (i.e., from H to L). For example, it is possible to make a secret reference public by overwriting its content with public data (e.g., swap p and s in Example 2.1). However, rule [WRITE-FS] is not completely unrestricted: it includes a security check known as *no-sensitive upgrade* (NSU) to avoid leaking information through *implicit flows* (Austin and Flanagan, 2009). Intuitively, this check (shaded in Figure 2.8b) prohibits updates in which the decision to update a particular reference depends on information that is *more sensitive* than the label of the reference itself. The following example shows why this check is needed to avoid leaks.

Example 2.2. Consider the following program, which branches on secret data, updates an initially public (flow-sensitive) reference in one branch, and then reads back the reference.

```

let  $r = \text{new}(p)$  in
  if  $s$  then  $r := s$  else ();
  ! $r$ 

```

Without the no-sensitive upgrade check described above, this program *leaks* information through an implicit flow. Formally, we show that the program above breaks non-interference (Theorem 1). Consider two executions of the program with program counter label public, i.e., $pc = L$, and with two L -equivalent input environments, e.g., $\theta_1 = [p \mapsto \text{true}^L, s \mapsto \text{false}^H] \approx_L [p \mapsto \text{true}^L, s \mapsto \text{true}^H] = \theta_2$. In the first

execution with environment θ_1 , reference r is *not* updated, and thus the result of the program is public value $p = \mathbf{true}^L$. Instead, in the second execution with environment θ_2 , the program updates reference r (rule [WRITE-FS] without the NSU check) and thus we obtain the secret value $s = \mathbf{true}^H$ as result. However, this breaks Theorem 1: the input environments are L -equivalent, i.e., $\theta_1 \approx_L \theta_2$, but the results of the program are not, i.e., $p = \mathbf{true}^L \not\approx_L \mathbf{true}^H = s$, because $L \sqsubseteq L$, but $H \not\sqsubseteq L$ (neither rule [VALUE_H] nor [VALUE_L] from Figure 2.5 apply).¹¹

To avoid such leaks, we include the no-sensitive upgrade check in rule [WRITE-FS] through the constraint $\ell \sqsubseteq \ell_1$, which ensures that the decision of updating reference n depends on information at some security level ℓ below the current label of the reference, i.e., ℓ_1 from $\mu''[n] = r_1^{\ell_1}$. This constraint causes the IFC monitor to abort the second execution of the program above ($H \not\sqsubseteq L$), which gets stuck and thus satisfies (trivially) Theorem 1.

2.3.3 Security

λ^{dFG} extended with flow-sensitive references is also secure, i.e., it satisfies termination insensitive non-interference. However, the theorem for the extended calculus is more complicated than Theorem 1: it requires a relaxed notion of L -equivalence up to a *bijection*, which relates observable flow-sensitive references with different heap addresses, and a side-condition on program configurations to rule out dangling references (Banerjee and Naumann, 2005). In this section, we make these differences rigorous and we reestablish non-interference.

The need for Bijections. Since the label of flow-sensitive references can change throughout program execution, our semantics allocates cells for both public and secret references in a single, unlabeled heap. An unfortunate consequence of this design choice is that allocations made in secret contexts can influence the addresses of public references allocated in the rest of the program, which get shifted by the number of secret

¹¹The mismatch in the label of the results can also be propagated to raw values through the label introspection primitives. We refer to (Austin and Flanagan, 2009) for a more elaborated example that leaks directly through raw values.

references previously allocated. This can be problematic for security because attackers can exploit these deterministic side-effects to leak information through the address of public references (Hedin and Sands, 2006). However, the addresses considered in this work are *opaque*, i.e., programs cannot inspect and compare them, and so attackers cannot use them to leak secrets. To show that, we extend the L -equivalence relation from Section 2.2 with a *bijection* (Banerjee and Naumann, 2005), a one-to-one correspondence between heap addresses, which we use to relate the addresses of references allocated in public contexts.

Definition 1 (Bijection). A bijection $\beta : Addr \rightarrow Addr$ is a bijective finite partial function between heap addresses. Formally, for all addresses n and n' , $\beta(n) = n' \iff \beta^{-1}(n') = n$, where β^{-1} is the inverse function of β .

Notation and Auxiliary Definitions. In the following, we treat partial functions as sets of input-output pairs and for clarity we write $(n_1, n_2) \in \beta$ for $\beta(n_1) = n_2$, i.e., if the partial function β is defined on input n_1 and has output n_2 . We say that a bijection β' *extends* another bijection β , written $\beta \subseteq \beta'$, if and only if for all input-output pairs (n_1, n_2) , $(n_1, n_2) \in \beta$ implies $(n_1, n_2) \in \beta'$. Bijection composition is written $\beta_2 \circ \beta_1 = \{(n_1, n_3) \mid (n_1, n_2) \in \beta_1 \wedge (n_2, n_3) \in \beta_2\}$ and the inverse operator as $\beta^{-1} = \{(n_2, n_1) \mid (n_1, n_2) \in \beta\}$. We also define the domain and range of bijections in the obvious way, i.e., $dom(\beta) = \{n_1 \mid (n_1, n_2) \in \beta\}$, and $rng(\beta) = \{n_2 \mid (n_1, n_2) \in \beta\}$. Furthermore, we write ι_n for the *finite*, partial identity bijection defined up to n , i.e., $\iota_n = \{(n', n') \mid n' \in \{0, \dots, n-1\}\}$. Identity bijections satisfy the following laws.

Property 3 (Identity Laws). For all bijections β and $n \in \mathbb{N}$:

1. **Inverse Identity.** $\iota_n^{-1} \equiv \iota_n$.
2. **Absorb Left.** If $rng(\beta) \subseteq dom(\iota_n)$, then $\iota_n \circ \beta \equiv \beta$.
3. **Absorb Right.** If $dom(\beta) \subseteq rng(\iota_n)$, then $\beta \circ \iota_n \equiv \beta$.

L -Equivalence up to Bijection. Formally, we redefine L -equivalence as a relation $\approx_L^\beta \subseteq \text{Value} \times \text{Value}$ and write $v_1 \approx_L^\beta v_2$ to indicate that values v_1 and v_2 are indistinguishable to an attacker at security level L up to bijection β , which relates the heap addresses of corresponding observable flow-sensitive references in v_1 and v_2 . We extend the L -equivalence relation for all other syntactic categories (raw values, memories, and stores) with a bijection in the same way. Besides the bijection β , which we only add to \approx_L , all the rules from Figure 2.5 remain unchanged. In particular, rules [REF $_L$] and [REF $_H$] for flow-insensitive references simply ignore the bijection. Since we keep labeled memories partitioned in the store and isolated from the heap, the problem of mismatching addresses does not arise for flow-insensitive references. Intuitively, the *address space* of public memories is completely independent from the address space of secret memories, and thus allocations made in secret contexts cannot influence the memory addresses of public references. As a result, rule [REF $_L$] guarantees that L -equivalent public references have exactly the same *memory address*. In contrast, heap allocations made in secret contexts can shift the *heap addresses* of observable flow-sensitive references (as explained above), which then may not be necessarily the same. Therefore, we relate observable flow-sensitive references with possibly different addresses through the following new rule:

$$\begin{array}{c} \text{(REF-FS)} \\ (n_1, n_2) \in \beta \\ \hline n_1 \approx_L^\beta n_2 \end{array}$$

Rule [REF-FS] relates references with arbitrary heap addresses n_1 and n_2 , as long as their addresses are “matched” by the bijection β , i.e., $(n_1, n_2) \in \beta$. Similarly, we define heap L -equivalence by relating heap cells that are at corresponding addresses according to the bijection.

Definition 2 (Heap L -equivalence). Two heaps μ_1 and μ_2 are L -equivalent up to bijection β , written $\mu_1 \approx_L^\beta \mu_2$, if and only if:

1. $\text{dom}(\beta) \subseteq \{0, \dots, |\mu_1| - 1\}$,
2. $\text{rng}(\beta) \subseteq \{0, \dots, |\mu_2| - 1\}$, and
3. For all addresses n_1 and n_2 , if $(n_1, n_2) \in \beta$, then $\mu_1[n_1] \approx_L^\beta \mu_2[n_2]$.

In the definition above, the first two conditions require that the domain and the range of β are contained in the address space of μ_1 and μ_2 , respectively.¹² These side-conditions ensure that the heap addresses considered in the third condition are *valid*, i.e., they point to some cell in each heap. The definition of L -equivalence for heaps is complicated by the fact that heaps can contain both public and secret cells, i.e., cells allocated in public and secret contexts, respectively. In particular, the L -equivalence relation needs to relate the values contained in corresponding public cells, which are at possibly different addresses in the heaps.¹³ This task is simplified by the bijection β , which identifies exactly the addresses of corresponding public cells. Therefore, the third condition above ensures that the values contained in the cells at addresses related by the bijection β , i.e., $(n_1, n_2) \in \beta$, are themselves L -equivalent, i.e., $\mu_1[n_1] \approx_L^\beta \mu_2[n_2]$. Finally, we extend L -equivalence for configurations with a bijection and additionally require their heaps to be related. Formally, $c_1 \approx_L^\beta c_2$ iff $c_1 = \langle \Sigma_1, \mu_1, v_1 \rangle$, $c_2 = \langle \Sigma_2, \mu_2, v_2 \rangle$, and all their components are related, i.e., $\Sigma_1 \approx_L^\beta \Sigma_2$, $\mu_1 \approx_L^\beta \mu_2$, and $v_1 \approx_L^\beta v_2$.

Next, we discuss some side-conditions over program configurations and inputs that are needed for reasoning with L -equivalence relations up to bijection.

Valid References. In Section 2.2, we proved that the L -equivalence relation is reflexive, symmetric, and transitive, and showed how these properties helped us structure and simplify the proof of non-interference. Does the L -equivalence extended with bijection also enjoy these properties? Unfortunately, reflexivity does not hold *unconditionally* anymore for L -equivalence relations up to bijection. The culprit of the problem is in the definition of L -equivalence for heaps. In this case, proving reflexivity requires showing that any heap is L -equivalent to itself up to the identity bijection defined over its address space, i.e., $\forall \mu. \mu \approx_L^{|\mu|} \mu$. Technically, this property does not hold in general for any *arbitrary* heap μ . To see this, consider for example the ill-formed heap $\mu' = 1^L : []$, which contains the *dangling* flow-sensitive reference 1^L . In order to

¹²Heap addresses start at 0 just like memory addresses.

¹³Secret cells in each heap are simply disregarded as they may not necessarily have a counterpart in the other heap.

prove $\mu' \approx_L^{\iota_1} \mu'$, we must show that this dangling reference is L -equivalent to itself up to bijection ι_1 (Definition 2.3), i.e., $1^L \approx_L^{\iota_1} 1^L$, which by rule [REF-FS] would require a proof for the false statement $(1, 1) \in \iota_1 = \{(0, 0)\}$. One may be tempted to repair reflexivity by choosing a sufficiently large domain for the identity bijection, so that also dangling references can be related (e.g., for μ' we could pick ι_2 and prove $1^L \approx_L^{\iota_2} 1^L$). Unfortunately, this is not possible without breaking the first and the second condition of Definition 2. For any heap μ , the identity bijection $\iota_{|\mu|}$ has the largest domain and range that satisfies the side-conditions $\text{dom}(\iota_{|\mu|}) \subseteq \{0 \dots |\mu| - 1\}$ (Definition 2.1) and $\text{rng}(\iota_{|\mu|}) \subseteq \{0 \dots |\mu| - 1\}$ (Definition 2.2).

The solution to this technical nuisance is to simply prove a weaker version of reflexivity, which holds only for *well-formed* heaps (and similarly values, stores, memories, and configurations) that contain only *valid* (i.e., not dangling) references. Luckily, this restriction is inconsequential in our setting because references are always valid in λ^{dFG} . Intuitively, references are *unforgeable*: programs can only create fresh, valid references through the **new**(\cdot) construct. After creation references simply remain valid throughout program execution: they can be accessed, passed around, or stored, but never deallocated (we simply provide no construct to do so) or tempered. We now formalize this insight in the form of a *valid judgment*, which rules out dangling reference from all the categories of the calculus together with the proof that this judgment is an *invariant* of the operational semantics of λ^{dFG} .

Valid Judgment. Figure 2.9 defines a valid judgment for λ^{dFG} values and configurations. These judgments ensure that all flow-sensitive references values of a program configuration are valid (i.e., not dangling) in a heap of a given size. For values (Figure 2.9a), this judgment takes the form of $n \vdash \mathbf{Valid}(v)$, which indicates that all heap addresses contained in value v are valid in a heap of size n . This judgment is mutually defined with similar judgments for raw values, i.e., $n \vdash \mathbf{Valid}(r)$, and environments $n \vdash \mathbf{Valid}(\theta)$. Importantly, rule [VALID-FS-REF] disallows *dangling references*: a flow-sensitive reference is valid if and only if its heap address n' is strictly smaller than the size n of the heap. Notice that this judgment holds unconditionally for *flow-insensitive*

values through rule [VALID-FI-REF]. Intuitively, these references contain memory addresses, which are not affected by the technical issue described above. Most of the remaining rules are fairly straightforward: they simply require that all sub-values are homogeneously valid in heaps of the same size. For example, values r^ℓ is valid only if and only if the raw value r is also valid (rule [VALID-VALUE]) and closures $(x.e, \theta)$ are valid if and only if the value environment θ is also valid (rule [VALID-CLOSURE]).¹⁴ Environments θ are valid if they contain only valid values, i.e., $n \vdash \mathbf{Valid}(\theta)$ iff $\forall x \in \text{dom}(\theta). n \vdash \mathbf{Valid}(\theta(x))$. Similarly, valid stores Σ contain only valid memories, i.e., $n \vdash \mathbf{Valid}(\Sigma)$ iff $\forall \ell \in \mathcal{L}. n \vdash \mathbf{Valid}(\Sigma(\ell))$, and so on for memories and heaps, i.e., $n \vdash \mathbf{Valid}(M)$ iff $\forall n \in \{0, \dots, |M| - 1\}. n \vdash \mathbf{Valid}(M[n])$ and $n \vdash \mathbf{Valid}(\mu)$ iff $\forall n \in \{0, \dots, |\mu| - 1\}. n \vdash \mathbf{Valid}(\mu[n])$, respectively. Finally, the judgments for initial and final configurations (Figure 2.9b) instantiate the parameter n of the judgments above with the size of the heap. For example, a configuration $c = \langle \Sigma, \mu, e \rangle$ and input environment θ are valid through rule [VALID-INPUTS] if and only if all the components of the configuration and the value environment are valid in the heap of size $n = |\mu|$, i.e., $\vdash \mathbf{Valid}(c, \theta)$ iff $n \vdash \mathbf{Valid}(\Sigma)$, $n \vdash \mathbf{Valid}(\mu)$, and $n \vdash \mathbf{Valid}(\theta)$. Rule [VALID-FINAL] is for final configurations and is similar.

Before showing that this valid judgment is an invariant of the semantics of λ^{dFG} , we prove two simple helping lemmas. The first lemma shows that values that are valid in a heap of a certain size, are also valid in any larger heap, while the second lemma shows that heaps can only grow larger in the semantics of λ^{dFG} .

Lemma 2.5 (Valid Weakening). For all values v , environments θ , raw values r and natural numbers n and n' :

1. If $n \vdash \mathbf{Valid}(v)$ and $n \leq n'$, then $n' \vdash \mathbf{Valid}(v)$,
2. If $n \vdash \mathbf{Valid}(r)$ and $n \leq n'$, then $n' \vdash \mathbf{Valid}(r)$
3. If $n \vdash \mathbf{Valid}(\theta)$ and $n \leq n'$, then $n' \vdash \mathbf{Valid}(\theta)$

¹⁴Since expressions do not contain heap addresses or values (they belong to distinct syntactic categories), we do not need to define a valid judgment for expressions.

$$\begin{array}{c}
\text{(VALID-VALUE)} \\
\frac{n \vdash \mathbf{Valid}(r)}{n \vdash \mathbf{Valid}(r^\ell)} \\
\\
\text{(VALID-FS-REF)} \\
\frac{n' < n}{n \vdash \mathbf{Valid}(n')} \\
\\
\text{(VALID-FI-REF)} \\
n \vdash \mathbf{Valid}(n'_\ell) \\
\\
\text{(VALID-CLOSURE)} \\
\frac{n \vdash \mathbf{Valid}(\theta)}{n \vdash \mathbf{Valid}((x.e, \theta))}
\end{array}$$

(a) Judgment for values $n \vdash \mathbf{Valid}(v)$ and raw values $n \vdash \mathbf{Valid}(r)$ (selected cases).

$$\begin{array}{c}
\text{(VALID-INPUTS)} \\
\frac{n = |\mu| \quad n \vdash \mathbf{Valid}(\Sigma) \quad c = \langle \Sigma, \mu, e \rangle \quad n \vdash \mathbf{Valid}(\mu) \quad n \vdash \mathbf{Valid}(\theta)}{\vdash \mathbf{Valid}(c, \theta)}
\end{array}$$

$$\begin{array}{c}
\text{(VALID-OUTPUTS)} \\
\frac{n = |\mu| \quad n \vdash \mathbf{Valid}(\Sigma) \quad c = \langle \Sigma, \mu, v \rangle \quad n \vdash \mathbf{Valid}(\mu) \quad n \vdash \mathbf{Valid}(v)}{\vdash \mathbf{Valid}(c)}
\end{array}$$

(b) Judgments for valid program inputs and outputs in λ^{dFG} .

Figure 2.9: Some valid judgments for λ^{dFG} .

Proof. By mutual induction over the judgments and using transitivity of \leq for the case [VALID-FS-REF]. \square

Lemma 2.6 (Heaps Only Grow). Let $c = \langle \Sigma, \mu, e \rangle$, $c' = \langle \Sigma', \mu', v \rangle$, if $c \Downarrow_{pc}^\theta c'$, then $|\mu| \leq |\mu'|$.

Proof. By induction on the big-step reduction. In particular, notice that no rule deallocates cells from the heap, which can only grow through rule [NEW-FS]. \square

Property 4 (Valid Invariant). If $c \Downarrow_{pc}^\theta c'$ and $\vdash \mathbf{Valid}(c, \theta)$, then $\vdash \mathbf{Valid}(c')$.

Proof. By induction on the big-step reduction and using Lemmas 2.5 and 2.6. \square

We now reconsider the properties of L -equivalence up to bijection.

Property 5. For all values, raw values, environments, memories, and stores x, y, z , and $n \in \mathbb{N}$:

1. **Restricted Reflexivity.** If $n \vdash \mathbf{Valid}(x)$, then $x \approx_L^{t_n} x$.
2. **Symmetry.** If $x \approx_L^\beta y$, then $y \approx_L^{\beta^{-1}} x$.
3. **Transitivity.** If $x \approx_L^{\beta_1} y$ and $y \approx_L^{\beta_2} z$, then $x \approx_L^{\beta_2 \circ \beta_1} z$.
4. **Weakening.** If $x \approx_L^\beta y$ and $\beta \subseteq \beta'$, then $x \approx_L^{\beta'} y$.

First, *reflexivity* up to a *identity bijection* t_n is restricted to program constructs that are valid in a heap of size n , as explained above. In a relation $x \approx_L^\beta y$, the bijection β maps the heap addresses of x to the corresponding heap addresses of y . Then, in order to obtain the mapping from y to x , the bijection must be inverted in the symmetric relation, i.e., $y \approx_L^{\beta^{-1}} x$. Similarly, when composing $x \approx_L^{\beta_1} y$ and $y \approx_L^{\beta_2} z$ through *transitivity*, the bijections must also be composed, i.e., $x \approx_L^{\beta_2 \circ \beta_1} z$, so that the composed bijection $\beta_2 \circ \beta_1$ first maps the addresses of x to the corresponding addresses of y through β_1 , and then the addresses of y to those of z through β_2 . Lastly, *weakening* allows to relax a relation $x \approx_L^\beta y$ with an extended bijection $\beta' \supseteq \beta$. All the properties above hold also for L -equivalence for heaps (Definition 2), with the exception for *weakening*, due to the side conditions on the domain and range of the bijection.

We conclude this part with the *square commutative digram* for heaps, outlined in Figure 2.10.¹⁵ In the diagram, the arrows connect L -equivalent heaps up to a given bijection, e.g., the arrow from μ_1 to μ'_1 labeled β_1 indicates that $\mu_1 \approx_L^{\beta_1} \mu'_1$. Compared to the square diagram without bijections from Figure 2.6, this diagram is complicated by the fact that the bijections must be composed and inverted appropriately, as needed by *symmetry* and *transitivity*, in order to obtain the

¹⁵We omit the analogous square commutative diagram for stores.

$$\begin{array}{ccc}
\mu_1 & \xrightarrow{\beta_1} & \mu'_1 \\
\beta \downarrow & & \downarrow \beta' = \beta_2 \circ \beta \circ \beta_1^{-1} \\
\mu_2 & \xrightarrow{\beta_2} & \mu'_2
\end{array}$$

Figure 2.10: Square commutative diagram for heaps (Lemma 2.7). Solid arrows represent the assumptions of the lemma and the dashed arrow represents the conclusion. If $\beta_1 = \iota_{|\mu_1|}$, $\beta_2 = \iota_{|\mu_2|}$, then β' is simply equal to β .

dashed arrow that completes the square. However, when β_1 and β_2 are (appropriate) identity bijections, the bijection between the final heaps can be simplified and shown to be equal to the bijection between the initial heaps.

Lemma 2.7 (Square Commutative Diagram for Heaps). If $\mu_1 \approx_L^\beta \mu_2$, $\mu_1 \approx_L^{\beta_1} \mu'_1$, and $\mu_2 \approx_L^{\beta_2} \mu'_2$, then $\mu'_1 \approx_L^{\beta'} \mu'_2$, where $\beta' = \beta_2 \circ \beta \circ \beta_1^{-1}$. Furthermore, if $\beta_1 = \iota_{|\mu_1|}$ and $\beta_2 = \iota_{|\mu_2|}$, then $\mu'_1 \approx_L^\beta \mu'_2$.

Proof. By applying *symmetricity* (Property 5.2) and *transitivity* (Property 5.3) like in Lemma 2.1, we obtain directly $\mu'_1 \approx_L^{\beta'} \mu'_2$ and $\beta' = \beta_2 \circ \beta \circ \beta_1^{-1}$. If $\beta_1 = \iota_{|\mu_1|}$ and $\beta_2 = \iota_{|\mu_2|}$, then $\beta' = \iota_{|\mu_2|} \circ \beta \circ \iota_{|\mu_1|}^{-1}$. Using the *identity laws* (Property 3), we show that $\beta' = \iota_{|\mu_2|} \circ \beta \circ \iota_{|\mu_1|}^{-1} = \beta$ and thus we have $\mu'_1 \approx_L^\beta \mu'_2$. \square

Termination-Insensitive Non-Interference. We now turn our attention to the termination-insensitive non-interference theorem for λ^{dFG} extended with flow-sensitive references. First, the theorem assumes that the initial configurations are L -equivalent up to a bijection, which we use to account for secret-dependent heap allocations and hence relate the heap addresses of corresponding public flow-sensitive references, as explained above. Then, the theorem guarantees that the final configurations are also L -equivalent, but up to some *extended bijection*. Intuitively, programs may dynamically allocate new cells in the heap, therefore the heap addresses contained in their final configurations must be shown to be related with respect to the address space of the final heaps. Additionally, the theorem explicitly assumes that the initial

program configurations are valid and thus do not contain dangling references. As explained before, this extra assumption is needed for technical reasons and, besides having to explicitly propagate it as needed in our mechanized proof scripts, it does not impact the security guarantees of λ^{dFG} . The proof of this theorem is also structured on two key lemmas: *store and heap confinement* and *L-equivalence preservation*. In the following, we focus on the changes needed to adapt these lemmas for flow-sensitive references, heaps, and *L*-equivalence up to bijection.

Confinement guarantees that programs running in a secret context cannot leak secret data implicitly through the program state, i.e., the store and the heap. Formally, we show that if the program counter label is above the attacker's label ($pc \not\sqsubseteq L$), then the initial store and heap are *L*-equivalent to the final store and heap obtained at the end of the execution, up to the *identity bijection*. In the lemma, *L*-equivalence is up to the identity bijection defined over the address space of the initial heap. Intuitively, this is because the references allocated in secret contexts cannot be observed by the attacker, and therefore must be disregarded by the bijection, which instead keeps track only of observable references, i.e., those allocated in public contexts.

Lemma 2.8 (Store and Heap Confinement). For all configurations $c = \langle \Sigma, \mu, e \rangle$, $c' = \langle \Sigma', \mu', v \rangle$, program counter labels $pc \not\sqsubseteq L$, if $\vdash \mathbf{Valid}(c, \theta)$ and $c \Downarrow_{pc}^{\theta} c'$, then $\Sigma \approx_L^{|\mu|} \Sigma'$ and $\mu \approx_L^{|\mu|} \mu'$.

Proof. By induction on the big-step reduction and using the fact the semantics preserves valid references (Property 4) to propagate the valid judgment through the intermediate configurations. In particular, we apply *restricted reflexivity* (Property 5.1) in the base cases, and *transitivity* (Property 5.3) and the bijection *identity laws* (Property 3) in the inductive cases. Intuitively, when we apply transitivity, we need to show that the composition of identity bijections defined over different heaps results in the identity bijection over the initial heap required by the lemma. For example, to show that $\mu \approx_L^{|\mu|} \mu''$ through some intermediate heap μ' such that $\mu \approx_L^{|\mu|} \mu'$ and $\mu' \approx_L^{|\mu'|} \mu''$, we apply transitivity and obtain $\mu \approx_L^{|\mu'| \circ |\mu|} \mu''$. Since heaps only grow in size (Lemma 2.6), we have that $|\mu| \leq |\mu'|$, thus $rng(\iota_{|\mu|}) \subseteq dom(\iota_{|\mu'|})$

and $\iota_{|\mu'|} \circ \iota_{|\mu|} = \iota_{|\mu|}$ by Property 3.2 and so we have $\mu \approx_L^{\iota_{|\mu|}} \mu''$, as required. \square

For L -equivalence preservation, we assume that the program inputs are L -equivalent up to some bijection β and must show that the final configurations are L -equivalent up to some extended bijection $\beta' \supseteq \beta$. Since we cannot predict, in general, how arbitrary programs allocate references at run-time, in these lemmas we *existentially quantify* the bijection β' that relates the heap addresses in final configurations, but guarantee that the final bijection β' is at least as large as the initial bijection β , i.e., $\beta \subseteq \beta'$. This invariant is critical for the proof, where the inductive steps require combining L -equivalent values and environments with different bijections. In particular, the fact that bijections only get extended, e.g., $\beta \subseteq \beta'$, lets us lift L -equivalence relations up to a small bijection β into L -equivalence relations up to a larger bijection β' via *weakening* (Property 5.4).

Like before, we prove two L -equivalence preservation lemmas, for secret and public contexts, respectively, which we then combine in the proof of termination-insensitive non-interference. For preservation in secret contexts, we observe that the bijection that relates the final configurations is *exactly the same* bijection that relates the initial configurations.¹⁶ Since the program executions occur in secret contexts, their allocations are secret-dependent and therefore not observable by the attacker, which is reflected in the lemma by relating the initial and final configurations with the same bijection. The proof technique for this lemma is similar to that of the previous L -equivalence preservation lemma. We consider each program execution individually and reason independently about program state (i.e., heap and store) and values in the final configurations. In particular, we apply *store and heap confinement* to each execution and relate the initial and final program states with the *identity bijection*, which are then combined in a square commutative diagram (Figure 2.10) to relate the final configurations.

¹⁶This is a special case of L -equivalence preservation, where we can predict the final bijection β , and therefore we omit the existential quantification. Notice that in this case the invariant $\beta \subseteq \beta$ is trivially satisfied.

Lemma 2.9 (*L -Equivalence Preservation in Secret Contexts*). For all program counter labels $pc_1 \not\sqsubseteq L$ and $pc_2 \not\sqsubseteq L$, valid inputs $\vdash \mathbf{Valid}(c_1, \theta_1)$ and $\vdash \mathbf{Valid}(c_2, \theta_2)$, such that $c_1 = \langle \Sigma_1, \mu_1, e_1 \rangle$, $c_2 = \langle \Sigma_2, \mu_2, e_2 \rangle$, $\Sigma_1 \approx_L^\beta \Sigma_2$ and $\mu_1 \approx_L^\beta \mu_2$, if $\langle \Sigma_1, \mu_1, e_1 \rangle \Downarrow_{pc_1}^{\theta_1} c_1$ and $\langle \Sigma_2, \mu_2, e_2 \rangle \Downarrow_{pc_2}^{\theta_2} c_2$, then $c_1 \approx_L^\beta c_2$.

Proof. Assume $pc_1 \not\sqsubseteq L$, $pc_2 \not\sqsubseteq L$, $\Sigma_1 \approx_L^\beta \Sigma_2$, $\mu_1 \approx_L^\beta \mu_2$, and let the final configurations be $c_1 = \langle \Sigma'_1, \mu'_1, v_1 \rangle$ and $c_2 = \langle \Sigma'_2, \mu'_2, v_2 \rangle$. First, we apply *Store and Heap Confinement* (Lemma 2.8) to $\langle \Sigma_1, \mu_1, e_1 \rangle \Downarrow_{pc_1}^{\theta_1} \langle \Sigma'_1, \mu'_1, v_1 \rangle$ and $\langle \Sigma_2, \mu_2, e_2 \rangle \Downarrow_{pc_2}^{\theta_2} \langle \Sigma'_2, \mu'_2, v_2 \rangle$ and obtain $\mu_1 \approx_L^{|\mu_1|} \mu'_1$ and $\mu_2 \approx_L^{|\mu_2|} \mu'_2$, respectively. Then, we construct the *Square Commutative Diagram for Heaps* (Lemma 2.7) using $\mu_1 \approx_L^\beta \mu_2$, $\mu_1 \approx_L^{|\mu_1|} \mu'_1$ and $\mu_2 \approx_L^{|\mu_2|} \mu'_2$, and obtain $\mu'_1 \approx_L^\beta \mu'_2$. We derive $\Sigma'_1 \approx_L^\beta \Sigma'_2$ with a similar construction and prove $v_1 \approx_L^\beta v_2$ like in Lemma 2.3. \square

The proof of L -equivalence preservation in public contexts also follows the same structure of the proof without heaps. In particular, the proof is by simultaneous induction on the two big-step reductions and relies on *L -Equivalence Preservation in Secret Contexts* when the program executions deviate from each other at a secret-dependent control-flow point. Here, we discuss only the interesting case where the programs allocate a flow-sensitive reference and we must extend the bijection with a new pair of matching addresses.

Lemma 2.10 (*L -Equivalence Preservation in Public Contexts*). For all program counter labels $pc \sqsubseteq L$, valid inputs $\vdash \mathbf{Valid}(c_1, \theta_1)$ and $\vdash \mathbf{Valid}(c_2, \theta_2)$, such that $c_1 \approx_L^\beta c_2$, $\theta_1 \approx_L^\beta \theta_2$, if $c_1 \Downarrow_{pc}^{\theta_1} c'_1$, and $c_2 \Downarrow_{pc}^{\theta_2} c'_2$, then there exists an extended bijection $\beta' \supseteq \beta$, such that $c'_1 \approx_L^{\beta'} c'_2$.

Proof. From L -equivalence, i.e., $c_1 \approx_L^\beta c_2$, the program expressions are α -equivalent, therefore the two reductions step following the same rule and produce L -equivalent results and perform L -equivalent operations on the program store and heap. In particular, for rule [NEW-FS], the programs allocate L -equivalent heap cells, e.g., $v_1 \approx_L^{\beta'} v_2$ for some extended bijection $\beta' \supseteq \beta$ by induction hypothesis, at fresh addresses $n_1 = |\mu'_1|$ and $n_2 = |\mu'_2|$ of related heaps $\mu'_1 \approx_L^{\beta'} \mu'_2$, which remain

related under the *extended* bijection $\beta'' = \beta' \cup \{(n_1, n_2)\} \supseteq \beta$ by transitivity, i.e., $\mu_1[n_1 \mapsto v_1] \approx_L^{\beta''} \mu_2[n_2 \mapsto v_2]$, and so produce related references $n_1^L \approx_L^{\beta''} n_2^L$ by rule [VALUE_L] applied to rule [REF-FS]. \square

Lastly, we combine the L -equivalence preservation lemmas for public and secret contexts and prove termination-insensitive non-interference for λ^{dFG} extended with flow-sensitive references.

Theorem 2 (λ^{dFG} -TINI with Bijections). For all program counter labels pc and valid inputs $\vdash \mathbf{Valid}(c_1, \theta_1)$ and $\vdash \mathbf{Valid}(c_2, \theta_2)$, such that $c_1 \approx_L^\beta c_2$, $\theta_1 \approx_L^\beta \theta_2$, if $c_1 \Downarrow_{pc}^{\theta_1} c'_1$, and $c_2 \Downarrow_{pc}^{\theta_2} c'_2$, then there exists an extended bijection $\beta' \supseteq \beta$, such that $c'_1 \approx_L^{\beta'} c'_2$.

Conclusion. Dynamic language-based fine-grained IFC, of which λ^{dFG} is just a particular instance, represents an intuitive approach to tracking information flows in programs. Programmers annotate input values with labels that represent their sensitivity and a label-aware instrumented security monitor propagates those labels during execution and computes the result of the program together with a conservative approximation of its sensitivity. The next section describes an IFC monitor that tracks information flows at *coarse* granularity.

3

Coarse-Grained IFC Calculus

One of the drawbacks of dynamic fine-grained IFC is that the programming model requires all input values to be explicitly and fully annotated with their security labels. Imagine a program with many inputs and highly structured data: it quickly becomes cumbersome, if not impossible, for the programmer to specify all the labels. The label of some inputs may be sensitive (e.g., passwords, pin codes, etc.), but the sensitivity of the rest may probably be irrelevant for the computation, yet a programmer must come up with appropriate labels for them as well. The programmer is then torn between two opposing risks: over-approximating the actual sensitivity can negatively affect execution (the monitor might stop secure programs), under-approximating the sensitivity can endanger security. Even worse, specifying many labels manually is error-prone and assigning the wrong security label to a piece of sensitive data can be catastrophic for security and completely defeat the purpose of IFC. Dynamic coarse-grained IFC represents an attractive alternative that requires fewer annotations, in particular it allows the programmer to label only the inputs that need to be protected.

Types: $\tau ::= \mathbf{unit} \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 + \tau_2 \mid \tau_1 \times \tau_2 \mid \mathcal{L}$
 $\mid \mathbf{LIO} \tau \mid \mathbf{Labeled} \tau \mid \mathbf{Ref} \tau$

Labels: $\ell, pc \in \mathcal{L}$

Addresses: $n \in \mathbb{N}$

Environments: $\theta \in \mathit{Var} \rightarrow \mathit{Value}$

Values: $v ::= () \mid (x.e, \theta) \mid \mathbf{inl}(v) \mid \mathbf{inr}(v) \mid (v_1, v_2) \mid \ell$
 $\mid \mathbf{Labeled} \ell v \mid (t, \theta) \mid n_\ell$

Expressions: $e ::= x \mid \lambda x.e \mid e_1 e_2 \mid () \mid \ell \mid (e_1, e_2) \mid \mathbf{fst}(e) \mid \mathbf{snd}(e)$
 $\mid \mathbf{inl}(e_1) \mid \mathbf{inr}(e_2) \mid \mathbf{case}(e, x.e_1, x.e_2) \mid e_1 \sqsubseteq^? e_2 \mid t$

Thunks: $t ::= \mathbf{return}(e) \mid \mathbf{bind}(e, x.e) \mid \mathbf{unlabel}(e)$
 $\mid \mathbf{toLabeled}(e) \mid \mathbf{labelOf}(e) \mid \mathbf{getLabel} \mid \mathbf{taint}(e)$
 $\mid \mathbf{new}(e) \mid ! e \mid e_1 := e_2 \mid \mathbf{labelOfRef}(e)$

Type System: $\Gamma \vdash e : \tau$

Configurations: $c ::= \langle \Sigma, pc, e \rangle$

Stores: $\Sigma \in (\ell : \mathit{Label}) \rightarrow \mathit{Memory} \ell$

Memory ℓ : $M ::= [] \mid v : M$

Figure 3.1: Syntax of λ^{dCG} .

Syntax. Figure 3.1 shows the syntax of λ^{dCG} , a standard simply-typed λ -calculus extended with security primitives for dynamic coarse-grained IFC, inspired by Stefan *et al.* (2011) and adapted to use call-by-value instead of call-by-name to match λ^{dFG} . The λ^{dCG} -calculus features both standard (unlabeled) values and *explicitly labeled* values. For example, **Labeled** H **true** represents a secret boolean value of type **Labeled bool**.¹ The type constructor **LIO** encapsulates a security state monad, whose state consists of a labeled store and the program counter label. In addition to standard **return**(\cdot) and **bind**(\cdot) constructs, the

¹As in λ^{dFG} , we define **bool** = **unit** + **unit** and **if** e **then** e_1 **else** e_2 = **case**($e, _ . e_1, _ . e_2$). Unlike λ^{dFG} values, λ^{dCG} values are not intrinsically labeled, thus we encode boolean constants simply as **true** = **inl**($()$) and **false** = **inr**($()$).

$$\begin{array}{c}
\text{(THUNK)} \qquad \qquad \qquad \text{(FUN)} \qquad \qquad \qquad \text{(VAR)} \\
t \Downarrow^\theta (t, \theta) \qquad \lambda x.e \Downarrow^\theta (x.e, \theta) \qquad x \Downarrow^\theta \theta(x) \\
\\
\text{(APP)} \\
\frac{e_1 \Downarrow^\theta (x.e, \theta') \quad e_2 \Downarrow^\theta v_2 \quad e \Downarrow^{\theta'[x \mapsto v_2]} v}{e_1 e_2 \Downarrow^\theta v} \\
\\
\text{(CASE}_1\text{)} \qquad \qquad \qquad \text{(CASE}_2\text{)} \\
\frac{e_1 \Downarrow^\theta \mathbf{inl}(v_1) \quad e_1 \Downarrow^{\theta[x \mapsto v_1]} v}{\mathbf{case}(e, x.e_1, x.e_2) \Downarrow^\theta v} \quad \frac{e_1 \Downarrow^\theta \mathbf{inr}(v_2) \quad e_2 \Downarrow^{\theta[x \mapsto v_2]} v}{\mathbf{case}(e, x.e_1, x.e_2) \Downarrow^\theta v}
\end{array}$$

Figure 3.2: Pure semantics: $e \Downarrow^\theta v$ (selected rules).

monad provides primitives that regulate the creation and the inspection of labeled values, i.e., **toLabeled**(\cdot), **unlabel**(\cdot) and **labelOf**(\cdot), and the interaction with the labeled store, allowing the creation, reading and writing of labeled references n_ℓ through the constructs **new**(e), $!e$, $e_1 := e_2$, respectively.² The primitives of the **LIO** monad are listed in a separate sub-category of expressions called *think*. Intuitively, a think is just a description of a stateful computation, which only the top-level security monitor can execute—a *think closure*, i.e., (t, θ) , provides a way to suspend nested computations.

3.1 Dynamics

In order to track information flows dynamically at coarse granularity, λ^{dCG} employs a technique called *floating-label*, which was originally developed for IFC operating systems (e.g., Zeldovich *et al.*, 2006; Zeldovich *et al.*, 2008) and that was later applied in a language-based setting. In this technique, throughout a program’s execution, the program counter label *floats* above the label of any value observed during program execution and thus represents (an upper bound on) the sensitivity of all the values that are not explicitly labeled. For this reason, λ^{dCG} stores

²We extend λ^{dCG} with *flow-sensitive* references and *heaps* in Section 3.3.

$$\begin{array}{c}
\text{(FORCE)} \\
\frac{e \Downarrow^\theta (t, \theta') \quad \langle \Sigma, pc, t \rangle \Downarrow^{\theta'} \langle \Sigma', pc', v \rangle}{\langle \Sigma, pc, e \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle} \\
\text{(a) Forcing semantics: } \langle \Sigma, pc, e \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle. \\
\hline
\text{(RETURN)} \\
\frac{e \Downarrow^\theta v}{\langle \Sigma, pc, \mathbf{return}(e) \rangle \Downarrow^\theta \langle \Sigma, pc, v \rangle} \\
\text{(BIND)} \\
\frac{\langle \Sigma, pc, e_1 \rangle \Downarrow^\theta \langle \Sigma', pc', v_1 \rangle \quad \langle \Sigma', pc', e_2 \rangle \Downarrow^{\theta[x_1 \mapsto v_1]} \langle \Sigma'', pc'', v \rangle}{\langle \Sigma, pc, \mathbf{bind}(e_1, x.e_2) \rangle \Downarrow^\theta \langle \Sigma'', pc'', v \rangle} \\
\text{(TOLABELED)} \\
\frac{\langle \Sigma, pc, e \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle}{\langle \Sigma, pc, \mathbf{toLabeled}(e) \rangle \Downarrow^\theta \langle \Sigma', pc, \mathbf{Labeled } pc' v \rangle} \\
\text{(UNLABEL)} \\
\frac{e \Downarrow^\theta \mathbf{Labeled } \ell v}{\langle \Sigma, pc, \mathbf{unlabel}(e) \rangle \Downarrow^\theta \langle \Sigma, pc \sqcup \ell, v \rangle} \\
\text{(LABELOF)} \\
\frac{e \Downarrow^\theta \mathbf{Labeled } \ell v}{\langle \Sigma, pc, \mathbf{labelOf}(e) \rangle \Downarrow^\theta \langle \Sigma, pc \sqcup \ell, \ell \rangle} \\
\text{(GETLABEL)} \quad \text{(TAINT)} \\
\langle \Sigma, pc, \mathbf{getLabel} \rangle \Downarrow^\theta \langle \Sigma, pc, pc \rangle \quad \frac{e \Downarrow^\theta \ell}{\langle \Sigma, pc, \mathbf{taint}(e) \rangle \Downarrow^\theta \langle \Sigma, pc \sqcup \ell, () \rangle} \\
\text{(b) Thunk semantics: } \langle \Sigma, pc, t \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle. \\
\hline
\end{array}$$

Figure 3.3: Big-step semantics for λ^{dCG} .

the program counter label in the program configuration, so that the primitives of the **LIO** monad can control it explicitly. In technical terms the program counter is said to be *flow-sensitive*, i.e., it may assume different values in the final configuration depending on the control flow of the program.³

Like λ^{dFG} , the operational semantics of λ^{dCG} consists of a security monitor that fully evaluates secure programs but prevents the execution of insecure programs and similarly enforces *termination-insensitive* non-interference (Theorem 3). The big-step operational semantics of λ^{dCG} is structured in two parts: (i) a straightforward call-by-value side-effect-free semantics for pure expressions (Figure 3.2), and (ii) a top-level security monitor for monadic programs (Figure 3.3). The semantics of the security monitor is further split into two mutually recursive reduction relations, one for arbitrary expressions (Figure 3.3a) and one specific to thunks (Figure 3.3b). These constitute the *forcing* semantics of the monad, which reduces a thunk to a pure value and perform side-effects. In particular, given the initial store Σ , program counter label pc , expression e of type **LIO** τ for some type τ and input values θ (which may or may not be labeled), the monitor executes the program, i.e., $\langle \Sigma, pc, e \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle$ and gives an updated store Σ' , updated program counter pc' and a final value v of type τ , which also might not be labeled. The execution starts with rule [FORCE], which reduces the pure expression to a thunk closure, i.e., (t, θ') and then forces the thunk t in its environment θ' with the thunk semantics.

The pure semantics is fairly standard—we report some selected rules in Figure 3.2 for comparison with λ^{dFG} . A pure reduction, written $e \Downarrow^\theta v$, evaluates an expression e with an appropriate environment θ to a pure value v . Notice that, unlike λ^{dFG} , all reduction rules of the pure semantics ignore security, even those that affect the control flow of the program, like rules [APP], [CASE₁], and [CASE₂]: they do not feature the program counter label or label annotations. This is because these reductions are *pure*—they cannot perform side-effects and so leak sensitive data—and thus are inherently secure and need

³In contrast, we consider λ^{dFG} 's program counter *flow-insensitive* because it is part of the evaluation judgment and its value changes only inside nested judgments.

not to be monitored (Vassena *et al.*, 2017).⁴ For example, since pure programs do not have access to the store, they cannot leak through *implicit flows*, which are then a concern only for monadic programs.

How does the semantics prevent pure programs from performing side-effects, without getting the program stuck and raising a false alarm? If the pure evaluation reaches a side-effectful computation, i.e., *think* t , it *suspends* the computation by creating a *think* closure that captures the current environment θ (see rule [THUNK]).⁵ Notice that *think closures* and *function closures* are distinct values created by different rules, [THUNK] and [FUN] respectively.⁶ Function application succeeds only when the function evaluates to a function closure (rule [APP]). In the *think* semantics, rule [RETURN] evaluates a pure value embedded in the monad via **return**(\cdot) and leaves the state unchanged, while rule [BIND] executes the first computation with the forcing semantics, binds the result in the environment i.e., $\theta[x \mapsto v_1]$, passes it on to the second computation together with the updated state and returns the final result and state. Rule [UNLABEL] is interesting. Following the *floating-label* principle, it returns the value wrapped inside the labeled value, i.e., v , and raises the program counter with its label, i.e., $pc \sqcup \ell$, to reflect the fact that new data at security level ℓ is now in scope.

Floating-label based coarse-grained IFC systems like **LIO** suffer from the *label creep* problem, which occurs when the program counter gets over-tainted, e.g., because too many secrets have been unlabeled, to the point that no useful further computation can be performed. Primitive **toLabeled**(\cdot) provides a mechanism to address this problem by (i) creating a separate context where some sensitive computation can take place and (ii) restoring the original program counter label

⁴The strict separation between side-effect-free and side-effectful code is a distinctive feature of coarse-grained IFC, which is crucial to lightweight approaches to enforce security via software libraries (Russo *et al.*, 2009; Buiras *et al.*, 2015; Russo, 2015; Stefan *et al.*, 2012)

⁵Notice that *type preservation* for the pure semantics preserves types *exactly* i.e., if $\Gamma \vdash e : \tau$, $e \Downarrow^\theta v$ and $\vdash \theta : \Gamma$, then $\vdash v : \tau$, which reflects the *suspending* behavior for the monadic type **LIO** τ . In contrast, type preservation for the *think* and *forcing* semantics assumes that the expression (resp. *think*) has a monadic type, i.e., **LIO** τ for some type τ , and guarantees that the final value has type τ .

⁶It would have also been possible to define *think* values in terms of function closures using explicit suspension and an opaque constructor wrapper, e.g., **LIO** ($_ . t, \theta$).

afterwards. Rule [TOLABELED] formalizes this idea. Notice that the result of the nested sensitive computation, i.e., v , cannot be simply returned to the lower context—that would be a leak, so **toLabeled**(\cdot) wraps that piece of information in a labeled value protected by the final program counter of the sensitive computation, i.e., **Labeled** $pc' v$.⁷ Furthermore, notice that pc' , the label that tags the result v , is as sensitive as the result itself because the final program counter depends on all the **unlabel**(\cdot) operations performed to compute the result. This motivates why primitive **labelOf**(\cdot) does not simply project the label from a labeled value, but additionally taints the program counter with the label itself in rule [LABELOF]—a label in a labeled value has sensitivity equal to the label itself, thus the program counter label rises to accommodate reading new sensitive data.

Lastly, rule [GETLABEL] returns the value of the program counter, which does not rise (because $pc \sqcup pc = pc$), and rule [TAINT] simply taints the program counter with the given label and returns unit (this primitive matches the functionality of **taint**(\cdot) in λ^{dFG}). Note that, in λ^{dCG} , **taint**(\cdot) takes *only* the label with which the program counter must be tainted whereas, in λ^{dFG} , it additionally requires the expression that must be evaluated in the tainted environment. This difference highlights the *flow-sensitive* nature of the program counter label in λ^{dCG} .

3.1.1 References

λ^{dCG} features *flow-insensitive* labeled references similar to λ^{dFG} and allows programs to create, read, update and inspect the label inside the **LIO** monad (see Figure 3.4). The API of these primitives takes explicitly labeled values as arguments, by making explicit at the type level, the tagging that occurs in memory, which was left implicit in previous work (Stefan *et al.*, 2017). Rule [NEW] creates a reference labeled with the same label annotation as that of the labeled value it

⁷Stefan *et al.* (2017) have proposed an alternative flow-insensitive primitive, i.e., **toLabeled**(ℓ, e), which labels the result with the user-assigned label ℓ . In λ^{dCG} , we use primitive **toLabeled**(e) because its flow-sensitive labeling semantics aligns exactly with the semantics of λ^{dFG} , thus simplifying the translation between the languages.

$$\begin{array}{c}
\text{(NEW)} \\
\frac{e \Downarrow^\theta \mathbf{Labeled} \ell v \quad pc \sqsubseteq \ell \quad n = |\Sigma(\ell)|}{\langle \Sigma, pc, \mathbf{new}(e) \rangle \Downarrow^\theta \langle \Sigma[\ell \mapsto \Sigma(\ell)[n \mapsto v]], pc, n_\ell \rangle} \\
\\
\text{(READ)} \\
\frac{e \Downarrow^\theta n_\ell \quad \Sigma(\ell)[n] = v}{\langle \Sigma, pc, !e \rangle \Downarrow^\theta \langle \Sigma, pc \sqcup \ell, v \rangle} \\
\\
\text{(WRITE)} \\
\frac{e_1 \Downarrow^\theta n_{\ell_1} \quad e_2 \Downarrow^\theta \mathbf{Labeled} \ell_2 v \quad \ell_2 \sqsubseteq \ell_1 \quad pc \sqsubseteq \ell_1}{\langle \Sigma, pc, e_1 := e_2 \rangle \Downarrow^\theta \langle \Sigma[\ell_1 \mapsto \Sigma(\ell_1)[n \mapsto v]], pc, () \rangle} \\
\\
\text{(LABELOFREF)} \\
\frac{e \Downarrow^\theta n_\ell}{\langle \Sigma, pc, \mathbf{labelOfRef}(e) \rangle \Downarrow^\theta \langle \Sigma, pc \sqcup \ell, \ell \rangle}
\end{array}$$

Figure 3.4: Big-step semantics for λ^{dCG} (references).

receives as an argument, and checks that $pc \sqsubseteq \ell$ in order to avoid implicit flows. Rule [READ] retrieves the content of the reference from the ℓ -labeled memory and returns it. Since this brings data at security level ℓ in scope, the program counter is tainted accordingly, i.e., $pc \sqcup \ell$. Rule [WRITE] performs security checks analogous to those in λ^{dFG} and updates the content of a given reference and rule [LABELOFREF] returns the label on a reference and taints the context accordingly.

We conclude this section by noting that the forcing and the thunk semantics of λ^{dCG} satisfy Property 6 (“the final value of the program counter label of any λ^{dCG} program is always at least as sensitive as the initial value”).

Property 6.

- If $\langle \Sigma, pc, e \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle$ then $pc \sqsubseteq pc'$.
- If $\langle \Sigma, pc, t \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle$ then $pc \sqsubseteq pc'$.

Proof. By mutual induction on the given evaluation derivations. □

$$\begin{array}{c}
\text{(Labeled}_L\text{)} \\
\frac{\ell \sqsubseteq L \quad v_1 \approx_L v_2}{\text{Labeled } \ell \ v_1 \approx_L \text{Labeled } \ell \ v_2} \\
\\
\text{(Labeled}_H\text{)} \\
\frac{\ell_1 \not\sqsubseteq L \quad \ell_2 \not\sqsubseteq L}{\text{Labeled } \ell_1 \ v_1 \approx_L \text{Labeled } \ell_2 \ v_2} \\
\\
\text{(INL)} \qquad \qquad \text{(INR)} \qquad \qquad \text{(F-CLOSURE)} \\
\frac{v_1 \approx_L v_2}{\text{inl}(v_1) \approx_L \text{inl}(v_2)} \quad \frac{v_1 \approx_L v_2}{\text{inr}(v_1) \approx_L \text{inr}(v_2)} \quad \frac{e_1 \equiv_\alpha e_2 \quad \theta_1 \approx_L \theta_2}{(e_1, \theta_1) \approx_L (e_2, \theta_2)} \\
\\
\text{(T-CLOSURE)} \qquad \qquad \text{(REF}_L\text{)} \qquad \qquad \text{(REF}_H\text{)} \\
\frac{t_1 \equiv_\alpha t_2 \quad \theta_1 \approx_L \theta_2}{(t_1, \theta_1) \approx_L (t_2, \theta_2)} \quad \frac{\ell \sqsubseteq L}{n^\ell \approx_L n^\ell} \quad \frac{\ell_1 \not\sqsubseteq L \quad \ell_2 \not\sqsubseteq L}{n_1^{\ell_1} \approx_L n_2^{\ell_2}} \\
\\
\text{(PC}_H\text{)} \\
\frac{\Sigma_1 \approx_L \Sigma_2 \quad pc_1 \not\sqsubseteq L \quad pc_2 \not\sqsubseteq L}{\langle \Sigma_1, pc_1, v_1 \rangle \approx_L \langle \Sigma_2, pc_2, v_2 \rangle} \\
\\
\text{(PC}_L\text{)} \\
\frac{\Sigma_1 \approx_L \Sigma_2 \quad pc \sqsubseteq L \quad v_1 \approx_L v_2}{\langle \Sigma_1, pc, v_1 \rangle \approx_L \langle \Sigma_2, pc, v_2 \rangle}
\end{array}$$

Figure 3.5: L -equivalence for λ^{dCG} values (selected rules) and configurations.

3.2 Security

We now prove that λ^{dCG} is secure, i.e., it satisfies *termination-insensitive non-interference*. The meaning of the security condition is intuitively similar to that presented in Section 2.2 for λ^{dFG} —when secret inputs are changed, terminating programs do not produce any publicly observable effect—and based on a similar indistinguishability relation.

3.2.1 L -Equivalence

Figure 3.5 presents the definition of L -equivalence for the interesting cases only. Firstly, L -equivalence for λ^{dCG} labeled values relates public and secret values analogously to λ^{dFG} values. Specifically, rule

[LBELED_L] relates public labeled values that share the *same* observable label, i.e., $\ell \sqsubseteq L$, and contain related values, i.e., $v_1 \approx_L v_2$, while rule [LBELED_H] relates secret labeled values, with arbitrary sensitivity labels not below L , i.e., $\ell_1 \not\sqsubseteq L$ and $\ell_2 \not\sqsubseteq L$, and contents. Secondly, L -equivalence relates standard (unlabeled) values homomorphically. For example, values of the sum type are related only as follows: $\mathbf{inl}(v_1) \approx_L \mathbf{inl}(v'_1)$ iff $v_1 \approx_L v'_1$ through rule [INL] and $\mathbf{inr}(v_2) \approx_L \mathbf{inr}(v'_2)$ iff $v_2 \approx_L v'_2$ through rule [INR], i.e., we do not provide any rule to relate values with different injections, just like in λ^{dFG} . Function and thunk closures are related by rules [F-CLOSURE] and [T-CLOSURE], respectively. In the rules the function and the monadic computations are α -equivalent and their environments are related, i.e., $\theta_1 \approx_L \theta_2$ iff $\text{dom}(\theta_1) \equiv \text{dom}(\theta_2)$ and $\forall x. \theta_1(x) \approx_L \theta_2(x)$. Labeled references, memories and stores are related by L -equivalence analogously to λ^{dFG} . Lastly, L -equivalence relates *initial* configurations with related stores, equal program counters and α -equivalent expressions (resp. thunks), i.e., $c_1 \approx_L c_2$ iff $c_1 = \langle \Sigma_1, pc_1, e_1 \rangle$, $c_2 = \langle \Sigma_2, pc_2, e_2 \rangle$, $\Sigma_1 \approx_L \Sigma_2$, $pc_1 \equiv pc_2$, and $e_1 \equiv_\alpha e_2$ (resp. $t_1 \equiv_\alpha t_2$ for thunks t_1 and t_2), and *final* configurations with related stores and (i) equal public program counter label, i.e., $pc \sqsubseteq L$, and related values through rule [PC_L], or (ii) arbitrary secret program counter labels, i.e., $pc_1 \not\sqsubseteq L$ and $pc_2 \not\sqsubseteq L$, and arbitrary values through rule [PC_H].

L -equivalence for λ^{dCG} is *reflexive*, *symmetric*, and *transitive*, similarly to λ^{dFG} (Property 2), and so these properties can be combine to derive a *Square Commutative Diagram for λ^{dCG} Stores* as well, analogous to Figure 2.6.

3.2.2 Termination-Insensitive Non-Interference

We now formally prove that also the security monitor of λ^{dCG} is secure, i.e., it enforces *termination-insensitive non-interference* (TINI). The proof technique is the same used for λ^{dFG} and similarly based on *store confinement* and *L -equivalence preservation*. However, since the semantics of the security monitor of λ^{dCG} is defined by two mutually recursive relations, i.e., the forcing and the thunk semantics, each lemma is stated as a pair of lemmas (one for each semantics relation), which are then proved by mutual induction.

Store Confinement. Store confinement ensures that programs running in a secret context cannot leak data implicitly through the labeled store. To prevent these leaks, the security monitor of λ^{dCG} aborts programs that attempt to write public memories in secret contexts, as these changes may depend on secret data and would be observable by the attacker. Rules [NEW] and [WRITE] enforce exactly this security policy by allowing programs to write memories only if they are labeled above the program counter label.

Lemma 3.1 (Store Confinement). For all program counter labels $pc \not\sqsubseteq L$, initial configurations c , and final configurations $c' = \langle \Sigma', pc', v \rangle$:

- If $c = \langle \Sigma, pc, e \rangle$ and $c \Downarrow^\theta c'$, then $\Sigma \approx_L \Sigma'$.
- If $c = \langle \Sigma, pc, t \rangle$ and $c \Downarrow^\theta c'$, then $\Sigma \approx_L \Sigma'$.

Proof. By mutual induction, using *reflexivity* in most base cases and *transitivity* in case [BIND], where the program counter label remains above the attacker's level in the intermediate configuration by Property 6. In the base cases [NEW] and [WRITE], the programs run with program counter label secret, i.e., $pc \not\sqsubseteq L$, and write a secret memory labeled ℓ above pc , i.e., $pc \sqsubseteq \ell$, and thus also not observable by the attacker, i.e., $\ell \not\sqsubseteq L$. \square

L -Equivalence Preservation. Next, we prove L -equivalence preservation in secret contexts, i.e., programs running with program counter label secret ($pc \not\sqsubseteq L$) cannot leak secret data implicitly through their final value or observable changes to the store. In λ^{dCG} , values are not intrinsically labeled like λ^{dFG} values, therefore the final values computed by these programs may not be explicitly labeled. If these values are unlabeled, how can we establish if they depend on secret data, or if they are L -equivalent? Luckily, we can safely approximate the sensitivity of these values with the program counter label.

The program counter label always represents an upper bound over the sensitivity of all data not explicitly labeled in a program, including its final value. This is precisely why λ^{dCG} stores it in the program

configuration, so that the **LIO** monad can control it through the floating-label mechanism explained above. Therefore, the fact that the program counter label is secret in a final configuration indicates that the program has unlabeled secret data and thus the final value *may* depend on secrets. This approximation is sound, but also very *conservative*—in a secret context, the result of a program must always be considered secret even if the program has not actually used any secret data to compute it—and central to the design of the *coarse-grained* IFC approach embodied by λ^{dCG} .⁸

The L -equivalence relation for final configurations reflects this reading of the program counter label for unlabeled values. Specifically, if both program counter labels are secret, then rule $[PC_H]$ (Figure 3.5) simply ignores the values in the configurations (because they *may* depend on secret data) and only requires the stores to be L -equivalent.⁹ Then, to prove L -equivalence preservation in secret contexts, we simply observe that the program counter label can only increase during program execution (Property 6), thus programs that start in secret contexts must also necessarily terminate in secret contexts. Therefore, we can apply rule $[PC_H]$ provided that the final stores are L -equivalent, i.e., if we can show that programs cannot leak implicitly through the public memories of the store. To do that, we follow the same proof technique used for λ^{dFG} , i.e., we first prove that the final stores are L -equivalent to the initial stores by applying *store confinement* (Lemma 3.1) to each execution and then derive L -equivalence of the final stores by constructing a *square commutative diagram* (Figure 2.6).

Lemma 3.2 (*L -Equivalence Preservation in Secret Contexts*). For all program counter labels $pc_1 \not\sqsubseteq L$ and $pc_2 \not\sqsubseteq L$, and stores $\Sigma_1 \approx_L \Sigma_2$:

- If $\langle \Sigma_1, pc_1, e_1 \rangle \Downarrow^{\theta_1} c_1$ and $\langle \Sigma_2, pc_2, e_2 \rangle \Downarrow^{\theta_2} c_2$, then $c_1 \approx_L c_2$.
- If $\langle \Sigma_1, pc_1, t_1 \rangle \Downarrow^{\theta_1} c_1$ and $\langle \Sigma_2, pc_2, t_2 \rangle \Downarrow^{\theta_2} c_2$, then $c_1 \approx_L c_2$.

⁸Due to this conservative behavior, coarse-grained IFC languages like λ^{dCG} may seem inherently limited, compared to fine-grained languages like λ^{dFG} , which track data-dependencies more precisely, i.e., the result gets labeled secret *only if* secret data has been used to compute its value. In Section 5, we prove that λ^{dCG} can track data-dependencies as precisely as λ^{dFG} .

⁹In particular, rule $[PC_H]$ accepts arbitrary final values just like rule $[Labeled_H]$ for explicitly labeled values.

Proof. Assume $pc_1 \not\sqsubseteq L$, $pc_2 \not\sqsubseteq L$, $\Sigma_1 \approx_L \Sigma_2$, and let $c_1 = \langle \Sigma'_1, pc'_1, v_1 \rangle$ and $c_2 = \langle \Sigma'_2, pc'_2, v_2 \rangle$. First, we apply *store confinement* (Lemma 3.1) to each big-step reduction and obtain $\Sigma_1 \approx_L \Sigma'_1$ and $\Sigma_2 \approx_L \Sigma'_2$. These are then combined with the assumption $\Sigma_1 \approx_L \Sigma_2$ to form the commutative square diagram for λ^{dCG} stores (Figure 2.6), which gives $\Sigma'_1 \approx_L \Sigma'_2$. Then, we observe that the program counter labels in the final configurations are *secret*, i.e., $pc'_1 \not\sqsubseteq L$ and $pc'_2 \not\sqsubseteq L$ by Property 6, and therefore we have $c_1 \approx_L c_2$ by rule [PC_H]. \square

Now, we consider L -equivalence preservation in public contexts. Since λ^{dCG} separates *pure computations* from *monadic computations* in different semantics judgments, we first need to prove that L -equivalence is preserved also by the pure semantics. Notice that this separation is not only beneficial for security, as explained above, but simplifies the security analysis as well. For example, pure computations cannot inspect secret data in a public context, as this operation requires performing a side-effect. This is because, in public contexts, secret data is *explicitly* protected through the **Labeled** type, which prevents programs from inspecting secrets directly as they must first *extract* them via **unlabel**(\cdot): **Labeled** $\tau \rightarrow \mathbf{LIO} \tau$. The monadic type of this thunk indicates that this computation may perform side-effects (i.e., tainting the program counter label), which crucially, can be performed only by the security monitor, in the thunk semantics.¹⁰ As a result, the execution of pure programs in public contexts cannot depend on secret data—they simply cannot access secrets in the first place—and so their (unlabeled) results are always indistinguishable by the attacker.

Lemma 3.3 (Pure L -Equivalence Preservation). For all expressions $e_1 \equiv_\alpha e_2$ and environments $\theta_1 \approx_L \theta_2$, if $e_1 \Downarrow^{\theta_1} v_1$ and $e_2 \Downarrow^{\theta_2} v_2$, then $v_1 \approx_L v_2$.

Proof. By induction on the big-step reductions, which *must* always step according to the same rule. In the spurious cases, e.g., when expression **case**($e, x.e_1, x.e_2$) steps through rules [CASE₁] and [CASE₂], we show a contradiction. In particular, assume that the executions follow different paths, e.g., we have reductions $e \Downarrow^{\theta_1} \mathbf{inl}(v_1)$ and $e \Downarrow^{\theta_2} \mathbf{inr}(v_2)$ for the

¹⁰The pure semantics simply *suspends* the evaluation of thunks via rule [THUNK].

scrutinee e . Then, we apply the induction hypothesis to these reductions and obtain a proof for $\mathbf{inl}(v_1) \approx_L \mathbf{inr}(v_2)$. But this is impossible: L -equivalence relates values of sum types only if they have the same injection, i.e., rules [INL] and [INR] in Figure 3.5. Therefore, the two executions must evaluate the scrutinee to the same injection and follow the same path. \square

Similarly, the fact that λ^{dCG} requires programs to explicitly unlabel secret data simplifies also the analysis of *implicit flows* in monadic computations. This is because the control flow of λ^{dCG} programs can only depend on unlabeled data, whose sensitivity is *coarsely* approximated by the program counter label, as explained above. Since the program counter label only gets tainted in response to specific monadic actions (e.g., $\mathbf{unlabel}(e)$), and not by regular control-flow construct (e.g., $\mathbf{case}(e, x.e_1, x.e_2)$), the evaluation of pure expressions *cannot* cause implicit information flows. In particular, by encapsulating secret data in the **Labeled** data type, λ^{dCG} makes all data dependencies—even those implicit in the program control flow—*explicit* through the program counter label.

Lemma 3.4 (*L-Equivalence Preservation in Public Contexts*). For all public program counter labels $pc \sqsubseteq L$, environments $\theta_1 \approx_L \theta_2$, and stores $\Sigma_1 \approx_L \Sigma_2$:

- If $e_1 \equiv_\alpha e_2$, $\langle \Sigma_1, pc, e_1 \rangle \Downarrow^{\theta_1} c_1$, $\langle \Sigma_2, pc, e_2 \rangle \Downarrow^{\theta_2} c_2$, then $c_1 \approx_L c_2$;
- If $t_1 \equiv_\alpha t_2$, $\langle \Sigma_1, pc, t_1 \rangle \Downarrow^{\theta_1} c_1$, $\langle \Sigma_2, pc, t_2 \rangle \Downarrow^{\theta_2} c_2$, then $c_1 \approx_L c_2$.

Proof. The two lemmas are proved mutually, by simultaneous induction on the big-step reductions. The *forcing semantics* has only a single rule, i.e., [FORCE] in Figure 3.3a. Therefore, for the first lemma, we simply apply *Pure L-Equivalence Preservation* (Lemma 3.3) to the pure reductions $e_1 \Downarrow^{\theta_1} (t_1, \theta'_1)$ and $e_2 \Downarrow^{\theta_2} (t_2, \theta'_2)$, which gives $(t_1, \theta'_1) \approx_L (t_2, \theta'_2)$, i.e., $t_1 \equiv_\alpha t_2$ and $\theta'_1 \approx_L \theta'_2$, so we complete the proof by mutual induction with the lemma for thunks. For the second lemma, we observe that the thunks are α -equivalent, i.e., $t_1 \equiv_\alpha t_2$, therefore their reductions always step according to the same rule. For example, in case [UNLABEL], the labeled values that get unlabeled are L -equivalent,

i.e., **Labeled** $\ell_1 v_1 \approx_L$ **Labeled** $\ell_2 v_2$, by Lemma 3.3 applied to the pure reductions, and then we proceed by cases on the L -equivalence judgment. (Notice that the store does not change in rule [UNLABEL], so the final stores are L -equivalent by assumption). In case [LABELED $_L$], the two values are labeled public, i.e., $\ell_1 = \ell_2 \sqsubseteq L$, and the values are related, i.e., $v_1 \approx_L v_2$. Therefore, the program counter label in the final configurations remain observable by the attacker, i.e., $pc \sqcup \ell_1 \sqsubseteq L$ since $pc \sqsubseteq L$ and $\ell_1 \sqsubseteq L$, and the final configurations are L -equivalent by rule [PC $_L$]. In case [LABELED $_H$], the two values are labeled secret, i.e., $\ell_1 \not\sqsubseteq L$ and $\ell_2 \not\sqsubseteq L$, and the program counter label in the final configurations are secret, i.e., $pc \sqcup \ell_1 \not\sqsubseteq L$ and $pc \sqcup \ell_2 \not\sqsubseteq L$, and so the configurations are L -equivalent by rule [PC $_H$]. \square

Finally, we prove termination-insensitive non-interference for λ^{dCG} by combining the L -equivalence preservation lemmas above.

Theorem 3 (λ^{dCG} -TINI). If $c_1 \Downarrow^{\theta_1} c'_1$, $c_2 \Downarrow^{\theta_2} c'_2$, $\theta_1 \approx_L \theta_2$ and $c_1 \approx_L c_2$ then $c'_1 \approx_L c'_2$.

3.3 Flow-Sensitive References

We now continue our exploration of coarse-grained IFC by adding flow-sensitive references to λ^{dCG} and then showing that the extended language is secure. This extension is in many ways analogous to the extension presented in Section 2.3 for λ^{dFG} . Therefore, we focus mainly on the (apparently different) security checks performed by the security monitor of λ^{dCG} and how they relate to those performed by the monitor of λ^{dFG} . In Section 3.3.3, we formally establish the security guarantees of λ^{dCG} extended with flow-sensitive references. Though the security analysis of λ^{dCG} is also complicated by bijections like λ^{dFG} , we find that the separation between pure and monadic computations in λ^{dCG} limits the extra complexity of the analysis only to the monadic fragment of the semantics. In particular, this extension does not affect the pure fragment of the semantics and so adapting this part of the analysis is straightforward. Finally, we leverage our mechanized proof scripts to compare and analyze the security proofs of λ^{dFG} and λ^{dCG} and discuss our findings in Section 4.

3.3.1 Syntax

Figure 3.6 introduces the new syntactic constructs and semantics rules for flow-sensitive references. First, we annotate the reference type constructor with a tag $s \in \{I, S\}$, which indicates the sensitivity of a reference term. A plain address $n : \mathbf{Ref} S \tau$ represents a flow-sensitive reference pointing to a value of type τ , stored in the n -th cell of the heap. Notice that these references are *not* annotated with a label representing the sensitivity of their content, like flow-insensitive references. This is because flow-sensitive references are, by design, allowed to store data at different security levels, i.e., the label of these references can *change* throughout program execution. Thus, the references themselves are *unlabeled* and, instead, the security monitor *explicitly* labels their content in the heap when they are created and updated (see below).¹¹ To ensure that values in the heap are always labeled, heaps μ are syntactically defined as lists of explicitly labeled values $\mathbf{Labeled} \ell v$, whose label ℓ represents the sensitivity of value v and, at the same time, the label of any reference pointing to it. The heap is stored in program configurations $\langle \Sigma, \mu, pc, e \rangle$, where programs can access it through the same thunk constructs used for flow-insensitive references, i.e., $\mathbf{new}(e)$, $e_1 := e_2$, $!e$, and $\mathbf{labelOfRef}(e)$. As we explain next, these operations must be regulated by the security monitor of λ^{dCG} to enforce security.

3.3.2 Dynamics

Figure 3.6 extends the thunk semantics with new rules that allow programs to access the heap through flow-sensitive references without leaking data. Rule [NEW-FS] creates a new reference by allocating a new cell in the heap at fresh address $n = |\mu|$, which is initialized with the given labeled value argument, i.e., $\mu[n \mapsto \mathbf{Labeled} \ell v]$. Importantly, the rule requires the label of the value to be above the program counter label, i.e., $pc \sqsubseteq \ell$, to avoid leaks.¹² Without this constraint, a program

¹¹In contrast, flow-insensitive references are annotated with a *fixed* label, e.g., ℓ for $n_\ell : \mathbf{Ref} I \tau$, which represents (an upper bound over) the sensitivity of its content. Since this label does not change, the content of the reference can be stored directly in the memory labeled ℓ , without being explicitly labeled.

¹²This restriction, known as *no write-down* (Bell and La Padula, 1976), is a core design principles of several static IFC libraries (Russo, 2015; Vassena *et al.*, 2017),

Types: $\tau ::= \dots \mid \mathbf{Ref} \ s \ \tau$
 Values: $v ::= \dots \mid n$
 Heap: $\mu ::= [] \mid \mathbf{Labeled} \ \ell \ v : \mu$
 Configuration: $c ::= \langle \Sigma, \mu, pc, e \rangle$

(a) Syntax.

(NEW-FS)

$$\frac{e \Downarrow^\theta \ \mathbf{Labeled} \ \ell \ v \quad pc \sqsubseteq \ell \quad n = |\mu| \quad \mu' = \mu[n \mapsto \mathbf{Labeled} \ \ell \ v]}{\langle \Sigma, \mu, pc, \mathbf{new}(e) \rangle \Downarrow^\theta \langle \Sigma, \mu', pc, n \rangle}$$

(READ-FS)

$$\frac{e \Downarrow^\theta \ n \quad \mu[n] = \mathbf{Labeled} \ \ell \ v}{\langle \Sigma, \mu, pc, !e \rangle \Downarrow^\theta \langle \Sigma, \mu, pc \sqcup \ell, v \rangle}$$

(LABELOFREF-FS)

$$\frac{e \Downarrow^\theta \ n \quad \mu[n] = \mathbf{Labeled} \ \ell \ _}{\langle \Sigma, \mu, pc, \mathbf{labelOfRef}(e) \rangle \Downarrow^\theta \langle \Sigma, \mu, pc \sqcup \ell, \ell \rangle}$$

(WRITE-FS)

$$\frac{e_1 \Downarrow^\theta \ n \quad e_2 \Downarrow^\theta \ \mathbf{Labeled} \ \ell' \ v \quad \mu[n] = \mathbf{Labeled} \ \ell \ _ \quad pc \sqsubseteq \ell \quad \mu' = \mu[n \mapsto \mathbf{Labeled} \ (pc \sqcup \ell') \ v]}{\langle \Sigma, \mu, pc, e_1 := e_2 \rangle \Downarrow^\theta \langle \Sigma, \mu', pc, () \rangle}$$

(b) Dynamics. The shaded constraint corresponds to the *no-sensitive upgrade* security check in λ^{dFG} .

Figure 3.6: λ^{dCG} extended with flow-sensitive references.

could create a public reference in a secret context, which constitutes a leak. Compare this rule with the corresponding rule for λ^{dFG} , i.e., rule [NEW-FS] in Figure 2.8b. The security monitor of λ^{dFG} does not check that the label of the value is above the program counter label—in fact the rule does not contain any security check at all! Why must λ^{dCG} include that check and instead λ^{dFG} can skip it? This is because that check is *redundant* in λ^{dFG} : the constraint $pc \sqsubseteq \ell$ performed by the security monitor of λ^{dCG} is an invariant (Property 1) of the semantics of λ^{dFG} . Intuitively, λ^{dFG} programs cannot create public references in secret contexts because the result of a computation is always labeled above the current program counter label. Since λ^{dCG} does not enjoy this invariant, i.e., computations can return values labeled public even in secret contexts, that check *must* be included in all the rules that perform write side-effects (i.e., also in rules [NEW], [WRITE] for flow-insensitive references in Figure 3.4) to avoid leaks.

Rule [READ-FS] reads a flow-sensitive reference by retrieving its content from the heap, i.e., $\mu[n] = \mathbf{Labeled} \ell v$, and returning its value v . Importantly, the rule taints the program counter label, i.e., $pc \sqcup \ell$, to indicate that data at security level ℓ is now in scope. Rule [LABELOFREF-FS] is similar, but returns the label of the reference, i.e., the same label ℓ that also annotates its content, and so it taints the program counter label with the label itself, i.e., $pc \sqcup \ell$.

Lastly, rule [WRITE-FS] updates the content of a flow-sensitive reference n with a new value $\mathbf{Labeled} \ell' v$, which replaces the current value stored in the heap, i.e., $\mu[n] = \mathbf{Labeled} \ell _$. Notice that the new value can be less or *even more* sensitive than the old value, i.e., labels ℓ and ℓ' can be in any relation. This is exactly why flow-sensitive references are, in general, more permissive than flow-insensitive references also in λ^{dCG} .¹³

In the rule, the constraint $pc \sqsubseteq \ell$ is analogous to the NSU check for λ^{dFG} . (For convenience, these equivalent checks are shaded in Figure 2.8b and 3.6b). Intuitively, this check allows a reference update only if

which was not identified as such in dynamic IFC libraries (Stefan *et al.*, 2011; Stefan *et al.*, 2017).

¹³In particular, Example 2.1 can be adapted to λ^{dCG} as well, i.e., program $r \leftarrow \mathbf{new}(p); r := s; !r$ is aborted with flow-insensitive references, but accepted with flow-sensitive references.

the decision to update *that* reference depends on data less sensitive than the reference itself.¹⁴ Though the NSU checks in λ^{dFG} and λ^{dCG} protect against the same type of implicit data leaks, they involve different labels, i.e., the intrinsic label of the reference in λ^{dFG} and the program counter label in λ^{dCG} . In λ^{dFG} , the sensitivity of the reference is represented by its intrinsic label, which can then be used directly in the NSU check. In contrast, λ^{dCG} does not feature intrinsic labels, therefore, we use the program counter label as a coarse, but *sound*, approximation of the sensitivity of the reference, hence the NSU check $pc \sqsubseteq \ell$. Notice that when the new value is written in the heap, the rule taints its label with the program counter label, i.e., $\mu[n \mapsto \mathbf{Labeled}(pc \sqcup \ell') v]$, for the same reason. Intuitively, the sensitivity of the content is determined by the new value, which is explicitly labeled ℓ' , and by the sensitivity of the (unlabeled) reference, approximated by the program counter label pc , i.e., the new content has sensitivity at most $pc \sqcup \ell'$.¹⁵

3.3.3 Security

The security analysis of λ^{dCG} extended with flow-sensitive references is very much similar to the analysis presented in Section 2.3.3 for λ^{dFG} . In particular, the security monitor of λ^{dCG} allocates both public and secret data in the same linear heap and so secret-dependent allocations can influence the addresses of subsequent public references, just like in λ^{dFG} . Although this can create a dependency between concrete heap addresses and secret data, security is not at stake because references are *opaque* in λ^{dCG} as well. To formally show that, we need to adapt the *L*-equivalence relation, and consequently the security analysis, to use a *bijection* to reason about corresponding references with secret-dependent, yet indistinguishable, heap addresses.

***L*-Equivalence up to Bijection.** Formally, we redefine *L*-equivalence as a relation $\approx_L^\beta \subseteq \text{Value} \times \text{Value}$ and add the bijection β to \approx_L in all the rules previously defined for λ^{dCG} . These rules do not use the bijection,

¹⁴ λ^{dCG} would be insecure without this constraint. In particular, programs could leak data through secret-dependent reference updates, similarly to Example 2.2.

¹⁵In addition, by tainting the label ℓ' with the program counter label pc , the rule automatically respects the *no write-down* restriction, i.e., $pc \sqsubseteq pc \sqcup \ell'$, hence no explicit security check is needed like in rules [NEW], [WRITE], and [NEW-FS].

which is needed instead only in the rules for flow-sensitive references and heaps. In particular, two flow-sensitive references are indistinguishable only if their heap addresses are matched by the bijection, i.e., we add the following new rule:

$$\begin{array}{c} \text{(REF-FS)} \\ \frac{(n_1, n_2) \in \beta}{n_1 \approx_L^\beta n_2} \end{array}$$

Similarly, in the definition of L -equivalence for heaps, the addresses related by the bijection identify corresponding heap cells, which must be recursively related, as they can be read through related references.¹⁶

Definition 3 (Heap L -equivalence). Two heaps μ_1 and μ_2 are L -equivalent up to bijection β , written $\mu_1 \approx_L^\beta \mu_2$, if and only if:

1. $\text{dom}(\beta) \subseteq \{0, \dots, |\mu_1| - 1\}$,
2. $\text{rng}(\beta) \subseteq \{0, \dots, |\mu_2| - 1\}$, and
3. For all addresses n_1 and n_2 , if $(n_1, n_2) \in \beta$, then $\mu_1[n_1] \approx_L^\beta \mu_2[n_2]$.

In the definition above, only the rules for labeled values can relate heap cells $\mu_1[n_1] \approx_L^\beta \mu_2[n_2]$. This is because heaps are syntactically defined as a list of explicitly labeled values in Figure 3.6a, i.e., $\mu_1[n_1] = \mathbf{Labeled} \ell_1 v_1$ for some label ℓ_1 and value v_1 and similarly $\mu_2[n_2] = \mathbf{Labeled} \ell_2 v_2$, and the only rules that give $\mathbf{Labeled} \ell_1 v_1 \approx_L^\beta \mathbf{Labeled} \ell_2 v_2$ are [Labeled $_L$] and [Labeled $_H$] in Figure 3.5.

Similarly to λ^{dFG} , the L -equivalence relation up to bijection for λ^{dCG} satisfies *restricted reflexivity*, *symmetricity*, *transitivity*, and *weakening* (i.e., Property 5), and so we can construct square commutative diagrams for stores and heaps (i.e., Lemma 2.7) in λ^{dCG} as well. Notice that reflexivity is restricted only to *valid* program constructs, i.e., values, environments, memories, stores, heaps, and configurations that are *free* of dangling flow-sensitive references, just like it is in λ^{dFG} and for the same technical reasons. These side conditions are formalized

¹⁶Notice that these addresses are valid in the heap thanks to the first two side-conditions, which ensure that the domain and range of the bijection are compatible with the address space of the heaps. See also the explanation of Definition 2.

by straightforward judgments, e.g., $n \vdash \mathbf{Valid}(\theta)$ and $n \vdash \mathbf{Valid}(v)$ indicating that the references contained in value v and environment θ are valid in a heap of size n , which are mutually and recursively defined like in Figure 2.9. In these judgments, the size parameter n is instantiated by the top-level judgment for program inputs and outputs with the size of the current heap μ , i.e., $n = |\mu|$ as shown in Figure 3.7. The fact that the security lemmas of λ^{dCG} assume valid configurations does not actually weaken the security guarantees of the language: these judgments are naturally preserved by the pure and monadic semantics of λ^{dCG} .

$$\begin{array}{c}
 \text{(VALID-INPUTS)} \\
 \frac{n = |\mu| \quad n \vdash \mathbf{Valid}(\Sigma) \quad c = \langle \Sigma, \mu, pc, e \rangle \quad n \vdash \mathbf{Valid}(\mu) \quad n \vdash \mathbf{Valid}(\theta)}{\vdash \mathbf{Valid}(c, \theta)} \\
 \\
 \text{(VALID-OUTPUTS)} \\
 \frac{n = |\mu| \quad n \vdash \mathbf{Valid}(\Sigma) \quad c = \langle \Sigma, \mu, pc, v \rangle \quad n \vdash \mathbf{Valid}(\mu) \quad n \vdash \mathbf{Valid}(v)}{\vdash \mathbf{Valid}(c)}
 \end{array}$$

Figure 3.7: Judgments for valid program inputs and outputs in λ^{dCG} .

Property 7 (Valid Invariant).

1. If $e \Downarrow^\theta v$ and $n \vdash \mathbf{Valid}(\theta)$, then $n \vdash \mathbf{Valid}(v)$.
2. If $c \Downarrow^\theta c'$ and $\vdash \mathbf{Valid}(c, \theta)$, then $\vdash \mathbf{Valid}(c')$.

Termination-Insensitive Non-Interference. We now prove that λ^{dCG} extended with flow-sensitive references enforces termination-insensitive non-interference. This result is also based on two lemmas, i.e., *store and heap confinement* and *L-equivalence preservation*, which are complicated by the fact that *L-equivalence* is defined up to a bijection, similarly to λ^{dFG} . These lemmas rely on bijections to relate flow-sensitive references allocated in public contexts: these references are observable

by the attacker, but their heap addresses may differ due to previous, secret-dependent allocations. In contrast, references allocated in secret contexts cannot be observed by the attacker and must be ignored by the bijection. For example, in the store and heap confinement lemma, we show that initial and final stores and heaps are L -equivalent, for programs executed in secret contexts. Which bijection should we use to relate the addresses of flow-sensitive references in this lemma? Intuitively, the addresses of references already allocated in the initial configuration remain unchanged in the final configuration (because references cannot be tempered in λ^{dCG}), so they can be related by the *identity bijection*. Furthermore, since references allocated in a secret context are not observable by the attacker, the lemma shows L -equivalence up to a identity bijection that ignores any new heap allocation performed by the program, i.e., up to the identity bijection *restricted to the domain of the initial heap*.

Lemma 3.5 (Store and Heap Confinement). For all program counter labels $pc \not\sqsubseteq L$, valid initial configurations and environments $\vdash \mathbf{Valid}(c, \theta)$, and final configurations $c' = \langle \Sigma', \mu', pc', v \rangle$:

- If $c = \langle \Sigma, \mu, pc, e \rangle$ and $c \Downarrow^\theta c'$, then $\Sigma \approx_L^{|\mu|} \Sigma'$ and $\mu \approx_L^{|\mu|} \mu'$.
- If $c = \langle \Sigma, \mu, pc, t \rangle$ and $c \Downarrow^\theta c'$, then $\Sigma \approx_L^{|\mu|} \Sigma'$ and $\mu \approx_L^{|\mu|} \mu'$.

Proof. Analogous to Lemma 2.8. Since λ^{dCG} encapsulates side-effects in a monad, we need to explicitly propagate the valid judgment only in very few cases compared to λ^{dFG} . In particular, we only apply Property 7.1 in rule [FORCE] and Property 7.2 in rule [BIND]. \square

Next, we prove L -equivalence preservation in *secret contexts*. Since the programs considered in this lemma run in a secret context, the addresses of the flow-sensitive references allocated during these executions must be ignored by the bijection, i.e., the final configurations must be L -equivalent up to the *same bijection* used for the initial configurations. We prove that by constructing a square commutative diagram for heaps and stores, similar to Figure 2.10. (We explain the construction for

heaps, stores are treated analogously).¹⁷ In the figure, heaps μ_1 and μ_2 are from the initial configurations of the lemma, while μ'_1 and μ'_2 are the heaps from the final configurations. Since the initial configurations are L -equivalent, we have the vertical edge on the left, i.e., $\mu_1 \approx_L^\beta \mu_2$, by assumption. Instead, the horizontal edges of the square are obtained by applying store and heap confinement to each individual program execution. In particular, the lemma shows that the initial and final heaps in these executions are L -equivalent up to appropriate identity bijections, i.e., $\mu_1 \approx_L^{\iota_{|\mu_1|}} \mu'_1$ and $\mu_2 \approx_L^{\iota_{|\mu_2|}} \mu'_2$, which ignore new heap allocations, as explained above. Therefore, the bijection that relates the final heaps from the vertical edge on the right of the square can be simplified. Specifically, the identity bijections *cancel out* in the composition with the bijection from the initial L -equivalence relation, i.e., $\iota_{|\mu_2|} \circ \beta \circ \iota_{|\mu_1|}^{-1} = \beta$, and so the final heaps remain related up to the same bijection β used for the initial heaps, i.e., $\mu'_1 \approx_L^\beta \mu'_2$. In the following, we use dot-notation to access individual elements of a configuration, e.g., we write $c.pc$ to extract the program counter pc from the configuration $c = \langle -, -, pc, - \rangle$.

Lemma 3.6 (*L -Equivalence Preservation in Secret Contexts*). For all valid inputs $\vdash \mathbf{Valid}(c_1, \theta_1)$ and $\vdash \mathbf{Valid}(c_2, \theta_2)$ and bijections β , such that $c_1.pc \not\sqsubseteq L$, $c_2.pc \not\sqsubseteq L$, and $c_1 \approx_L^\beta c_2$, if $c_1 \Downarrow^{\theta_1} c'_1$ and $c_2 \Downarrow^{\theta_2} c'_2$, then $c'_1 \approx_L^\beta c'_2$.

We now adapt the proof of L -equivalence in *public contexts*. The references allocated by the program considered in this lemma are observable by the attacker and so we have to construct an appropriate bijection to relate their addresses in the final configurations. However, since these allocations occur at run time and may depend on the program inputs, we cannot, in general, predict the exact bijection needed to relate them. Therefore, the final bijection is *existentially quantified* in the lemma, so that the right bijection can be precisely constructed, depending on the program execution, in each case of the proof itself. Unfortunately, existential quantification also complicates the proof, which now involves

¹⁷Heaps and stores are very often treated homogeneously in proofs, therefore our mechanized proof scripts pairs them up in a *program state* helper data structure to shorten some proofs.

reasoning about terms that are L -equivalent, but up to arbitrary bijections, and so may not be combined together. To solve this issue, the lemma additionally requires to prove an ordering invariant about these bijections, i.e., that the final bijection *extends* the initial bijection. Intuitively, this extra property solves the issue described above because L -equivalent relations up to “smaller” bijections can be lifted to “larger” bijections via *weakening* (Property 5.4), and so combined together.

It is worth pointing out that this issue is not specific to λ^{dCG} : the security analysis of λ^{dFG} presents the same issue, which we solve in the same way in Section 2.4. However, only the analysis of the monadic fragment of the semantics of λ^{dCG} is affected by the issue described above: the final bijection does not need to be quantified when reasoning about the pure fragment of the semantics. This is because pure reductions cannot allocate new references (they do not even have access to a heap) and so the same bijection that relates their input values can also relate their output values. This simplifies our formal analysis, as previous lemmas and their proofs are largely unaffected by the addition of bijections. For example, to adapt Lemma 3.3, i.e., *Pure L-Equivalence Preservation*, for flow-sensitive references, we *only* need to add the bijection β to the L -equivalent relations in the statement of lemma: the proof requires no changes.

Lemma 3.7 (Pure L -Equivalence Preservation). For all expressions $e_1 \equiv_\alpha e_2$, bijections β , and environments $\theta_1 \approx_L^\beta \theta_2$, if $e_1 \Downarrow^{\theta_1} v_1$ and $e_2 \Downarrow^{\theta_2} v_2$, then $v_1 \approx_L^\beta v_2$.

The separation between pure and monadic semantics fragments of λ^{dCG} simplifies the L -equivalence preservation of monadic computations as well. This is because most of the rules of this fragment of the semantics only include pure reductions and so intermediate L -equivalent results can be combined directly, i.e., no weakening is required in most cases.

Lemma 3.8 (L -Equivalence Preservation in Public Contexts). For all valid initial inputs $\vdash \mathbf{Valid}(c_1, \theta_1)$ and $\vdash \mathbf{Valid}(c_2, \theta_2)$ and bijections β , such that $c_1.pc = c_2.pc \sqsubseteq L$, $c_1 \approx_L^\beta c_2$, and $\theta_1 \approx_L^\beta \theta_2$, if $c_1 \Downarrow^{\theta_1} c'_1$ and $c_2 \Downarrow^{\theta_2} c'_2$, then there exists an extended bijection $\beta' \supseteq \beta$ such that $c'_1 \approx_L^{\beta'} c'_2$.

Finally, we combine L -equivalence preservation in public and secret contexts and prove *termination-insensitive non-interference* for λ^{dCG} extended with flow-sensitive references.

Theorem 4 (λ^{dCG} -TINI with Bijections). For all valid inputs $\vdash \mathbf{Valid}(c_1, \theta_1)$ and $\vdash \mathbf{Valid}(c_2, \theta_2)$ and bijections β , such that $c_1 \approx_L^\beta c_2$, $\theta_1 \approx_L^\beta \theta_2$, if $c_1 \Downarrow^{\theta_1} c'_1$, and $c_2 \Downarrow^{\theta_2} c'_2$, then there exists an extended bijection $\beta' \supseteq \beta$, such that $c'_1 \approx_L^{\beta'} c'_2$.

Conclusion. At this point, we have formalized two calculi— λ^{dFG} and λ^{dCG} —that perform dynamic IFC at *fine* and *coarse* granularity, respectively. While they have some similarities, i.e., they are both functional languages that feature labeled annotated data, references and label introspection primitives, and ensure a termination-insensitive security condition, they also have striking differences. First and foremost, they differ in the number of label annotations—pervasive in λ^{dFG} and optional in λ^{dCG} —with significant implications for the programming model and usability. Second, they differ in the nature of the program counter, *flow-insensitive* in λ^{dFG} and *flow-sensitive* in λ^{dCG} . Third, they differ in the way they deal with side-effects— λ^{dCG} allows side-effectful computations exclusively inside the monad, while λ^{dFG} is *impure*, i.e., any λ^{dFG} expression can modify the state. This difference affects the effort required to implement a system that performs language-based fine- and coarse-grained dynamic IFC. In fact, several coarse-grained IFC languages (Schmitz *et al.*, 2018; Buiras *et al.*, 2015; Jaskelioff and Russo, 2011; Tsai *et al.*, 2007; Russo *et al.*, 2009; Russo, 2015) have been implemented as an embedded domain specific language (EDSL) in a Haskell library with little effort, exploiting the strict control that the host language provides on side-effects. Adapting an existing language to perform fine-grained IFC requires major engineering effort, because several components (all the way from the parser to the runtime system) must be adapted to be label-aware.

In the next section we discuss our verified artifacts of λ^{dFG} and λ^{dCG} and compare their mechanized proofs of non-interference. Then, in Section 5 and 6 we show that—despite their differences—these two calculi are, in fact, equally expressive.

4

Verified Artifacts

We now discuss our verified artifacts, in which we model λ^{dFG} and λ^{dCG} and provide machine-checked proofs of their security guarantees. We have formalized these languages using Agda (Norell, 2009; Bove *et al.*, 2009), a dependently typed functional language and an interactive proof assistant based on intuitionistic type theory. In our proof scripts, we embed the syntax, the type system, and the semantics of λ^{dFG} and λ^{dCG} into Agda data types, where we leverage dependent types to maintain additional assumptions about terms and judgments. For example, we use *well-typed syntax* and typed DeBruijn indexes to ensure that expressions, values, and environments are intrinsically well-scoped and well-typed (Abel *et al.*, 2019), which in turn let us define *type-preserving* big-step semantics judgments. In the following, we analyze our artifact and find that the security proofs for λ^{dFG} are longer (between 43% and 74%) than those for λ^{dCG} . These empirical results suggest that reasoning about the security of coarse-grained IFC languages is easier than for fine-grained languages.

Table 4.1: The table reports the size (Agda LOC excluding blank lines and comments) of the proof scripts that formalize λ^{dFG} and λ^{dCG} and their security proofs with flow-insensitive references (FI) and with also flow-sensitive references (FI + FS).

	FI		FI + FS	
	λ^{dFG}	λ^{dCG}	λ^{dFG}	λ^{dCG}
TYPES	15	16	16	17
SYNTAX	82	85	86	93
SEMANTICS	137	134	164	162
VALID	-	-	230	198
<i>L</i> -EQUIVALENCE	143	138	212	217
SECURITY	222	155	409	235
LATTICE	148		148	
STORE	152		112	
MEMORY			153	
HEAP			337	
BIJECTION			478	
OTHER	66		485	
TOTAL	965	894	2,830	2,635

4.1 Artifact Analysis

Table 4.1 summarizes the size of our proof scripts. In the table, we report the number of lines of Agda code needed to formalize different parts of the fine- and coarse-grained languages (λ^{dFG} and λ^{dCG}) featuring only flow-insensitive references (FI) and also with the addition of flow-sensitive references (FI + FS).¹

The bottom part of the table lists parts of the formalization that are shared and reused for both languages, where categories STORE, MEMORY and HEAP include also definitions and proofs relative to their

¹FI is the companion artifact of the conference version of this work (Vassena *et al.*, 2019), available at <https://hub.docker.com/r/marcovassena/granularity>. We developed artifact FI + FS by adding flow-sensitive references to artifact FI. The extended artifact is available at <https://hub.docker.com/r/marcovassena/granularity-ftpl>.

Table 4.2: The table reports the size of the main parts of the formal security analysis (SECURITY) of λ^{dFG} and λ^{dCG} with and without flow-sensitive references. CONFINEMENT refers to the store and heap confinement lemma, SECRET CONTEXT and PUBLIC CONTEXT to L -equivalence preservation in secret and public contexts, respectively, and PURE PRESERVATION to L -equivalence preservation for the pure semantics fragment of λ^{dCG} .

	FI		FI + FS	
	λ^{dFG}	λ^{dCG}	λ^{dFG}	λ^{dCG}
CONFINEMENT	51	24	89	44
SECRET CONTEXT	13	23	16	29
PURE PRESERVATION	-	35	-	35
PUBLIC CONTEXT	138	46	273	85
TINI	9	17	10	21
OTHER	11	10	21	21
TOTAL	222	155	409	235

L -equivalence relations and valid judgments.² The line counts for basic definitions (i.e., TYPES, SYNTAX, SEMANTICS, and L -EQUIVALENCE) are roughly the same for the two languages, in both artifacts. The only significant difference is in SECURITY, which represents the size of the scripts that state and prove the main security lemmas and theorems of the languages. In particular, we find that the security proofs in λ^{dFG} are about 43% longer than in λ^{dCG} in the artifact with only flow-insensitive references (FI) and 74% longer for the languages extended with flow-sensitive references (FI + FS). To understand better the reason behind this gap, we compare SECURITY in more detail in Table 4.2, which reports the size of the following lemmas and theorems:

- ▷ STORE AND HEAP CONFINEMENT
- ▷ L -EQUIVALENCE PRESERVATION IN SECRET CONTEXT
- ▷ PURE L -EQUIVALENCE PRESERVATION
- ▷ L -EQUIVALENCE PRESERVATION IN PUBLIC CONTEXT
- ▷ TERMINATION-INSENSITIVE NON-INTERFERENCE

²We report only one number for STORE and MEMORY because they are defined together in the same module in FI, but in separate modules in FI + FS. In the second artifact, MEMORY and HEAP are also defined by instantiating a generic container data structure for labeled data, which is counted in OTHER.

First of all, we observe that in λ^{dFG} the CONFINEMENT lemma is about twice as long as in λ^{dCG} . This is because side-effects can occur in every reduction rule in λ^{dFG} , while most constructs in λ^{dCG} are side-effect free and so we simply have to consider fewer cases in the second proof. For the languages extended with flow-sensitive reference (FI + FS), these proofs roughly *double* in size, as they additionally need to consider the validity of program configurations in order to reason about L -equivalence up to bijection. The size of SECRET CONTEXT and TINI is modest in both variants of each languages. Compared to λ^{dFG} , λ^{dCG} requires *twice* as many lines of code for these results simply because for this language each result has to be stated and proved *twice*, once for the forcing semantics and once for the thunk semantics.

PUBLIC CONTEXT is the lemma where we observe the greatest gap in the number of lines needed in λ^{dFG} and λ^{dCG} . In artifact FI, we find that PUBLIC CONTEXT for λ^{dFG} is 70% longer than PUBLIC CONTEXT and PURE PRESERVATION for λ^{dCG} combined, and 128% longer when considering also flow-sensitive references (FI + FS). Since all values are intrinsically labeled in λ^{dFG} , those proofs have to *explicitly* (i) reason about the sensitivity of program inputs and intermediate values in most cases of the proof, and (ii) rule out implicit flows for control-flow constructs. In contrast, most constructs in λ^{dCG} are security unaware and so evaluated by the pure fragment of the semantics, where the security analysis is straightforward. In particular, since pure reductions cannot perform side-effects or inspect labeled data, no implicit flows can arise in the proof of PURE PRESERVATION, which follows by simple induction (35 LOC). Only in the proof of PUBLIC CONTEXT for λ^{dCG} we need to reason about data flows explicitly using security labels, but even there implicit flows are less problematic, thanks to the monadic structure of computations. The line counts for PUBLIC CONTEXT in FI + FS show that the extra complexity of dealing with valid assumptions increases the size of the proofs in both languages, which grows by 85% for λ^{dCG} and 98% for λ^{dFG} .³ Not only the proofs of λ^{dFG} are longer than λ^{dCG} , but they are also more complicated because they

³Importantly though, the addition of flow-sensitive references did not require any change in the proof of PURE PRESERVATION.

often must (i) combine data related up to different bijections, and (ii) propagate the assumption that program configurations are valid in inductive cases. This results in an increased number of calls to helper lemmas and properties in λ^{dFG} , compared to λ^{dCG} . For example, the proof for PUBLIC CONTEXT in λ^{dFG} requires **22** calls to *weakening* (Property 5.4) and **52** calls to *valid invariant* (Property 4), while the same proof in λ^{dCG} requires only **4** calls in total.

5

Fine- to Coarse-Grained Program Translation

This section presents a provably semantics-preserving program translation from the fine-grained dynamic IFC calculus λ^{dFG} to the coarse-grained calculus λ^{dCG} . At a high level, the translation performs two tasks (i) it embeds the *intrinsic* label annotation of λ^{dFG} values into an *explicitly* labeled λ^{dCG} value via the **Labeled** type constructor and (ii) it restructures λ^{dFG} *side-effectful* expressions into *monadic operations* inside the **LIO** monad.

5.1 Types and Values

Our type-driven approach starts by formalizing this intuition in the function $\langle \cdot \rangle$, which maps the λ^{dFG} type τ to the corresponding λ^{dCG} type $\langle \tau \rangle$ (see Figure 5.1a). The function is defined by induction on types and recursively adds the **Labeled** type constructor to each existing λ^{dFG} type constructor. For the function type $\tau_1 \rightarrow \tau_2$, the result is additionally monadic, i.e., $\langle \tau_1 \rangle \rightarrow \mathbf{LIO} \langle \tau_2 \rangle$. This is because the function's body in λ^{dFG} may have side-effects. Furthermore, the translation for reference types **Ref** $s \tau$ preserves the sensitivity tag of the reference, i.e., **Ref** $s \langle \tau \rangle$.

$\langle\langle \mathbf{unit} \rangle\rangle = \mathbf{Labeled\ unit}$ $\langle\langle \mathcal{L} \rangle\rangle = \mathbf{Labeled\ } \mathcal{L}$ $\langle\langle \tau_1 \times \tau_2 \rangle\rangle = \mathbf{Labeled\ } (\langle\langle \tau_1 \rangle\rangle \times \langle\langle \tau_2 \rangle\rangle)$ $\langle\langle \tau_1 + \tau_2 \rangle\rangle = \mathbf{Labeled\ } (\langle\langle \tau_1 \rangle\rangle + \langle\langle \tau_2 \rangle\rangle)$ $\langle\langle \tau_1 \rightarrow \tau_2 \rangle\rangle = \mathbf{Labeled\ } (\langle\langle \tau_1 \rangle\rangle \rightarrow \mathbf{LIO}\ \langle\langle \tau_2 \rangle\rangle)$ $\langle\langle \mathbf{Ref\ } s\ \tau \rangle\rangle = \mathbf{Labeled\ } (\mathbf{Ref\ } s\ \langle\langle \tau \rangle\rangle)$	$\langle\langle r^\ell \rangle\rangle = \mathbf{Labeled\ } \ell\ \langle\langle r \rangle\rangle$ $\langle\langle () \rangle\rangle = ()$ $\langle\langle \ell \rangle\rangle = \ell$ $\langle\langle (v_1, v_2) \rangle\rangle = (\langle\langle v_1 \rangle\rangle, \langle\langle v_2 \rangle\rangle)$ $\langle\langle \mathbf{inl}(v) \rangle\rangle = \mathbf{inl}(\langle\langle v \rangle\rangle)$ $\langle\langle \mathbf{inr}(v) \rangle\rangle = \mathbf{inr}(\langle\langle v \rangle\rangle)$ $\langle\langle (x.e, \theta) \rangle\rangle = (x.\langle\langle e \rangle\rangle, \langle\langle \theta \rangle\rangle)$ $\langle\langle n_\ell \rangle\rangle = n_\ell$ $\langle\langle n \rangle\rangle = n$
(a) Types.	(b) Values.

Figure 5.1: Translation from λ^{dFG} to λ^{dCG} .

The translation for values (Figure 5.1b) is straightforward. Each λ^{dFG} label tag becomes the label annotation in a λ^{dCG} labeled value. The translation is homomorphic in the constructors on raw values. The translation converts a λ^{dFG} function closure into a λ^{dCG} thunk closure by translating the body of the function to a thunk, i.e., $\langle\langle e \rangle\rangle$ (see below), and translating the environment pointwise, i.e., $\langle\langle \theta \rangle\rangle = \lambda x.\langle\langle \theta(x) \rangle\rangle$. Finally, the translation preserves the memory address and the label for flow-insensitive references, i.e., $\langle\langle n_\ell \rangle\rangle = n_\ell$, and the heap address for flow-sensitive references, i.e., $\langle\langle n \rangle\rangle = n$.

5.2 Expressions

We show the translation of λ^{dFG} expressions to λ^{dCG} monadic thunks in Figure 5.2. We use the standard **do** notation for readability.¹ First, notice that the translation of all constructs occurs inside a **toLabeled**(\cdot) block. This achieves two goals, (i) it ensures that the value that results from a translated expression is *explicitly* labeled and (ii) it creates an isolated nested context where the translated thunk can execute without raising the program counter label at the top level. Inside the **toLabeled**(\cdot) block, the program counter label may rise, e.g., when

¹Syntax **do** $x \leftarrow e_1; e_2$ desugars to **bind**($e_1, x.e_2$) and syntax $e_1; e_2$ to **bind**($e_1, _ . e_2$).

$\langle\langle () \rangle\rangle = \mathbf{toLabeled}(\mathbf{return}(()))$ $\langle\langle \ell \rangle\rangle = \mathbf{toLabeled}(\mathbf{return}(\ell))$ $\langle\langle (\lambda x. e) \rangle\rangle =$ $\quad \mathbf{toLabeled}(\mathbf{return}(\lambda x. \langle\langle e \rangle\rangle))$ $\langle\langle \mathbf{inl}(e) \rangle\rangle = \mathbf{toLabeled}(\mathbf{do}$ $\quad lw \leftarrow \langle\langle e \rangle\rangle$ $\quad \mathbf{return}(\mathbf{inl}(lw)))$ $\langle\langle \mathbf{inr}(e) \rangle\rangle = \mathbf{toLabeled}(\mathbf{do}$ $\quad lw \leftarrow \langle\langle e \rangle\rangle$ $\quad \mathbf{return}(\mathbf{inr}(lw)))$ $\langle\langle (e_1, e_2) \rangle\rangle = \mathbf{toLabeled}(\mathbf{do}$ $\quad lw_1 \leftarrow \langle\langle e_1 \rangle\rangle$ $\quad lw_2 \leftarrow \langle\langle e_2 \rangle\rangle$ $\quad \mathbf{return}(lw_1, lw_2))$ $\langle\langle x \rangle\rangle = \mathbf{toLabeled}(\mathbf{unlabel}(x))$ $\langle\langle e_1 e_2 \rangle\rangle = \mathbf{toLabeled}(\mathbf{do}$ $\quad lw_1 \leftarrow \langle\langle e_1 \rangle\rangle$ $\quad lw_2 \leftarrow \langle\langle e_2 \rangle\rangle$ $\quad v_1 \leftarrow \mathbf{unlabel}(lw_1)$ $\quad lw \leftarrow v_1 lw_2$ $\quad \mathbf{unlabel}(lw))$ $\langle\langle \mathbf{case}(e, x. e_1, x. e_2) \rangle\rangle =$ $\quad \mathbf{toLabeled}(\mathbf{do}$ $\quad lw \leftarrow \langle\langle e \rangle\rangle$ $\quad v \leftarrow \mathbf{unlabel}(lw)$ $\quad lw' \leftarrow \mathbf{case}(v, x. \langle\langle e_1 \rangle\rangle, x. \langle\langle e_2 \rangle\rangle)$ $\quad \mathbf{unlabel}(lw'))$	$\langle\langle \mathbf{fst}(e) \rangle\rangle = \mathbf{toLabeled}(\mathbf{do}$ $\quad lw \leftarrow \langle\langle e \rangle\rangle$ $\quad v \leftarrow \mathbf{unlabel}(lw)$ $\quad \mathbf{unlabel}(\mathbf{fst}(v)))$ $\langle\langle \mathbf{snd}(e) \rangle\rangle = \mathbf{toLabeled}(\mathbf{do}$ $\quad lw \leftarrow \langle\langle e \rangle\rangle$ $\quad v \leftarrow \mathbf{unlabel}(lw)$ $\quad \mathbf{unlabel}(\mathbf{snd}(v)))$ $\langle\langle e_1 \sqsubseteq^? e_2 \rangle\rangle = \mathbf{toLabeled}(\mathbf{do}$ $\quad lw_1 \leftarrow \langle\langle e_1 \rangle\rangle$ $\quad lw_2 \leftarrow \langle\langle e_2 \rangle\rangle$ $\quad lu \leftarrow \mathbf{toLabeled}(\mathbf{return}(()))$ $\quad v_1 \leftarrow \mathbf{unlabel}(lw_1)$ $\quad v_2 \leftarrow \mathbf{unlabel}(lw_2)$ $\quad \mathbf{return}(\mathbf{if}$ $\quad \quad v_1 \sqsubseteq^? v_2$ $\quad \quad \mathbf{then} \mathbf{inl}(lu)$ $\quad \quad \mathbf{else} \mathbf{inr}(lu))$ $\langle\langle \mathbf{taint}(e_1, e_2) \rangle\rangle =$ $\quad \mathbf{toLabeled}(\mathbf{do}$ $\quad lw_1 \leftarrow \langle\langle e_1 \rangle\rangle$ $\quad v_1 \leftarrow \mathbf{unlabel}(lw_1)$ $\quad \mathbf{taint}(v_1)$ $\quad lw_2 \leftarrow \langle\langle e_2 \rangle\rangle$ $\quad \mathbf{unlabel}(lw_2))$ $\langle\langle \mathbf{labelOf}(e) \rangle\rangle =$ $\quad \mathbf{toLabeled}(\mathbf{do}$ $\quad lw \leftarrow \langle\langle e \rangle\rangle$ $\quad \mathbf{labelOf}(lw))$ $\langle\langle \mathbf{getLabel} \rangle\rangle =$ $\quad \mathbf{toLabeled}(\mathbf{getLabel}))$
--	---

Figure 5.2: Translation from λ^{dFG} to λ^{dCG} (expressions).

some intermediate result is unlabeled, and the translation relies on **LIO**'s floating-label mechanism to track dependencies between data of different security levels. In particular, we will show later that the value of the program counter label at the end of each nested block coincides with the label annotation of the λ^{dFG} value that the original expression evaluates to. For example, introduction forms of ground values (unit, labels, and functions) are simply returned inside the **toLabeled**(\cdot) block so that they get tagged with the current value of the program counter label just as in the corresponding λ^{dFG} introduction rules ([**LABEL**,**UNIT**,**FUN**]). Introduction forms of compounds values such as **inl**(e), **inr**(e) and (e_1, e_2) follow the same principle. The translation simply nests the translations of the nested expressions inside the same constructor, without raising the program counter label. This matches the behavior of the corresponding λ^{dFG} rules [**INL**,**INR**,**PAIR**].² For example, the λ^{dFG} reduction $(((), ())) \Downarrow_L^\emptyset ((()^L, ()^L)^L$ maps to a λ^{dCG} term that reduces to **Labeled** L (**Labeled** L ($()$), **Labeled** L ($()$)) when started with program counter label L .

The translation of variables gives some insight into how the λ^{dCG} floating-label mechanism can simulate λ^{dFG} 's tainting approach. First, the type-driven approach set out in Figure 5.1a demands that functions take only labeled values as arguments, so the variables in the source program are always associated to a labeled value in the translated program. The values that correspond to these variables are stored in the environment θ and translated separately, e.g., if $\theta(x) = r^\ell$ in λ^{dFG} , then x gets bound to $\langle\langle r^\ell \rangle\rangle = \mathbf{Labeled} \ell \langle\langle r \rangle\rangle$ when translated to λ^{dCG} . Thus, the translation converts a variable, say x , to **toLabeled**(**unlabel**(x)), so that its label gets tainted with the current program counter label. More precisely, **unlabel**(x) retrieves the labeled value associated with the variable, i.e., **Labeled** $\ell \langle\langle r \rangle\rangle$, taints the program counter with its label to make it $pc \sqcup \ell$, and returns the content, i.e., $\langle\langle r \rangle\rangle$. Since **unlabel**(x) occurs inside a **toLabeled**(\cdot) block, the code above results in **Labeled** $(pc \sqcup \ell) \langle\langle r \rangle\rangle$ when evaluated, matching precisely the tainting behavior of λ^{dFG} rule [**VAR**], i.e., $x \Downarrow_{pc}^{\theta[x \mapsto r^\ell]} r^{pc \sqcup \ell}$.

²We name a variable *lv* if it gets bound to a labeled value, i.e., to indicate that the variable has type **Labeled** τ .

The elimination forms for other types (function application, pair projections and case analysis) follow the same approach. For example, the code that translates a function application $e_1 e_2$ first executes the code that computes the translated function, i.e., $lv_1 \leftarrow \langle\langle e_1 \rangle\rangle$, then the code that computes the argument, i.e., $lv_2 \leftarrow \langle\langle e_2 \rangle\rangle$ and then retrieves the function from the first labeled value, i.e., $v_1 \leftarrow \mathbf{unlabel}(lv_1)$.³ The function v_1 applied to the labeled argument lv_2 gives a computation that gets executed and returns a labeled value lv that gets unlabeled to expose the final result (the surrounding $\mathbf{toLabeled}(\cdot)$ wraps it again right away). The translation of case analysis is analogous. The translation of pair projections first converts the λ^{dFG} pair into a computation that gives a λ^{dCG} labeled pair of labeled values, say $\mathbf{Labeled} \ell (\mathbf{Labeled} \ell_1 \langle\langle r_1 \rangle\rangle, \mathbf{Labeled} \ell_2 \langle\langle r_2 \rangle\rangle)$ and removes the label tag on the pair via $\mathbf{unlabel}$, thus raising the program counter label to $pc \sqcup \ell$. Then, it projects the appropriate component and unlabeled it, thus tainting the program counter label even further with the label of either the first or the second component. This coincides with the tainting mechanism of λ^{dFG} for projection rules, e.g., in rule [FST] where $\mathbf{fst}(e) \Downarrow_{pc}^{\theta} r_1^{pc \sqcup \ell \sqcup \ell_1}$ if $e \Downarrow_{pc}^{\theta} (r_1^{\ell_1}, r_2^{\ell_2})^{\ell}$.

Lastly, translating $\mathbf{taint}(e_1, e_2)$ requires (i) translating the expression e_1 that gives the label, (ii) using $\mathbf{taint}(\cdot)$ from λ^{dCG} to explicitly taint the program counter label with the label that e_1 gives, and (iii) translating the second argument e_2 to execute in the tainted context and unlabeled the result. The construct $\mathbf{labelOf}(e)$ of λ^{dFG} uses the corresponding λ^{dCG} primitive applied on the corresponding labeled value, say $\mathbf{Labeled} \ell \langle\langle r \rangle\rangle$, obtained from the translated expression. Notice that $\mathbf{labelOf}(\cdot)$ taints the program counter label in λ^{dCG} , which rises to $pc \sqcup \ell$, so the code just described results in $\mathbf{Labeled} (pc \sqcup \ell) \ell$, which corresponds to the translation of the result in λ^{dFG} , i.e., $\langle\langle \ell^\ell \rangle\rangle = \mathbf{Labeled} \ell \ell$ because $pc \sqcup \ell \equiv \ell$, since $pc \sqsubseteq \ell$ from Property 1. The translation of $\mathbf{getLabel}$ follows naturally

³Notice that it is incorrect to unlabeled the function before translating the argument, because that would taint the program counter label, which would raise at level, say $pc \sqcup \ell$, and affect the code that translates the argument, which was to be evaluated with program counter label equal to pc by the original *flow-insensitive* λ^{dFG} rule [APP] for function application.

$$\begin{array}{c}
(\text{WKENTYPE}) \\
\frac{\Gamma \setminus \bar{x} \vdash e : \tau}{\Gamma \vdash \mathbf{wken}(\bar{x}, e) : \tau}
\end{array}
\qquad
\begin{array}{c}
(\text{WKEN}) \\
\frac{e \Downarrow^{\theta} \setminus \bar{x} \ v}{\mathbf{wken}(\bar{x}, e) \Downarrow^{\theta} \ v}
\end{array}$$

Figure 5.3: Typing and semantics rules of **wken** for λ^{dCG} .

by simply wrapping λ^{dCG} 's **getLabel** inside a **toLabeled**(\cdot), which correctly returns the program counter label labeled with itself, i.e., **Labeled** *pc*.

5.2.1 Note on Environments and Weakening

The semantics rules of λ^{dFG} and λ^{dCG} feature an environment θ for input values that gets extended with intermediate values during program evaluation and that may be captured inside a closure. Unfortunately, this capturing behavior is undesirable for our program translation. The program translation defined above introduces temporary auxiliary variables that carry the value of intermediate results, e.g., the labeled value obtained from running a computation that translates some λ^{dFG} expression. When the translated program is executed, these values end up in the environment, e.g., by means of rules [APP] and [BIND], and mix with the input values of the source program and output values as well, thus complicating the correctness statement of the translation, which now has to account for those extra variables as well. In order to avoid this nuisance, we employ a special form of weakening that allows shrinking the environment at run-time and removing spurious values that are not needed in the rest of the program. In particular, expression **wken**(\bar{x} , e) has the same type as e if variables \bar{x} are not free in e , see the formal typing rule [WKENTYPE] in Figure 5.3. At run-time, the expression **wken**(\bar{x} , e) evaluates e in an environment from which variables \bar{x} have been dropped, so that they do not get captured in any closure created during the execution of e . Rule [WKEN] is part of the pure semantics of λ^{dCG} —the semantics of λ^{dFG} includes an analogous rule.

$\langle\langle \mathbf{new}(e) \rangle\rangle =$ $\mathbf{toLabeled}(\mathbf{do}$ $ lv \leftarrow \langle\langle e \rangle\rangle$ $ \mathbf{new}(lv))$ $\langle\langle ! e \rangle\rangle =$ $\mathbf{toLabeled}(\mathbf{do}$ $ lr \leftarrow \langle\langle e \rangle\rangle$ $ r \leftarrow \mathbf{unlabel}(lr)$ $! r)$	$\langle\langle e_1 := e_2 \rangle\rangle =$ $\mathbf{toLabeled}(\mathbf{do}$ $ lr \leftarrow \langle\langle e_1 \rangle\rangle$ $ lv \leftarrow \langle\langle e_2 \rangle\rangle$ $ r \leftarrow \mathbf{unlabel}(lr)$ $ r := lv)$ $\mathbf{toLabeled}(\mathbf{return}())$ $\langle\langle \mathbf{labelOfRef}(e) \rangle\rangle =$ $\mathbf{toLabeled}(\mathbf{do}$ $ lr \leftarrow \langle\langle e \rangle\rangle$ $ r \leftarrow \mathbf{unlabel}(lr)$ $ \mathbf{labelOfRef}(r))$
--	---

Figure 5.4: Translation λ^{dFG} to λ^{dCG} (references).

We remark that this expedient is not essential—we can avoid it by slightly complicating the correctness statement to explicitly account for those extra variables. Nor is this expedient particularly interesting. In fact, we omit **wken** from the code of the program translations to avoid clutter (our mechanization includes **wken** in the appropriate places).

5.3 References

Figure 5.4 shows the program translation of λ^{dFG} primitives that access the store and the heap via references. Notice that these translations are the same for flow-insensitive and flow-sensitive references: they replicate the behavior of λ^{dFG} primitives for each kind of reference using the corresponding primitives of λ^{dCG} . Below we explain the translation for flow-insensitive references, the discussion for flow-sensitive references is analogous.

The translation of λ^{dFG} values wraps references in λ^{dCG} labeled values (Figure 5.1b), so the translations of Figure 5.4 take care of boxing and unboxing references. The translation of $\mathbf{new}(e)$ has a top-level $\mathbf{toLabeled}(\cdot)$ block that simply translates the content ($lv \leftarrow \langle\langle e \rangle\rangle$) and puts it in a new reference ($\mathbf{new}(lv)$). The λ^{dCG} rule [NEW] (Figure

3.4) assigns the label of the translated content to the new reference, which also gets labeled with the original program counter label⁴, just as in the λ^{dFG} rule [NEW] (Figure 2.4). In λ^{dFG} , rule [READ] reads from a reference $n_{\ell'}$ at security level ℓ' that points to the ℓ -labeled memory, and returns the content $\Sigma(\ell)[n]^{\ell \sqcup \ell'}$ at level $\ell \sqcup \ell'$. Similarly, the translation creates a **toLabeled**(\cdot) block that executes to get a labeled reference $lr = \mathbf{Labeled} \ell' n_{\ell}$, extracts the reference n_{ℓ} ($r \leftarrow \mathbf{unlabel}(lr)$) tainting the program counter label with ℓ' , and then reads the reference's content further tainting the program counter label with ℓ as well. The code that translates and updates a reference consists of two **toLabeled**(\cdot) blocks. The first block is responsible for the update: it extracts the labeled reference and the labeled new content (lr and lw resp.), extracts the reference from the labeled value ($r \leftarrow \mathbf{unlabel}(lr)$) and updates it ($r := lw$). The second block, **toLabeled**(**return**(\cdot)), returns unit at security level pc , i.e., **Labeled** pc (\cdot), similar to the λ^{dFG} rule [WRITE]. The translation of **labelOfRef**(e) extracts the reference and projects its label via the λ^{dCG} primitive **labelOfRef**(\cdot), which additionally taints the program counter with the label itself, similar to the λ^{dFG} rule [LABELOFREF].

5.4 Correctness

In this section, we establish some desirable properties of the λ^{dFG} -to- λ^{dCG} translation defined above. These properties include type and semantics preservation as well as recovery of non-interference—a meta criterion that rules out a class of semantically correct (semantics preserving), yet elusive translations that do not preserve the meaning of security labels (Rajani and Garg, 2018; Barthe *et al.*, 2007).

We start by showing that the program translation preserves typing. The type translation for typing contexts Γ is pointwise, i.e., $\langle\langle \Gamma \rangle\rangle = \lambda x. \langle\langle \Gamma(x) \rangle\rangle$.

Lemma 5.1 (Type Preservation). Given a well-typed λ^{dFG} expression, i.e., $\Gamma \vdash e : \tau$, the translated λ^{dCG} expression is also well-typed, i.e., $\langle\langle \Gamma \rangle\rangle \vdash \langle\langle e \rangle\rangle : \mathbf{LIO} \langle\langle \tau \rangle\rangle$.

⁴The nested block does not execute any **unlabel**(\cdot) nor **taint**(\cdot).

Proof. By induction on the given typing derivation. \square

The main correctness criterion for the translation is semantics preservation. Intuitively, proving this theorem ensures that the program translation preserves the meaning of secure λ^{dFG} programs when translated and executed with λ^{dCG} semantics (under a translated environment). In the theorem below, the translation of stores, memories, and heaps is pointwise, i.e., $\langle\langle\Sigma\rangle\rangle = \lambda\ell.\langle\langle\Sigma(\ell)\rangle\rangle$, and $\langle\langle[]\rangle\rangle = []$ and $\langle\langle r : M \rangle\rangle = \langle\langle r \rangle\rangle : \langle\langle M \rangle\rangle$ for each ℓ -labeled memory M , and $\langle\langle[]\rangle\rangle = []$ and $\langle\langle r^\ell : \mu \rangle\rangle = \langle\langle r^\ell \rangle\rangle : \langle\langle \mu \rangle\rangle$ for heaps μ . Furthermore, notice that in the translation, the initial and final program counter labels are the same. This establishes that the program translation preserves the flow-insensitive program counter label of λ^{dFG} (by means of primitive **toLabeled**(\cdot)).

Theorem 5 (Semantics Preservation of $\langle\langle \cdot \rangle\rangle : \lambda^{dFG} \rightarrow \lambda^{dCG}$).

For all well-typed λ^{dFG} programs e , if $\langle\Sigma, \mu, e\rangle \Downarrow_{pc}^\theta \langle\Sigma', \mu', v\rangle$, then $\langle\langle\Sigma\rangle\rangle, \langle\langle\mu\rangle\rangle, pc, \langle\langle e \rangle\rangle \Downarrow^{\langle\theta\rangle} \langle\langle\Sigma'\rangle\rangle, \langle\langle\mu'\rangle\rangle, pc, \langle\langle v \rangle\rangle$.

Proof. By induction on the given evaluation derivation using basic properties of the security lattice and of the translation function.⁵ \square

5.5 Recovery of Non-Interference

We conclude this section by constructing a proof of termination-insensitive non-interference for λ^{dFG} (Theorem 2) from the corresponding theorem for λ^{dCG} (Theorem 4), using the semantics preserving translation (Theorem 5), together with a property that the translation preserves L -equivalence (Lemma 5.2), the validity of references (Lemma 5.4), as well as a property to recover source L -equivalence from target L -equivalence (Lemma 5.3). This result ensures that the meaning of labels is preserved by the translation (Rajani and Garg, 2018; Barthe *et al.*, 2007). In the absence of such an artifact, one could devise a semantics-preserving translation that simply does not use the security features of the target language. While technically correct (i.e., semantics preserving), the

⁵In our mechanized proofs, this proof also requires the (often used) axiom of functional extensionality.

translation would not be meaningful from the perspective of security.⁶ The following lemma shows that the translation of λ^{dFG} *initial* configurations, defined as $\langle\langle c \rangle\rangle^{pc} = \langle\langle \Sigma \rangle\rangle, \langle\langle \mu \rangle\rangle, pc, \langle\langle e \rangle\rangle$ if $c = \langle \Sigma, \mu, e \rangle$, preserves L -equivalence by lifting L -equivalence from source to target and back. Since the translation preserves the address of references and the size of the heap, the lemma relates the target configurations using the same bijection that relates the source configurations.

Lemma 5.2 ($\langle\langle \cdot \rangle\rangle$ preserves \approx_L^β). For all program counter labels pc and bijections β , $c_1 \approx_L^\beta c_2$ if and only if $\langle\langle c_1 \rangle\rangle^{pc} \approx_L^\beta \langle\langle c_2 \rangle\rangle^{pc}$.

Proof. By definition of L -equivalence for initial configurations in both directions (Sections 2.2 and 3.2), using injectivity of the translation function, i.e., if $\langle\langle e_1 \rangle\rangle \equiv_\alpha \langle\langle e_2 \rangle\rangle$ then $e_1 \equiv_\alpha e_2$, in the *if* direction, and by mutually proving similar lemmas for all categories. \square

The following lemma recovers L -equivalence of *source* final configurations by back-translating L -equivalence of *target* final configurations. We define the translation for λ^{dFG} *final* configurations as $\langle\langle c \rangle\rangle^{pc} = \langle\langle \Sigma \rangle\rangle, \langle\langle \mu \rangle\rangle, pc, \langle\langle v \rangle\rangle$ if $c = \langle \Sigma, \mu, v \rangle$.

Lemma 5.3 (\approx_L^β recovery via $\langle\langle \cdot \rangle\rangle$). Let $c_1 = \langle \Sigma_1, \mu_1, r_1^{\ell_1} \rangle$, $c_2 = \langle \Sigma_2, \mu_2, r_2^{\ell_2} \rangle$ be λ^{dFG} final configurations. For all bijections β and program counter label pc , such that $pc \sqsubseteq \ell_1$ and $pc \sqsubseteq \ell_2$, if $\langle\langle c_1 \rangle\rangle^{pc} \approx_L^\beta \langle\langle c_2 \rangle\rangle^{pc}$ then $c_1 \approx_L^\beta c_2$.

Proof. By case analysis on the L -equivalence relation of the target final configurations, two cases follow. First, through $\langle\langle c_1 \rangle\rangle \approx_L^\beta \langle\langle c_2 \rangle\rangle$ we recover L -equivalence of the source stores, i.e., $\Sigma_1 \approx_L^\beta \Sigma_2$, from that of the target stores, i.e., $\langle\langle \Sigma_1 \rangle\rangle \approx_L^\beta \langle\langle \Sigma_2 \rangle\rangle$, and of the source heaps, i.e., $\mu_1 \approx_L^\beta \mu_2$, from the target heaps $\langle\langle \mu_1 \rangle\rangle \approx_L^\beta \langle\langle \mu_2 \rangle\rangle$. Then, the program counter in the target configurations is either (i) *above* the attacker's level $[PC_H]$, i.e., $pc \not\sqsubseteq L$, and the source values are L -equivalent, i.e., $r_1^{\ell_1} \approx_L^\beta r_2^{\ell_2}$ by

⁶Note that such bogus translations are also ruled out due to the need to preserve the outcome of any label introspection. Nonetheless, building this proof artifact increases our confidence in the robustness of our translation. In contrast, if the enforcement of IFC is *static*, then there is no label introspection, and this proof artifact is extremely important, as argued in prior work (Rajani and Garg, 2018; Barthe *et al.*, 2007).

rule [VALUE_H] applied to $\ell_1 \not\sqsubseteq L$ and $\ell_2 \not\sqsubseteq L$ (from $pc \not\sqsubseteq L$ and, respectively, $pc \sqsubseteq \ell_1$ and $pc \sqsubseteq \ell_2$), or (ii) *below* the attacker's level [PC_L], i.e., $pc \sqsubseteq L$, then $\langle r_1^{\ell_1} \rangle \approx_L^\beta \langle r_2^{\ell_2} \rangle$ and the source values are L -equivalent, i.e., $r_1^{\ell_1} \approx_L^\beta r_2^{\ell_2}$, by Lemma 5.2 for values. \square

Recall that non-interference for λ^{dCG} extended with flow-sensitive (Theorem 4) requires side conditions about the validity of program inputs, which must contain only valid heap addresses (see the $\vdash \mathbf{Valid}(c, \theta)$ judgment in Figure 3.7). Thus, to recover non-interference for λ^{dFG} through non-interference for λ^{dCG} , we need to instantiate these judgments in the proof. Luckily, non-interference for λ^{dFG} also requires similar side conditions (Figure 2.9), so we can fulfill these assumptions by lifting the valid judgment for the source inputs into a valid judgment for the translated inputs.

Lemma 5.4 ($\langle \cdot \rangle$ preserves $\vdash \mathbf{Valid}(\cdot)$). For all program counters pc , initial configurations c and environments θ , if $\vdash \mathbf{Valid}(c, \theta)$, then $\vdash \mathbf{Valid}(\langle c \rangle^{pc}, \langle \theta \rangle)$.

Finally, we recover termination insensitive non-interference for λ^{dFG} through our semantics preserving translation to λ^{dCG} . Notice that the theorem statement below is identical to Theorem 2, what is new is the proof of the theorem, which relies on non-interference for λ^{dCG} and our verified translation.

Theorem 6 (λ^{dFG} -TINI with Bijections via $\langle \cdot \rangle$). For all program counter labels pc and valid inputs $\vdash \mathbf{Valid}(c_1, \theta_1)$ and $\vdash \mathbf{Valid}(c_2, \theta_2)$, such that $c_1 \approx_L^\beta c_2$, $\theta_1 \approx_L^\beta \theta_2$, if $c_1 \Downarrow_{pc}^{\theta_1} c'_1$, and $c_2 \Downarrow_{pc}^{\theta_2} c'_2$, then there exists an extended bijection $\beta' \supseteq \beta$, such that $c'_1 \approx_L^{\beta'} c'_2$.

Proof. We start by applying the fine to coarse grained program translation to the initial configurations and environments. By Theorem 5 (semantics preservation), we derive the corresponding λ^{dCG} reductions, i.e., $\langle c_1 \rangle^{pc} \Downarrow^{\langle \theta_1 \rangle} \langle c'_1 \rangle^{pc}$ and $\langle c_2 \rangle^{pc} \Downarrow^{\langle \theta_2 \rangle} \langle c'_2 \rangle^{pc}$. Then, we lift L -equivalence of the initial configurations and environments from *source* to *target*, i.e., from $c_1 \approx_L^\beta c_2$ to $\langle c_1 \rangle^{pc} \approx_L^\beta \langle c_2 \rangle^{pc}$ and from $\theta_1 \approx_L^\beta \theta_2$ to $\langle \theta_1 \rangle \approx_L^\beta \langle \theta_2 \rangle$ (Lemma 5.2), and similarly we lift the validity judgments (Lemma 5.4), i.e., $\vdash \mathbf{Valid}(\langle c_1 \rangle, \langle \theta_1 \rangle)$ and $\vdash \mathbf{Valid}(\langle c_2 \rangle, \langle \theta_2 \rangle)$

from $\vdash \mathbf{Valid}(c_1, \theta_1)$ and $\vdash \mathbf{Valid}(c_2, \theta_2)$, respectively. Then, we apply λ^{dCG} -*TINI with Bijections* (Theorem 4) and obtain L -equivalence of the *target* final configurations, i.e., $\langle\langle c'_1 \rangle\rangle^{pc} \approx_L^{\beta'} \langle\langle c'_2 \rangle\rangle^{pc}$ for some bijection $\beta' \supseteq \beta$. Finally, we recover L -equivalence of the final configurations from *target* to *source*, i.e., from $\langle\langle c'_1 \rangle\rangle^{pc} \approx_L^{\beta'} \langle\langle c'_2 \rangle\rangle^{pc}$ to $c'_1 \approx_L^{\beta'} c'_2$, via Lemma 5.3, applied to $c'_1 = \langle -, -, r_1^{\ell_1} \rangle$ and $c'_2 = \langle -, -, r_2^{\ell_2} \rangle$, and where $pc \sqsubseteq \ell_1$ and $pc \sqsubseteq \ell_2$ by Property 1 applied to the source reductions, i.e., $c_1 \Downarrow_{pc}^{\theta_1} c'_1$ and $c_2 \Downarrow_{pc}^{\theta_2} c'_2$. \square

6

Coarse- to Fine-Grained Program Translation

We now show a verified program translation in the opposite direction—from the coarse grained calculus λ^{dCG} to the fine grained calculus λ^{dFG} . The translation in this direction is more involved—a program in λ^{dFG} contains strictly more information than its counterpart in λ^{dCG} , namely the extra *intrinsic* label annotations that tag every value. The challenge in constructing this translation is two-fold. On one hand, the translation must come up with labels for all values. However, it is not always possible to do this statically during the translation: Often, the labels depend on input values and arise at run-time with intermediate results since the λ^{dFG} calculus is designed to compute and attach labels at run-time. On the other hand, the translation cannot conservatively under-approximate the values of labels¹— λ^{dCG} and λ^{dFG} have label introspection so, in order to get semantics preservation, labels must be preserved precisely. Intuitively, we solve this impasse by crafting a program translation that (i) preserves the labels that can be inspected by λ^{dCG} and (ii) lets the λ^{dFG} semantics compute the remaining label

¹In contrast, previous work on *static* type-based fine-to-coarse grained translation safely under-approximates the label annotations in types with the bottom label of the lattice (Rajani and Garg, 2018). The proof of type preservation of the translation recovers the actual labels via *subtyping*.

annotations automatically—we account for those labels with a *cross-language* relation that represents semantic equivalence between λ^{dFG} and λ^{dCG} modulo extra annotations (Section 6.4). The fact that the source program in λ^{dCG} cannot inspect those labels—they have no value counterpart in the source λ^{dCG} program—facilitates this aspect of the translation. We elaborate more on the technical details later.

At a high level, an interesting aspect of the translation (that informally attests that it is indeed semantics-preserving) is that it encodes the *flow-sensitive* program counter of the source λ^{dCG} program into the label annotation of the λ^{dFG} value that results from executing the translated program. For example, if a λ^{dCG} monadic expression starts with program counter label pc and results in some value, say **true**, and final program counter pc' , then the translated λ^{dFG} expression, starting with the same program counter label pc , computes the *same* value (modulo extra label annotations) at the same security level pc' , i.e., the value **true** ^{pc'} . This encoding requires keeping the value of the program counter label in the source program synchronized with the program counter label in the target program, by loosening the fine-grained precision of λ^{dFG} at run-time in a controlled way.

6.1 Types and Values

Types. The λ^{dCG} -to- λ^{dFG} translation follows the same type-driven approach used in the other direction, starting from the function $\llbracket \cdot \rrbracket$ in Figure 6.1a, that translates a λ^{dFG} type τ into the corresponding λ^{dCG} type $\llbracket \tau \rrbracket$. The translation is defined by induction on τ and preserves all the type constructors standard types. Only the cases corresponding to λ^{dCG} -specific types are interesting. In particular, it converts *explicitly* labeled types, i.e., **Labeled** τ , to a standard pair type in λ^{dFG} , i.e., $\mathcal{L} \times \llbracket \tau \rrbracket$, where the first component is the label and the second component the content of type τ . Type **LIO** τ becomes a *suspension* in λ^{dFG} , i.e., the function type **unit** $\rightarrow \llbracket \tau \rrbracket$ that delays a computation and that can be forced by simply applying it to the unit value $()$.

Values. The translation of values follows the type translation, as shown in Figure 6.1b. Notice that the translation is indexed by the program

$\llbracket \mathcal{L} \rrbracket = \mathcal{L}$ $\llbracket \mathbf{unit} \rrbracket = \mathbf{unit}$ $\llbracket \tau_1 \rightarrow \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \rightarrow \llbracket \tau_2 \rrbracket$ $\llbracket \tau_1 + \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket + \llbracket \tau_2 \rrbracket$ $\llbracket \tau_1 \times \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \times \llbracket \tau_2 \rrbracket$ $\llbracket \mathbf{Ref} \ s \ \tau \rrbracket = \mathbf{Ref} \ s \ \llbracket \tau \rrbracket$ $\llbracket \mathbf{Labeled} \ \tau \rrbracket = \mathcal{L} \times \llbracket \tau \rrbracket$ $\llbracket \mathbf{LIO} \ \tau \rrbracket = \mathbf{unit} \rightarrow \llbracket \tau \rrbracket$	$\llbracket () \rrbracket^{pc} = ()^{pc}$ $\llbracket \ell \rrbracket^{pc} = \ell^{pc}$ $\llbracket \mathbf{inl}(v) \rrbracket^{pc} = \mathbf{inl}(\llbracket v \rrbracket^{pc})^{pc}$ $\llbracket \mathbf{inr}(v) \rrbracket^{pc} = \mathbf{inr}(\llbracket v \rrbracket^{pc})^{pc}$ $\llbracket (v_1, v_2) \rrbracket^{pc} = (\llbracket v_1 \rrbracket^{pc}, \llbracket v_2 \rrbracket^{pc})^{pc}$ $\llbracket (x.e, \theta) \rrbracket^{pc} = (x.\llbracket e \rrbracket, \llbracket \theta \rrbracket^{pc})^{pc}$ $\llbracket (t, \theta) \rrbracket^{pc} = (_.\llbracket t \rrbracket, \llbracket \theta \rrbracket^{pc})^{pc}$ $\llbracket \mathbf{Labeled} \ \ell \ v \rrbracket^{pc} = (\ell^\ell, \llbracket v \rrbracket^\ell)^{pc}$ $\llbracket n_\ell \rrbracket^{pc} = (n_\ell)^{pc}$ $\llbracket n \rrbracket^{pc} = n^{pc}$
(a) Types.	(b) Values.

Figure 6.1: Translation from λ^{dCG} to λ^{dFG} (part I).

counter label (the translation is written $\llbracket v \rrbracket^{pc}$), which converts the λ^{dCG} value v in scope of a computation protected by security level pc to the corresponding fully label-annotated λ^{dFG} value. The translation is pretty straightforward and uses the program counter label to tag each value, following the λ^{dCG} principle that the program counter label protects every value in scope that is not explicitly labeled. The translation converts a λ^{dCG} function closure into a corresponding λ^{dFG} function closure by translating the body of the function to a λ^{dFG} expression (see below) and translating the environment pointwise, i.e., $\llbracket \theta \rrbracket^{pc} = \lambda x.\llbracket \theta(x) \rrbracket^{pc}$. A *think* value or a *think closure*, which denotes a suspended side-effectful computation, is also converted into a λ^{dFG} function closure. Technically, the translation would need to introduce a *fresh variable* that would get bound to \mathbf{unit} when the suspension gets forced. However, the argument to the suspension does not have any purpose, so we do not bother with giving a name to it and write $_.\llbracket t \rrbracket$ instead. (In our mechanized proofs we employ unnamed De Bruijn indexes and this issue does not arise.) The translation converts an explicitly labeled value **Labeled** $\ell \ v$, into a labeled pair at security level pc , i.e., $(\ell^\ell, \llbracket v \rrbracket^\ell)^{pc}$. The pair consists of the label ℓ tagged with itself, and the value translated at a security level equal to the label

<pre> [[()]] = () [[ℓ]] = ℓ [[x]] = x [[λx.e]] = λx. [[e]] [[e₁ e₂]] = [[e₁]] [[e₂]] [[(e₁, e₂)] = ([[e₁]], [[e₂]]) [[fst(e)]] = fst([[e]]) [[snd(e)]] = snd([[e]]) [[inl(e)]] = inl([[e]]) [[inr(e)]] = inr([[e]]) [[case (e, x.e₁, x.e₂)] = case ([[e]], x. [[e₁]], x. [[e₂]]) [[e₁ ⊑[?] e₂]] = [[e₁]] ⊑[?] [[e₂]] [[t]] = λ_. [[t]] </pre>	<pre> [[return(e)]] = [[e]] [[bind(e₁, x.e₂)] = let x = [[e₁]] () in taint(labelOf(x), [[e₂]]()) [[unlabel(e)]] = let x = [[e]] in taint(fst(x), snd(x)) [[toLabeled(e)]] = let x = [[e]] () in (labelOf(x), x) [[labelOf(e)]] = fst([[e]]) [[getLabel]] = getLabel [[taint(e)]] = taint([[e]], ()) </pre>
(a) Expressions.	(b) Thunks.

Figure 6.2: Translation from λ^{dCG} to λ^{dFG} (part II).

annotation, i.e., $[[v]]^\ell$. Notice that tagging the label with itself allows us to translate the λ^{dCG} (label introspection) primitive **labelOf**(\cdot) by simply projecting the first component, thus preserving the label and its security level across the translation.

6.2 Expressions and Thunks

The translation of pure expressions (Figure 6.2a) is trivial: it is homomorphic in all constructs, mirroring the type translation. The translation of a thunk expression t builds a suspension explicitly with a λ -abstraction, i.e., $\lambda_. [[t]]$, (the name of the variable is again irrelevant, thus we omit it as explained above), and translates the thunk itself according to the definition in Figure 6.2b. The thunk **return**(e) becomes $[[e]]$, since **return**(\cdot) does not have any side-effect. When two monadic computations are combined via **bind**($e_1, x.e_2$), the translation (i) converts the first computation to a suspension and forces it by applying unit ($[[e_1]] ()$),

(ii) binds the result to x and passes it to the second computation,² which is also converted, forced, and, *importantly*, iii) executed with a program counter label tainted with the security level of the result of the first computation ($\mathbf{taint}(\mathbf{labelOf}(x), \llbracket e_2 \rrbracket ())$). Notice that $\mathbf{taint}(\cdot)$ is essential to ensure that the second computation executes with the program counter label set to the correct value—the precision of the fine-grained system would otherwise retain the initial lower program counter label according to rule [APP] (Figure 2.2) and the value of the program counter labels in the source and target programs would differ in the remaining execution.

Similarly, the translation of $\mathbf{unlabel}(e)$ first translates the labeled expression e (the translated expression does not need to be forced because it is not of a monadic type), binds its result to x and then projects the content in a context tainted with its label, as in $\mathbf{taint}(\mathbf{fst}(x), \mathbf{snd}(x))$. This closely follows λ^{dCG} 's [UNLABEL] rule in Figure 3.3b. The translation of $\mathbf{toLabeled}(e)$ forces the nested computation with $\llbracket e \rrbracket ()$, binds its result to x and creates the pair $(\mathbf{labelOf}(x), x)$, which corresponds to the labeled value obtained in λ^{dCG} via rule [TOLABELED]. Intuitively, the translation guarantees that the value of the final program counter label in the nested computation coincides with the security level of the translated result (bound to x). Therefore, the first component contains the correct label and it is furthermore at the right security level, because $\mathbf{labelOf}(\cdot)$ protects the projected label with the label itself in λ^{dFG} . Primitive $\mathbf{labelOf}(e)$ simply projects the first component of the pair that encodes the labeled value in λ^{dFG} as explained above. Lastly, $\mathbf{getLabel}$ in λ^{dCG} maps directly to $\mathbf{getLabel}$ in λ^{dFG} —rule [GETLABEL] in λ^{dCG} simply returns the program counter label and does not raise its value, so it corresponds exactly to rule [GETLABEL] in λ^{dFG} , which returns label pc at security level pc . Similarly, $\mathbf{taint}(e)$ translates to $\mathbf{taint}(\llbracket e \rrbracket, ())$, since rule [TAINT] in λ^{dCG} taints the program counter with the label that e evaluates to, say ℓ and returns unit with program counter label equal to $pc \sqcup \ell$, which corresponds to the result of the translated program, i.e., $()^{pc \sqcup \ell}$.

²Syntax $\mathbf{let } x = e_1 \mathbf{ in } e_2$ where x is free in e_2 is a shorthand for $(\lambda x. e_2) e_1$.

$$\begin{aligned}
\llbracket \mathbf{new}(e) \rrbracket &= \\
&\quad \mathbf{let} \ x = \llbracket e \rrbracket \ \mathbf{in} \\
&\quad \quad \mathbf{new}(\mathbf{taint}(\mathbf{fst}(x), \mathbf{snd}(x))) \\
\llbracket e_1 := e_2 \rrbracket &= \\
&\quad \llbracket e_1 \rrbracket := (\mathbf{let} \ x = \llbracket e_2 \rrbracket \ \mathbf{in} \ \mathbf{taint}(\mathbf{fst}(x), \mathbf{snd}(x))) \\
\llbracket !e \rrbracket &= !\llbracket e \rrbracket \\
\llbracket \mathbf{labelOfRef}(e) \rrbracket &= \mathbf{labelOfRef}(\llbracket e \rrbracket)
\end{aligned}$$

Figure 6.3: Translation from λ^{dCG} to λ^{dFG} (references).

6.3 References

Figure 6.3 shows the translation of primitives that access the store and the heap via references. Like the translation in the opposite direction, these translations work homogeneously for flow-insensitive and flow-sensitive references, therefore we focus our discussion on flow-insensitive references. Since λ^{dCG} 's rule [NEW] in Figure 3.4 creates a new reference labeled with the label of the argument (which must be a labeled value), the translation converts $\mathbf{new}(e)$ to an expression that first binds $\llbracket e \rrbracket$ to x and then creates a new reference with the same content as the source, i.e., $\mathbf{snd}(x)$, but tainted with the label in x , i.e., $\mathbf{fst}(x)$. Notice that the use of $\mathbf{taint}(\cdot)$ is essential to ensure that λ^{dFG} 's rule [NEW] in Figure 2.4 assigns the correct label to the new reference. Due to its *fine-grained* precision, λ^{dFG} might have labeled the content with a different label that is less sensitive than the explicit label that *coarsely* approximates the security level in λ^{dCG} . Similarly, when updating a reference we translate the new labeled value into a pair, i.e., $\mathbf{let} \ x = \llbracket e_2 \rrbracket$, and taint the content with the label of the labeled value projected from the pair, i.e., $\mathbf{taint}(\mathbf{fst}(x), \mathbf{snd}(x))$.³ The translation of the primitives that read and query the label of a reference is trivial.

³Technically, tainting is redundant for flow-insensitive references, i.e., the translation $\llbracket e_1 \rrbracket := \mathbf{snd}(\llbracket e_2 \rrbracket)$ from the conference version of this work would preserve the semantics for references $\Gamma \vdash e_1 : \mathbf{Ref} \ I \ \tau$. Since flow-insensitive references have a *fixed* label, rule [WRITE] in both λ^{dFG} and λ^{dCG} accepts values less sensitive than the reference and stores them in the memory labeled like the reference. In contrast, the label of flow-sensitive references is not fixed and tainting is necessary to set the correct label for the value written in the heap.

6.4 Cross-Language Equivalence Relation

When a λ^{dCG} program is translated to λ^{dFG} via the program translation described above, the λ^{dFG} result contains strictly more information than the original λ^{dCG} result. This happens because the semantics of λ^{dFG} tracks flows of information at fine granularity, in contrast with λ^{dCG} , which instead coarsely approximates the security level of all values in scope of a computation with the program counter label. When translating a λ^{dCG} program, we can capture this condition precisely for input values θ by *homogeneously* tagging all standard (unlabeled) values with the initial program counter label, i.e., $\llbracket \theta \rrbracket^{pc}$. However, a λ^{dCG} program handles, creates and mixes unlabeled data that originated at different security levels at run-time, e.g., when a secret is unlabeled and combined with previously public (unlabeled) data. Crucially, when the translated program executes, the fine-grained semantics of λ^{dFG} tracks those flows of information precisely and thus new labels appear (these labels do not correspond to the label of any labeled value in the source value nor to the program counter label). Intuitively, this implies that the λ^{dFG} result will *not* be homogeneously labeled with the final program counter label of the λ^{dCG} computation, i.e., if a λ^{dCG} program terminates with value v and program counter label pc' , the translated λ^{dFG} program does *not* necessarily result in $\llbracket v \rrbracket^{pc'}$. The following example illustrates this issue, which we then address via a *cross-language* equivalence relation that correctly approximates the additional labels computed by λ^{dFG} .

Example 6.1. Consider the execution of λ^{dCG} program $e = \mathbf{taint}(H); \mathbf{return}(x)$, i.e., $\langle \Sigma, L, \mathbf{taint}(H); \mathbf{return}(x) \rangle \Downarrow_{[x \mapsto \mathbf{true}]} \langle \Sigma, H, \mathbf{true} \rangle$, which taints the program counter label with H , and then returns $\mathbf{true} = \mathbf{inl}()$ and the store Σ unchanged.

Let $\llbracket e \rrbracket$ be the expression obtained by applying the program translation from Figure 6.2 to the example program:

$$\begin{aligned} \llbracket e \rrbracket &= \lambda_{-}. \\ &\quad \mathbf{let } y = \mathbf{taint}(H, ()) \mathbf{ in} \\ &\quad \quad \mathbf{taint}(\mathbf{labelOf}(y), x) \end{aligned}$$

$$\begin{array}{c}
\text{(VALUE)} \\
\frac{\ell_1 \sqsubseteq pc \quad r_1 \downarrow_{pc} v_2}{r_1^{\ell_1} \downarrow_{pc} v_2}
\end{array}
\quad
\begin{array}{c}
\text{(UNIT)} \\
() \downarrow_{pc} ()
\end{array}
\quad
\begin{array}{c}
\text{(LABEL)} \\
\ell \downarrow_{pc} \ell
\end{array}
\quad
\begin{array}{c}
\text{(REF)} \\
n_\ell \downarrow_{pc} n_\ell
\end{array}$$

$$\begin{array}{c}
\text{(REF-FS)} \\
n \downarrow_{pc} n
\end{array}
\quad
\begin{array}{c}
\text{(INL)} \\
\frac{v_1 \downarrow_{pc} v'_1}{\mathbf{inl}(v_1) \downarrow_{pc} \mathbf{inl}(v'_1)}
\end{array}
\quad
\begin{array}{c}
\text{(INR)} \\
\frac{v_2 \downarrow_{pc} v'_2}{\mathbf{inr}(v_2) \downarrow_{pc} \mathbf{inr}(v'_2)}
\end{array}$$

$$\begin{array}{c}
\text{(PAIR)} \\
\frac{v_1 \downarrow_{pc} v'_1 \quad v_2 \downarrow_{pc} v'_2}{(v_1, v_2) \downarrow_{pc} (v'_1, v'_2)}
\end{array}
\quad
\begin{array}{c}
\text{(FUN)} \\
\frac{\theta_1 \downarrow_{pc} \theta_2}{(x. \llbracket e \rrbracket, \theta_1) \downarrow_{pc} (x.e, \theta_2)}
\end{array}$$

$$\begin{array}{c}
\text{(THUNK)} \\
\frac{\theta_1 \downarrow_{pc} \theta_2}{(-. \llbracket t \rrbracket, \theta_1) \downarrow_{pc} (t, \theta_2)}
\end{array}
\quad
\begin{array}{c}
\text{(LABELED)} \\
\frac{v_1 \downarrow_{\ell} v_2}{(\ell^\ell, v_1) \downarrow_{pc} (\mathbf{Labeled} \ell v_2)}
\end{array}$$

Figure 6.4: Cross-language value equivalence modulo label annotations.

When we force the program $\llbracket e \rrbracket$ and execute it starting from program counter label equal to L , and an input environment translated according to the initial program counter label, i.e., $x \mapsto \llbracket \mathbf{true} \rrbracket^L = \mathbf{inl}((\)^L)^L = \mathbf{true}^L$, we do *not* obtain the translated result homogeneously labeled with H :

$$\begin{aligned}
\langle \llbracket \Sigma \rrbracket, \llbracket e \rrbracket () \rangle \downarrow_L^{x \mapsto [\mathbf{true}^L]} & \langle \llbracket \Sigma \rrbracket, \mathbf{true}^H \rangle \\
& = \langle \llbracket \Sigma \rrbracket, \mathbf{inl}((\)^L)^H \rangle \\
& \neq \langle \llbracket \Sigma \rrbracket, \mathbf{inl}((\)^H)^H \rangle \\
& = \langle \llbracket \Sigma \rrbracket, \llbracket \mathbf{true} \rrbracket^H \rangle
\end{aligned}$$

In particular, λ^{dFG} preserves the public label tag on data nested inside the left injection, i.e., $(\)^L$ in $\mathbf{inl}((\)^L)^H$ above. This happens because λ^{dFG} 's rule [VAR] taints only the *outer* label of the value \mathbf{true}^L when it looks up variable x in program counter label H .

Solution. In order to recover a notion of semantics preservation, we introduce a key contribution of this work, a *cross-language* binary relation that associates values of the two calculi that, in the scope of a computation at a given security level, are semantically equivalent up to the extra annotations present in the λ^{dFG} value.⁴ Technically, we use this equivalence in the semantics preservation theorem in Section 6.5, which *existentially* quantifies over the result of the translated λ^{dFG} program, but guarantees that it is semantically equivalent to the result obtained in the source program.

Concretely, for a λ^{dFG} value v_1 and a λ^{dCG} value v_2 , we write $v_1 \downarrow_{\approx pc} v_2$ if the label annotations (including those nested inside compound values) of v_1 are no more sensitive than label pc and its raw value corresponds to v_2 . Figure 6.4 formalizes this intuition by means of two mutually inductive relations, one for λ^{dFG} values and one for λ^{dFG} raw values. In particular, rule [VALUE] relates λ^{dFG} value $r_1^{\ell_1}$ and λ^{dCG} value v_2 at security level pc if the label annotation on the raw value r_1 flows to the program counter label, i.e., $\ell_1 \sqsubseteq pc$, and if the raw value is in relation with the standard value, i.e., $r_1 \downarrow_{\approx pc} v_2$. The relation between raw values and standard values relates only semantically equivalent values, i.e., it is syntactic equality for ground types ([UNIT, LABEL, REF, REF-FS]), requires the same injection for values of the sum type ([INL, INR]) and requires related components for pairs ([PAIR]).

Rules [FUN] (resp. [THUNK]) relates function (resp. thunk) closures only when environments are related pointwise, i.e., $\theta_1 \downarrow_{\approx pc} \theta_2$ iff $dom(\theta_1) \equiv dom(\theta_2)$ and $\forall x. \theta_1(x) \downarrow_{\approx pc} \theta_2(x)$, and the λ^{dFG} function body $x.\llbracket e \rrbracket$ (resp. thunk body $_.\llbracket t \rrbracket$) is obtained from the λ^{dCG} function body e (resp. thunk t) via the program translation defined above. Lastly, rule [LABELED] relates a λ^{dCG} labeled value **Labeled** ℓv_1 to a pair (ℓ^ℓ, v_2) , consisting of the label ℓ protected by itself in the first component and value v_2 related with the content v_1 at security level ℓ ($v_1 \downarrow_{\approx \ell} v_2$) in the second component. This rule follows the principle of **LIO** that for explicitly labeled values, the label annotation represents an upper

⁴This relation is conceptually similar to the logical relation developed by Rajani and Garg (2018) for their translations with *static* IFC enforcement, but technically different in the treatment of labeled values.

bound on the sensitivity of the content. Stores are related pointwise, i.e., $\Sigma_1 \downarrow \approx \Sigma_2$ iff $\Sigma_1(\ell) \downarrow \approx \Sigma_2(\ell)$ for $\ell \in \mathcal{L}$, and ℓ -labeled memories relate their contents respectively at security level ℓ , i.e., $[\] \downarrow \approx [\]$ and $(r : M_1) \downarrow \approx (v : M_2)$ iff $r \downarrow_{\approx \ell} v$ and $M_1 \downarrow \approx M_2$ for λ^{dFG} and λ^{dCG} memories $M_1, M_2 : \text{Memory } \ell$. Heaps are also related pointwise, i.e., $[\] \downarrow \approx [\]$ and $(r^\ell : \mu_1) \downarrow \approx (\mathbf{Labeled } \ell v : \mu_2)$ iff $r \downarrow_{\approx \ell} v$ and $\mu_1 \downarrow \approx \mu_2$, where values at corresponding positions must additionally have the same label. Lastly, we lift the relation to initial and final configurations.

Definition 4 (Cross-Language Equivalence of Configurations). For all initial and final configurations:

- ▷ $\langle \Sigma_1, \mu_1, [\![e]\!] \ () \rangle \downarrow \approx \langle \Sigma_2, \mu_2, pc, e \rangle$ iff $\Sigma_1 \downarrow \approx \Sigma_2$ and $\mu_1 \downarrow \approx \mu_2$,
- ▷ $\langle \Sigma_1, \mu_1, [\![t]\!] \rangle \downarrow \approx \langle \Sigma_2, \mu_2, pc, t \rangle$ iff $\Sigma_1 \downarrow \approx \Sigma_2$ and $\mu_1 \downarrow \approx \mu_2$.
- ▷ $\langle \Sigma_1, \mu_1, r^{pc} \rangle \downarrow \approx \langle \Sigma_2, \mu_2, pc, v \rangle$ iff $\Sigma_1 \downarrow \approx \Sigma_2$, $\mu_1 \downarrow \approx \mu_2$, and $r \downarrow_{\approx pc} v$.

First, the relation requires the stores and heaps of initial and final configurations to be related. Additionally, for initial configurations, the relation requires the λ^{dFG} code to be obtained from the λ^{dCG} expression (resp. thunk) via the program translation function $[\![\cdot]\!]$ defined above (similar to rules [FUN] and [THUNK] in Figure 6.4). Furthermore, in the first case (expressions), the relation additionally forces the translated suspension $[\![e]\!]$ by applying it to $()$, so that when the λ^{dFG} security monitor executes the translated program, it obtains the result that corresponds to the λ^{dCG} monadic program e . Finally, in the definition for final configurations, the security level of the final λ^{dFG} result must match the program counter label pc of the final λ^{dCG} configuration, and the final λ^{dCG} result must correspond to the λ^{dFG} result up to extra annotations at security level pc , i.e., $r \downarrow_{\approx pc} v$.

Before showing semantics preservation, we prove some basic properties of the cross-language equivalence relation that will be useful later. The following property allows instantiating the semantics preservation theorem with the λ^{dCG} initial configuration. The translation for initial configurations is per-component, i.e., $[\![\langle \Sigma, \mu, pc, t \rangle]\!] = \langle [\![\Sigma]\!], [\![\mu]\!], [\![t]\!] \rangle$ and forcing for suspensions, i.e., $[\![\langle \Sigma, \mu, pc, e \rangle]\!] = \langle [\![\Sigma]\!], [\![\mu]\!], [\![e]\!] \ () \rangle$, pointwise for stores, i.e., $[\![\Sigma]\!] = \lambda \ell. [\![\Sigma(\ell)]\!]$, memories, i.e., $[\![[]]\!] = []$

and $\llbracket v : M \rrbracket = \llbracket v \rrbracket^\ell : \llbracket M \rrbracket$ for ℓ -labeled memory M , and heaps, i.e., $\llbracket [] \rrbracket = []$ and $\llbracket \mathbf{Labeled} \ell v : \mu \rrbracket = \llbracket v \rrbracket^\ell : \llbracket \mu \rrbracket$.

Property 8 (Reflexivity of $\downarrow \approx$). For all λ^{dCG} initial configurations c , $\llbracket c \rrbracket \downarrow \approx c$.

Proof. The proof is by induction and relies on analogous properties for all syntactic categories: for stores, $\llbracket \Sigma \rrbracket \downarrow \approx \Sigma$, for memories, $\llbracket M \rrbracket \downarrow \approx M$, for heaps $\llbracket \mu \rrbracket \downarrow \approx \mu$, for environments $\llbracket \theta \rrbracket^{pc} \downarrow \approx_{pc} \theta$, for values $\llbracket v \rrbracket^{pc} \downarrow \approx_{pc} v$, for any label pc . \square

The next property guarantees that values and environments remain in the relation when the program counter label rises.

Property 9 (Weakening). For all labels pc and pc' such that $pc \sqsubseteq pc'$, and for all λ^{dFG} raw values r_1 , values v_1 and environments θ_1 , and λ^{dCG} values v_2 and environments θ_2 :

- ▷ If $r_1 \downarrow \approx_{pc} v_2$ then $r_1 \downarrow \approx_{pc'} v_2$
- ▷ If $v_1 \downarrow \approx_{pc} v_2$ then $v_1 \downarrow \approx_{pc'} v_2$
- ▷ If $\theta_1 \downarrow \approx_{pc} \theta_2$ then $\theta_1 \downarrow \approx_{pc'} \theta_2$

Proof. By mutual induction on the cross-language equivalence relation. \square

6.5 Correctness

With the help of the cross-language relation defined above, we can now state and prove that the λ^{dCG} -to- λ^{dFG} translation is correct, i.e., it satisfies a semantics-preservation theorem analogous to that proved for the translation in the opposite direction. At a high level, the theorem ensures that the translation preserves the meaning of a secure terminating λ^{dCG} program when executed under λ^{dFG} semantics, with the same program counter label and translated input values. Since the translated λ^{dFG} program computes strictly more information than the original λ^{dCG} program, the theorem existentially quantify over the λ^{dFG} result, but insists that it is semantically equivalent to the original λ^{dCG} result at a security level equal to the final value of the program counter label, using the cross-language relation just defined.

We start by proving that the program translation preserves typing.

Lemma 6.1 (Type Preservation). If $\Gamma \vdash e : \tau$ then $\llbracket \Gamma \rrbracket \vdash \llbracket e \rrbracket : \llbracket \tau \rrbracket$.

Proof. By induction on the typing judgment. \square

Next, we prove semantics preservation of λ^{dCG} pure reductions. Since these reductions do not perform any security-relevant operation (they do not read or write state), they can be executed with *any* program counter label in λ^{dFG} and do not change the state in λ^{dFG} .

Lemma 6.2 ($\llbracket \cdot \rrbracket : \lambda^{dCG} \rightarrow \lambda^{dFG}$ preserves Pure Semantics). If $e \Downarrow^\theta v$ then for any program counter label pc , λ^{dFG} store Σ , heap μ , and environment θ' such that $\theta' \downarrow_{pc} \theta$, there exists a raw value r , such that $\langle \Sigma, \mu, \llbracket e \rrbracket \rangle \Downarrow_{pc}^{\theta'} \langle \Sigma, \mu, r^{pc} \rangle$ and $r \downarrow_{pc} v$.

Proof. By induction on the given evaluation derivation and using basic properties of the lattice. \square

Notice that the lemma holds for *any* input target environment θ' in relation with the source environment θ at security level pc rather than just for the translated environment $\llbracket \theta \rrbracket^{pc}$. Intuitively, we needed to generalize the lemma so that the induction principle is strong enough to discharge cases where (i) we need to prove reductions that use an existentially quantified environment, e.g., [APP], and (ii) when some intermediate result at a security level other than pc gets added to the environment, so the environment is no longer homogeneously labeled with pc . While the second condition does not arise in pure reductions, it does occur in the reduction of monadic expressions considered in the following theorem.

Theorem 7 (Thunk and Forcing Semantics Preservation via \downarrow_{\approx}). For all λ^{dFG} environments θ_1 , initial configurations c_1 and λ^{dCG} environments θ_2 and initial configuration c_2 , such that $\theta_1 \downarrow_{c_2.pc} \theta_2$ and $c_1 \downarrow_{\approx} c_2$, if $c_2 \Downarrow^{\theta_2} c'_2$, then there exists a final configuration c'_1 , such that $c_1 \Downarrow_{c_2.pc}^{\theta_1} c'_1$ and $c'_1 \downarrow_{\approx} c'_2$.

Proof. By mutual induction on the given derivation for expressions and thunks, using Lemma 6.2 for pure reductions, *Weakening* (Property 9), and basic properties for operations on related stores and heaps. \square

We finally instantiate the semantics-preservation theorem (Theorem 7) with the translated input environment at security level pc , via *reflexivity* of the cross-language relation (Property 8).

Corollary 1 (Semantics Preservation of $\llbracket \cdot \rrbracket : \lambda^{dCG} \rightarrow \lambda^{dFG}$). For all λ^{dCG} well-typed initial configurations c_2 and environments θ , if $c_2 \Downarrow^\theta c'_2$, then there exists a final λ^{dFG} configuration c'_1 such that $c'_1 \Downarrow \approx c'_2$ and $\llbracket c_2 \rrbracket \Downarrow_{pc}^{\llbracket \theta \rrbracket pc} c'_1$, where $pc = c_2.pc$.

In the corollary above, the flow-sensitive program counter of the source λ^{dCG} program gets encoded in the security level of the result of the λ^{dFG} translated program. For example, if $\langle \Sigma_2, \mu_2, pc, e \rangle \Downarrow^\theta \langle \Sigma'_2, \mu'_2, pc', v \rangle$ then, by Corollary 1 and unrolling Definition 4, there exists a raw value r at security level pc' , a store Σ'_1 , and a heap μ'_1 such that $\langle \llbracket \Sigma_2 \rrbracket, \llbracket \mu_2 \rrbracket, \llbracket e \rrbracket \rangle \Downarrow_{pc}^{\llbracket \theta \rrbracket pc} \langle \Sigma'_1, \mu'_1, r^{pc'} \rangle$, $r \Downarrow_{pc'} \approx v$, $\Sigma'_1 \Downarrow \approx \Sigma'_2$, and $\mu'_1 \Downarrow \approx \mu'_2$.

6.6 Recovery of Non-Interference

Similarly to our presentation of Theorem 6 for the translation in the opposite direction, we conclude this section with a sanity check—recovering a proof of termination-insensitive non-interference (*TINI*) for λ^{dCG} through the program translation defined above, semantics preservation (Corollary 1), and λ^{dFG} non-interference (Theorem 2). By re-proving non-interference of the source language from the target language, we show that our program translation is authentic.

To prove this result, we first need to prove that the translation preserves L -equivalence (Lemma 6.3) and the validity of references (Lemma 6.7), as well as a property for recovering source L -equivalence from target L -equivalence through the cross-language relation (Lemma 6.6). The following lemma ensures that the translation of initial configurations preserves L -equivalence.

Lemma 6.3 ($\llbracket \cdot \rrbracket$ preserves \approx_L^β). For all bijections β and initial configurations c_1 and c_2 , if $c_1 \approx_L^\beta c_2$, then $\llbracket c_1 \rrbracket \approx_L^\beta \llbracket c_2 \rrbracket$.

Proof. By induction on the L -equivalence judgment and proving similar lemmas for all syntactic categories. \square

The following lemmas recovers λ^{dCG} L -equivalence from λ^{dFG} L -equivalence using the cross-language equivalence relation for values and environments in public contexts.

Lemma 6.4 (\approx_L^β recovery from $\downarrow \approx_L$ for values and envs). For all bijections β and public program counter labels $pc \sqsubseteq L$, for all λ^{dFG} values v_1, v_2 , raw values r_1, r_2 , environments θ_1, θ_2 , and corresponding λ^{dCG} values v'_1, v'_2 , and environments θ'_1, θ'_2 :

- ▷ If $v_1 \approx_L^\beta v_2$, $v_1 \downarrow \approx_{pc} v'_1$ and $v_2 \downarrow \approx_{pc} v'_2$, then $v'_1 \approx_L^\beta v'_2$,
- ▷ If $r_1 \approx_L^\beta r_2$, $r_1 \downarrow \approx_{pc} v'_1$ and $r_2 \downarrow \approx_{pc} v'_2$, then $v'_1 \approx_L^\beta v'_2$,
- ▷ If $\theta_1 \approx_L^\beta \theta_2$, $\theta_1 \downarrow \approx_{pc} \theta'_1$ and $\theta_2 \downarrow \approx_{pc} \theta'_2$, then $\theta'_1 \approx_L^\beta \theta'_2$.

Proof. The lemmas are proved mutually, by induction on the L -equivalence relation and the cross-language equivalence relations and using injectivity of the translation function $\llbracket \cdot \rrbracket$ for closure values.⁵ \square

Next, we extend this result to program state, i.e., stores, memories, and heaps.

Lemma 6.5 (\approx_L^β recovery from $\downarrow \approx$ for state). For all bijections β , λ^{dFG} memories M_1, M_2 , stores Σ_1, Σ_2 , heaps μ_1, μ_2 , and corresponding λ^{dCG} memories M'_1, M'_2 , stores Σ'_1, Σ'_2 , heaps μ'_1 and μ'_2 :

- ▷ If $M_1 \approx_L^\beta M_2$, $M_1 \downarrow \approx M'_1$ and $M_2 \downarrow \approx M'_2$, then $M'_1 \approx_L^\beta M'_2$,
- ▷ If $\Sigma_1 \approx_L^\beta \Sigma_2$, $\Sigma_1 \downarrow \approx \Sigma'_1$ and $\Sigma_2 \downarrow \approx \Sigma'_2$, then $\Sigma'_1 \approx_L^\beta \Sigma'_2$,
- ▷ If $\mu_1 \approx_L^\beta \mu_2$, $\mu_1 \downarrow \approx \mu'_1$ and $\mu_2 \downarrow \approx \mu'_2$, then $\mu'_1 \approx_L^\beta \mu'_2$.

Proof. By induction on the L -equivalence relation and the cross-language equivalence relations and using Lemma 6.4. \square

The next lemma lifts the previous lemmas to final configurations.

⁵Technically, the function $\llbracket \cdot \rrbracket$ presented in Section 6 is not injective. For example, consider the type translation function from Figure 6.1a: $\llbracket \mathbf{Labeled\ unit} \rrbracket = \mathcal{L} \times \mathbf{unit} = \llbracket \mathcal{L} \times \mathbf{unit} \rrbracket$ but $\mathbf{Labeled\ unit} \neq \mathcal{L} \times \mathbf{unit}$, and $\llbracket \mathbf{LIO\ unit} \rrbracket = \mathbf{unit} \rightarrow \mathbf{unit} = \llbracket \mathbf{unit} \rightarrow \mathbf{unit} \rrbracket$ but $\mathbf{LIO\ unit} \neq \mathbf{unit} \rightarrow \mathbf{unit}$. We make the translation injective by (i) adding a wrapper type $\mathbf{Id}\ \tau$ to λ^{dFG} , together with constructor $\mathbf{Id}(e)$, a deconstructor $\mathbf{unId}(e)$ and raw value $\mathbf{Id}(v)$, and (ii) tagging security-relevant types and terms with the wrapper, i.e., $\llbracket \mathbf{Labeled}\ \tau \rrbracket = \mathbf{Id}\ (\mathcal{L} \times \llbracket \tau \rrbracket)$ and $\mathbf{LIO}\ \tau = \mathbf{Id}\ \mathbf{unit} \rightarrow \llbracket \tau \rrbracket$. Adapting the translations in both directions is tedious but straightforward; we refer the interested reader to our mechanized proofs for details.

Lemma 6.6 (\approx_L^β recovery from $\downarrow \approx$). Let c_1 and c_2 be λ^{dFG} final configurations, let c'_1 and c'_2 be λ^{dCG} final configurations. If $c_1 \approx_L^\beta c_2$, $c_1 \downarrow \approx c'_1$ and $c_2 \downarrow \approx c'_2$, then $c'_1 \approx_L^\beta c'_2$.

Proof. Let $c_1 = \langle \Sigma_1, \mu_1, v_1 \rangle$, $c_2 = \langle \Sigma_2, \mu_2, v_2 \rangle$, $c'_1 = \langle \Sigma'_1, \mu'_1, pc_1, v'_1 \rangle$, $c'_2 = \langle \Sigma'_2, \mu'_2, pc_2, v'_2 \rangle$. From L -equivalence of the λ^{dFG} final configurations, we have L -equivalence for the stores and the values, i.e., $\Sigma_1 \approx_L^\beta \Sigma_2$ and $v_1 \approx_L^\beta v_2$ from $c_1 \approx_L^\beta c_2$ (Section 2.2). Similarly, we derive cross-language equivalence relations for the components of the final configurations, i.e., respectively $\Sigma_1 \downarrow \approx \Sigma'_1$, $\mu_1 \downarrow \approx \mu'_1$, and $v_1 \downarrow \approx_{pc_1} v'_1$ from $c_1 \downarrow \approx c'_1$, and $\Sigma_2 \downarrow \approx \Sigma'_2$, $\mu_2 \downarrow \approx \mu'_2$, and $v_2 \downarrow \approx_{pc_2} v'_2$ from $c_2 \downarrow \approx c'_2$ (Definition 4). First, the λ^{dCG} stores and heaps are L -equivalent, i.e., $\Sigma'_1 \approx_L^\beta \Sigma'_2$ and $\mu'_1 \approx_L^\beta \mu'_2$ by Lemma 6.5 for stores and heaps, respectively. Then, two cases follow by case split on $v_1 \approx_L^\beta v_2$. Either (i) both label annotations on the values are not observable ($[VALUE_H]$), then the program counter labels are also not observable ($pc_1 \not\sqsubseteq L$ and $pc_2 \not\sqsubseteq L$ from $v_1 \downarrow \approx_{pc_1} v'_1$ and $v_2 \downarrow \approx_{pc_2} v'_2$) and $c'_1 \approx_L^\beta c'_2$ by rule $[PC_H]$ or (ii) the label annotations are equal and observable by the attacker ($[VALUE_L]$), i.e., $pc_1 \equiv pc_2 \sqsubseteq L$, then $v'_1 \approx_L^\beta v'_2$ by Lemma 6.4 for values and $c'_1 \approx_L^\beta c'_2$ by rule $[PC_L]$. \square

Before recovering non-interference, we show that the translation preserves validity of references, i.e., the assumption of the *TINI* theorem for λ^{dFG} extended with flow-sensitive references.

Lemma 6.7 ($\llbracket \cdot \rrbracket$ preserves $\vdash \mathbf{Valid}(\cdot)$). For all initial configurations c and environments θ , if $\vdash \mathbf{Valid}(c, \theta)$, then $\vdash \mathbf{Valid}(\llbracket c \rrbracket, \llbracket \theta \rrbracket^{c.pc})$.

Finally, we combine these lemmas and prove *TINI* for λ^{dCG} through our verified translation.

Theorem 8 (λ^{dCG} -*TINI* with Bijections via $\llbracket \cdot \rrbracket$). For all valid inputs $\vdash \mathbf{Valid}(c_1, \theta_1)$ and $\vdash \mathbf{Valid}(c_2, \theta_2)$ and bijections β , such that $c_1 \approx_L^\beta c_2$, $\theta_1 \approx_L^\beta \theta_2$, if $c_1 \Downarrow^{\theta_1} c'_1$, and $c_2 \Downarrow^{\theta_2} c'_2$, then there exists an extended bijection $\beta' \supseteq \beta$, such that $c'_1 \approx_L^{\beta'} c'_2$.

Proof. First, we apply the translation $\llbracket \cdot \rrbracket : \lambda^{dCG} \rightarrow \lambda^{dFG}$ to the initial configurations c_1 and c_2 and the respective environments θ_1 and θ_2 . Let

pc be the initial program counter label common to configurations c_1 and c_2 (it is the same because $c_1 \approx_L c_2$). Then, *Semantics Preservation of $\llbracket \cdot \rrbracket$* (Corollary 1) ensures that there exist two λ^{dFG} configurations c_1'' and c_2'' , such that $\llbracket c_1 \rrbracket \Downarrow_{pc}^{\llbracket \theta_1 \rrbracket^{pc}} c_1''$ and $c_1'' \Downarrow \approx c_1'$, and $\llbracket c_2 \rrbracket \Downarrow_{pc}^{\llbracket \theta_2 \rrbracket^{pc}} c_2''$ and $c_2'' \Downarrow \approx c_2'$. We then lift L -equivalence of source configurations and environments to L -equivalence in the target language via Lemma 6.3, i.e., $\llbracket \theta_1 \rrbracket^{pc} \approx_L^\beta \llbracket \theta_2 \rrbracket^{pc}$ from $\theta_1 \approx_L^\beta \theta_2$ and $\llbracket c_1 \rrbracket \approx_L^\beta \llbracket c_2 \rrbracket$ from $c_1 \approx_L^\beta c_2$. Similarly, we lift the valid judgments for λ^{dCG} inputs to λ^{dFG} via Lemma 6.7, i.e., $\vdash \mathbf{Valid}(\llbracket c_1 \rrbracket, \llbracket \theta_1 \rrbracket)$ and $\vdash \mathbf{Valid}(\llbracket c_2 \rrbracket, \llbracket \theta_2 \rrbracket)$ from $\vdash \mathbf{Valid}(c_1, \theta_1)$ and $\vdash \mathbf{Valid}(c_2, \theta_2)$, respectively. Then, we apply λ^{dFG} -*TINI with Bijections* (Theorem 2) to the reductions i.e., $\llbracket c_1 \rrbracket \Downarrow_{pc}^{\llbracket \theta_1 \rrbracket^{pc}} c_1''$ and $\llbracket c_2 \rrbracket \Downarrow_{pc}^{\llbracket \theta_2 \rrbracket^{pc}} c_2''$, which gives L -equivalence of the resulting configurations, i.e., $c_1'' \approx_L^{\beta'} c_2''$, up to some bijection $\beta' \supseteq \beta$. Then, we apply Lemma 6.6 to $c_1'' \approx_L^{\beta'} c_2''$, $c_1'' \Downarrow \approx c_1'$, and $c_2'' \Downarrow \approx c_2'$, and recover L -equivalence for the source configurations, i.e., $c_1' \approx_L^{\beta'} c_2'$, up to the same bijection $\beta' \supseteq \beta$. \square

7

Related work

7.1 Relative Expressiveness of IFC Systems

Fine- and Coarse-Grained IFC. Systematic study of the relative expressiveness of fine- and coarse-grained information flow control (IFC) systems has started only recently. Rajani *et al.* (2017) initiated this study in the context of *static* coarse- and fine-grained IFC, enforced via type systems. In more recent work, Rajani and Garg (2018) show that a fine-grained IFC type system, which they call FG, and two variants of a coarse-grained IFC type system, which they call CG, are equally expressive. Their approach is based on type-directed translations, which are type- and semantics-preserving. For proofs, they develop logical relations models of FG and the two variants of CG, as well as cross-language logical relations. Our work and some of our techniques are directly inspired by their work, but we examine *dynamic* IFC systems based on runtime monitors. As a result, our technical development is completely different. In particular, in our work we handle label introspection, which has no counterpart in the earlier work on static IFC systems, and which also requires significant care in translations. Our dynamic setting also necessitated the use of tainting operators in both the fine-grained and the coarse-grained systems. Furthermore, our

languages and translations support flow-sensitive references, which are not considered by Rajani *et al.* (2017), as the type-system of FG and CG is only flow-insensitive.

Our coarse-grained system λ^{dCG} is the dynamic analogue of the second variant of Rajani and Garg (2018)’s CG type system. This variant is described only briefly in their paper (in Section 4, paragraph “Original HLIO”) but covered extensively in Part-II of the paper’s appendix. Rajani and Garg (2018) argue that translating their fine-grained system FG to this variant of CG is very difficult and requires significant use of parametric label polymorphism. The astute reader may wonder why we do not encounter the same difficulty in translating our fine-grained system λ^{dFG} to λ^{dCG} . The reason for this is that our fine-grained system λ^{dFG} is not a direct dynamic analogue of Rajani and Garg (2018)’s FG. In λ^{dFG} , a value constructed in a context with program counter label pc automatically receives the security label pc . In contrast, in Rajani and Garg (2018)’s FG, all introduction rules create values (statically) labeled \perp . Hence, leaving aside the static-vs-dynamic difference, FG’s labels are more precise than λ^{dFG} ’s, and this makes Rajani and Garg (2018)’s FG to CG translation more difficult than our λ^{dFG} to λ^{dCG} translation. In fact, in earlier work, Rajani *et al.* (2017) introduced a different type system called FG^- , a static analogue of λ^{dFG} that labels all constructed values with pc (statically), and noted that translating it to the second variant of CG is much easier (in the static setting).

Flow-Insensitive and Flow-Sensitive IFC. Researchers have explored the relative permissiveness of static and dynamic IFC systems with respect to the flow-sensitivity of the analysis. Hunt and Sands (2006) show that flow-sensitive *static* analysis are a natural generalization of flow-insensitive analysis. In the dynamic settings, Buiras *et al.* (2014) provide a semantics-preserving translation that embeds flow-sensitive references into flow-insensitive references in LIO (Stefan *et al.*, 2012). Their embedding tracks the mutable label of references through an extra level of indirection, i.e., they translate each flow-sensitive references into a flow-insensitive reference, which points to another flow-insensitive reference that stores the content. Interestingly, this result seems to

suggest that flow-sensitive references do not fundamentally increase the expressiveness and permissiveness of dynamic coarse-grained IFC systems. In contrast to our translations, however, their embedding is not local (or macro-expressible using the terminology of Felleisen (1991)): it relies on the flow-sensitive heap to assign a fixed label to flow-insensitive references.

Some works also study the relative permissiveness of static and dynamic IFC systems. As one would expect, dynamic flow-insensitive analysis are more permissive than their static counterpart (Sabelfeld and Russo, 2009). Russo and Sabelfeld (2010) show that, perhaps surprisingly, this is not the case for *flow-sensitive* analysis, i.e., purely dynamic and static flow-sensitive analysis are incomparable in terms of permissiveness, and propose a more permissive class of hybrid IFC monitors.

Other IFC Techniques. Balliu *et al.* (2017) develop a formal framework to study the soundness and permissiveness trade-offs of dynamic information-flow trackers, but their analysis is only with respect to explicit and implicit flows. Bielova and Rezk (2016a) compare dynamic and hybrid information-flow monitors for imperative languages with respect to soundness and transparency. *Secure multi-execution* (SME) is a dynamic IFC mechanisms that enforces security precisely, at the cost of executing a program multiple times (Devriese and Piessens, 2010). *Multiple facets* (MF) simulates SME through a single execution that maintains multiple views on data, but provides weaker security guarantees than SME (Austin and Flanagan, 2012). Bielova and Rezk (2016b) show that MF is not equivalent to SME (even when SME is relaxed to enforce the same security condition as MF) and adapt MF to provide the same guarantees of SME. Schmitz *et al.* (2018) unify MF and SME in a single framework that combines their best features by allowing programs to switch between the two mechanisms at run-time.

7.2 Coarse-Grained Dynamic IFC

Coarse-grained dynamic IFC systems are prevalent in security research in operating systems (Efstathopoulos *et al.*, 2005; Krohn *et al.*, 2007;

Zeldovich *et al.*, 2006). These ideas have also been successfully applied to other domains, e.g., the web (Giffin *et al.*, 2012; Stefan *et al.*, 2014; Yip *et al.*, 2009; Bauer *et al.*, 2015), mobile applications (Jia *et al.*, 2013; Nadkarni *et al.*, 2016), IoT (Fernandes *et al.*, 2016), and distributed systems (Zeldovich *et al.*, 2008; Pedersen and Chong, 2019; Cheng *et al.*, 2012). Our λ^{dCG} calculus is based on LIO, a domain-specific language embedded in Haskell that rephrases OS-like IFC enforcement into a language-based setting (Stefan *et al.*, 2011; Stefan *et al.*, 2012). Heule *et al.* (2015) introduce a general framework for retrofitting coarse-grained IFC in any programming language in which external effects can be controlled. Co-Inflow (Xiang and Chong, 2021) extends Java with coarse-grained dynamic IFC, which is implemented via compilation, similar to our λ^{dFG} -to- λ^{dCG} translation. Laminar (Roy *et al.*, 2009) unifies mechanisms for IFC in programming languages and operating systems, resulting in a mix of dynamic fine- and coarse-grained enforcement.

7.3 Fine-Grained Dynamic IFC

The dangerous combination of highly dynamic scripting languages and third-party code in web pages (Nikiforakis *et al.*, 2012) and IoT platforms (Surbatovich *et al.*, 2017; Ahmadpanah *et al.*, 2021) has stirred a line of work on dynamic fine-grained IFC systems for JavaScript interpreters (Hedin *et al.*, 2014), engines (Bichhawat *et al.*, 2014b; Rajani *et al.*, 2015), and IoT apps (Bastys *et al.*, 2018). Our λ^{dFG} calculus is inspired by the calculus of Austin and Flanagan (2009), which we have extended with flow-insensitive references and label introspection. In a follow up work, Austin and Flanagan (2010) replace the *no-sensitive upgrade* check with a *permissive upgrade* strategy, which tracks partially leaked data to ensure it is not completely leaked. Breeze is conceptually similar to our λ^{dFG} , except for the `taint(\cdot)` primitive (Hritcu *et al.*, 2013). Breeze (Hritcu *et al.*, 2013), JSFlow (Hedin and Sabelfeld, 2012), and other dynamic fine-grained IFC languages (Bichhawat *et al.*, 2021; Austin *et al.*, 2017) feature exception handling primitives, which allow programs to recover from IFC violations without leaking data. Since LIO (Stefan *et al.*, 2017) features similar primitives, we believe that our results extend also to IFC languages with exceptions.

7.4 Label Introspection and Flow-Sensitive References

In general, dynamic fine-grained IFC systems often do not support label introspection, with Breeze (Hritcu *et al.*, 2013) as notable exception. Stefan *et al.* (2017) show that careless label introspection can leak data and discuss alternative flow-sensitive and flow-insensitive APIs (see Footnote 7). Xiang and Chong (2021) argue that the flow-sensitive API does not provide a usable programming model (because inspecting a label can unpredictably taint the program counter label) and use *opaque labeled values* instead. To support label introspection securely, our calculi protect each label with the label itself. Kozyri *et al.* (2019) generalizes this mechanism to chains of labels of arbitrary length (where each label defines the sensitivity of its predecessor) and study the trade-offs between permissiveness and storage.

Several dynamic fine-grained IFC systems support references with flow-sensitive labels (Hedin *et al.*, 2014; Austin and Flanagan, 2010; Austin and Flanagan, 2009; Bichhawat *et al.*, 2014b). This design choice, however, allows label changes to be exploited as a covert channel for information leaks (Russo and Sabelfeld, 2010; Austin and Flanagan, 2010; Austin and Flanagan, 2009). There are many approaches to preventing such leaks—from using static analysis techniques (Sabelfeld and Myers, 2006), to disallowing label upgrades depending on sensitive data (i.e., no-sensitive-upgrades (Zdancewic, 2002; Austin and Flanagan, 2009)), to avoiding branching on data whose labels have been upgraded (i.e., permissive-upgrades (Austin and Flanagan, 2010; Bichhawat *et al.*, 2021; Bichhawat *et al.*, 2014a)). Buiras *et al.* (2014) extend LIO with flow-sensitive references and explicitly protect the label of these references with the program counter label at creation time. The semantics of λ^{dCG} avoids keeping track of this extra label by using the label of the value itself as a sound approximation (see rules [NEW-FS] and [WRITE-FS] in Figure 3.6b), which corresponds precisely to the semantics of λ^{dFG} .

7.5 Proof Techniques for Termination-Insensitive Non-Interference

Since Goguen and Meseguer (1982) introduced the notion of *non-interference*, different proof techniques for IFC languages have emerged.

Our proof technique based on *L-equivalence preservation* and *confinement* dates back to the seminal work by Volpano *et al.* (1996) and is similar to the proof by Austin and Flanagan (2009), although we do not make any assumptions about the heap allocator. Instead, our treatment of heap addresses is inspired by Banerjee and Naumann (2005): we extend the *L-equivalence* relation with a bijection, which accounts precisely for different, yet indistinguishable addresses. Although bijections can complicate the formal analysis, Vassena *et al.* (2017) argue that they can be avoided by partitioning data structures per security level (as we do here for the store). Moreover, the separation between pure computations and side-effects further simplifies the security analysis of monadic languages like λ^{dCG} and, as we explain in Section 4, it leads to shorter proofs than in impure languages like λ^{dFG} . Hirsch and Cecchetti (2021) generalize this insight to other effects (non-termination and exceptions) through a new proof technique for pure languages that provide effects through a monad. In their fine-grained static IFC λ -calculus, Pottier and Simonet (2003) represent secret values and expressions explicitly, through a syntactic bracketed pair construct. This representation simplifies the non-interference proof (which is derived from subject reduction), but requires reasoning about a non-standard semantics.

Logical Relations. The first proofs of non interference that use logical relations are for the pure fragment of the fine-grained static IFC languages by Zdancewic (2002) and Heintze and Riecke (1998). Rajani *et al.* (2017) extend this proof technique for FG and CG, but their logical relation require step-indexed Kripke worlds (Birkedal *et al.*, 2011) to avoid circular arguments when reasoning about state. Recently, Gregersen *et al.* (2021) develop a mechanized semantic model based on logical relations on top of the Iris framework (Jung *et al.*, 2018) for an expressive fine-grained static IFC language. Proofs based on logical relations for stateful languages feature two types of logical relations: a *binary* relation for observable values (similar to *L-equivalence*), and a *unary* relation for secret values, which provides a semantics interpretation of the *confinement* lemma.

PER and Parametricity. Abadi *et al.* (1999) suggests a connection between parametricity and non-interference, which is formalized through a domain-theoretic semantics for the Dependency Core Calculus (DCC). Sabelfeld and Sands (2001) develop this idea further with a general semantics model of information flow based on partial equivalence relations (PER). Bowman and Ahmed (2015) prove non-interference for DCC from parametricity by proving that their translation from DCC into System F is fully abstract. Algehed and Bernardy (2019) simplifies this proof by embedding DCC into the Calculus of Construction and applies the same technique to derive a shorter proof for the core of LIO (Stefan *et al.*, 2017).

8

Conclusion

This tutorial presents a detailed and homogeneous account of dynamic fine- and coarse-grained IFC security and unifies these paradigms, which were considered fundamentally at odds with respect to precision and permissiveness. To this end, we formalized two representative IFC languages that track information flows with fine and coarse granularity, established their security guarantees using standard proof techniques, and devised verified semantics- and security-preserving translations between them. These results formally establish a connection between dynamic fine- and coarse-grained enforcement for IFC, showing that these paradigms are equally expressive under reasonable assumptions. Indeed, this work provides a systematic way to bridging the gap between a wide range of dynamic IFC techniques often proposed by the programming languages (fine-grained) and operating systems (coarse-grained) communities. As a consequence, this allows future designs of dynamic IFC to choose a coarse-grained approach, which is easier to implement and use, without giving up on the precision of fine-grained IFC.

References

- Abadi, M., A. Banerjee, N. Heintze, and J. Riecke. (1999). “A Core Calculus of Dependency”. In: *Proc. ACM Symp. on Principles of Programming Languages*. 147–160.
- Abel, A., G. Allais, A. Hameer, B. Pientka, A. Momigliano, S. Schäfer, and K. Stark. (2019). “POPLMark reloaded: Mechanizing proofs by logical relations”. *Journal of Functional Programming*. 29: e19. DOI: [10.1017/S0956796819000170](https://doi.org/10.1017/S0956796819000170).
- Ahmadpanah, M. M., D. Hedin, M. Balliu, L. E. Olsson, and A. Sabelfeld. (2021). “SandTrap: Securing JavaScript-driven Trigger-Action Platforms”. In: *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association. 2899–2916. URL: <https://www.usenix.org/conference/usenixsecurity21/presentation/ahmadpanah>.
- Alghed, M. and J.-P. Bernardy. (2019). “Simple Noninterference from Parametricity”. *Proc. ACM Program. Lang.* 3(ICFP). DOI: [10.1145/3341693](https://doi.org/10.1145/3341693).
- Austin, T. H. and C. Flanagan. (2009). “Efficient Purely-Dynamic Information Flow Analysis”. In: *Proc. of the 9th ACM Workshop on Programming Languages and Analysis for Security (PLAS '09)*.
- Austin, T. H. and C. Flanagan. (2010). “Permissive Dynamic Information Flow Analysis”. In: *Proc. of the 5th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security. PLAS '10*.

- Austin, T. H. and C. Flanagan. (2012). “Multiple Facets for Dynamic Information Flow”. In: *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL '12*. Philadelphia, PA, USA: Association for Computing Machinery. 165–178. DOI: [10.1145/2103656.2103677](https://doi.org/10.1145/2103656.2103677).
- Austin, T. H., T. Schmitz, and C. Flanagan. (2017). “Multiple Facets for Dynamic Information Flow with Exceptions”. *ACM Trans. Program. Lang. Syst.* 39(3). DOI: [10.1145/3024086](https://doi.org/10.1145/3024086).
- Balliu, M., D. Schoepe, and A. Sabelfeld. (2017). “We Are Family: Relating Information-Flow Trackers”. In: *ESORICS*.
- Banerjee, A. and D. A. Naumann. (2005). “Stack-based access control and secure information flow”. *Journal Functional Programming*. 15(2): 131–177.
- Barthe, G., T. Rezk, and A. Basu. (2007). “Security Types Preserving Compilation”. *Computer Languages, Systems & Structures*. 33(2): 35–59. DOI: [10.1016/j.cl.2005.05.002](https://doi.org/10.1016/j.cl.2005.05.002).
- Bastys, I., M. Balliu, and A. Sabelfeld. (2018). “If This Then What? Controlling Flows in IoT Apps”. In: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security. CCS '18*. Toronto, Canada: Association for Computing Machinery. 1102–1119. DOI: [10.1145/3243734.3243841](https://doi.org/10.1145/3243734.3243841).
- Bauer, L., S. Cai, L. Jia, T. Passaro, M. Stroucken, and Y. Tian. (2015). “Run-time Monitoring and Formal Analysis of Information Flows in Chromium”. In: *Proc. of the 22nd Annual Network & Distributed System Security Symposium*. Internet Society.
- Bell, E. D. and J. L. La Padula. (1976). “Secure computer system: Unified exposition and Multics interpretation”. Bedford, MA. URL: <http://csrc.nist.gov/publications/history/bell76.pdf>.
- Bichhawat, A., V. Rajani, D. Garg, and C. Hammer. (2014a). “Generalizing Permissive-Upgrade in Dynamic Information Flow Analysis”. In: *Proceedings of the Ninth Workshop on Programming Languages and Analysis for Security. PLAS'14*. Uppsala, Sweden: Association for Computing Machinery. 15–24. DOI: [10.1145/2637113.2637116](https://doi.org/10.1145/2637113.2637116).

- Bichhawat, A., V. Rajani, D. Garg, and C. Hammer. (2014b). “Information Flow Control in WebKit’s JavaScript Bytecode”. In: *International Conference on Principles of Security and Trust (POST)*. 159–178.
- Bichhawat, A., V. Rajani, D. Garg, and C. Hammer. (2021). “Permissive runtime information flow control in the presence of exceptions”. *Journal of Computer Security*. 29: 361–401. DOI: [10.3233/JCS-211385](https://doi.org/10.3233/JCS-211385).
- Bielova, N. and T. Rezk. (2016a). “A Taxonomy of Information Flow Monitors”. In: *Principles of Security and Trust*. Ed. by F. Piessens and L. Viganò. Berlin, Heidelberg: Springer Berlin Heidelberg. 46–67.
- Bielova, N. and T. Rezk. (2016b). “Spot the Difference: Secure Multi-execution and Multiple Facets”. In: *Computer Security – ESORICS 2016*. Cham: Springer International Publishing. 501–519.
- Birkedal, L., B. Reus, J. Schwinghammer, K. Støvring, J. Thamsborg, and H. Yang. (2011). “Step-Indexed Kripke Models over Recursive Worlds”. In: *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL ’11*. Austin, Texas, USA: Association for Computing Machinery. 119–132. DOI: [10.1145/1926385.1926401](https://doi.org/10.1145/1926385.1926401).
- Bove, A., P. Dybjer, and U. Norell. (2009). “A Brief Overview of Agda – A Functional Language with Dependent Types”. In: *Theorem Proving in Higher Order Logics*. Ed. by S. Berghofer, T. Nipkow, C. Urban, and M. Wenzel. Berlin, Heidelberg: Springer Berlin Heidelberg. 73–78.
- Bowman, W. J. and A. Ahmed. (2015). “Noninterference for Free”. In: *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming. ICFP 2015*. Vancouver, BC, Canada: Association for Computing Machinery. 101–113. DOI: [10.1145/2784731.2784733](https://doi.org/10.1145/2784731.2784733).
- Broberg, N., B. van Delft, and D. Sands. (2013). “Paragon for Practical Programming with Information-Flow Control”. In: *Proc. of the 11th Asian Symposium on Programming Languages and Systems. APLAS ’13*. 217–232.

- Buiras, P., D. Stefan, and A. Russo. (2014). “On Dynamic Flow-Sensitive Floating-Label Systems”. In: *Proc. of the 2014 IEEE 27th Computer Security Foundations Symposium. CSF '14*. Washington, DC, USA: IEEE Computer Society. 65–79. DOI: [10.1109/CSF.2014.13](https://doi.org/10.1109/CSF.2014.13).
- Buiras, P., D. Vytiniotis, and A. Russo. (2015). “HLIO: Mixing Static and Dynamic Typing for Information-Flow Control in Haskell”. In: *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming. ICFP 2015*. Vancouver, BC, Canada: Association for Computing Machinery. 289–301. DOI: [10.1145/2784731.2784758](https://doi.org/10.1145/2784731.2784758).
- Cheng, W., D. R. Ports, D. Schultz, V. Popic, A. Blankstein, J. Cowling, D. Curtis, L. Shriram, and B. Liskov. (2012). “Abstractions for Usable Information Flow Control in Aeolus”. In: *2012 USENIX Annual Technical Conference (USENIX ATC 12)*. Boston, MA: USENIX Association. 139–151. URL: <https://www.usenix.org/conference/atc12/technical-sessions/presentation/cheng>.
- Devriese, D. and F. Piessens. (2010). “Noninterference through Secure Multi-execution”. In: *Proc. of the 2010 IEEE Symposium on Security and Privacy. SP '10*. IEEE Computer Society.
- Efstathopoulos, P., M. Krohn, S. VanDeBogart, C. Frey, D. Ziegler, E. Kohler, D. Mazières, F. Kaashoek, and R. Morris. (2005). “Labels and Event Processes in the Asbestos Operating System”. In: *Proc. of the 20th ACM symp. on Operating systems principles. SOSP '05*.
- Felleisen, M. (1991). “On the Expressive Power of Programming Languages”. *Sci. Comput. Program.* 17(1-3): 35–75. DOI: [10.1016/0167-6423\(91\)90036-W](https://doi.org/10.1016/0167-6423(91)90036-W).
- Fernandes, E., J. Paupore, A. Rahmati, D. Simionato, M. Conti, and A. Prakash. (2016). “FlowFence: Practical Data Protection for Emerging IoT Application Frameworks”. In: *USENIX Security Symposium*. 531–548.
- Giffin, D. B., A. Levy, D. Stefan, D. Terei, D. Mazières, J. C. Mitchell, and A. Russo. (2012). “Hails: Protecting Data Privacy in Untrusted Web Applications”. In: *10th USENIX Symposium on Operating Systems Design and Implementation, OSDI '12*.

- Goguen, J. and J. Meseguer. (1982). “Security Policies and Security Models”. In: *Proc. of IEEE Symposium on Security and Privacy*. IEEE Computer Society.
- Gregersen, S. O., J. Bay, A. Timany, and L. Birkedal. (2021). “Mechanized Logical Relations for Termination-Insensitive Noninterference”. *Proc. ACM Program. Lang.* 5(POPL). DOI: [10.1145/3434291](https://doi.org/10.1145/3434291).
- Hedin, D., A. Birgisson, L. Bello, and A. Sabelfeld. (2014). “JSFlow: Tracking Information Flow in JavaScript and its APIs”. In: *Proc. of the ACM Symposium on Applied Computing (SAC '14)*.
- Hedin, D. and D. Sands. (2006). “Noninterference in the presence of non-opaque pointers”. In: *Proc. of the 19th IEEE Computer Security Foundations Workshop*. IEEE Computer Society Press.
- Hedin, D. and A. Sabelfeld. (2012). “Information-Flow Security for a Core of JavaScript”. In: *Proc. IEEE Computer Sec. Foundations Symposium*. IEEE Computer Society.
- Heintze, N. and J. G. Riecke. (1998). “The SLam calculus: programming with secrecy and integrity”. In: *Proc. ACM Symp. on Principles of Programming Languages*. 365–377.
- Heule, S., D. Stefan, E. Z. Yang, J. C. Mitchell, and A. Russo. (2015). “IFC Inside: Retrofitting Languages with Dynamic Information Flow Control”. In: *Proc. of the Conference on Principles of Security and Trust (POST '15)*. Springer.
- Hirsch, A. K. and E. Cecchetti. (2021). “Giving Semantics to Program-Counter Labels via Secure Effects”. *Proc. ACM Program. Lang.* 5(POPL). DOI: [10.1145/3434316](https://doi.org/10.1145/3434316).
- Hritcu, C., M. Greenberg, B. Karel, B. C. Pierce, and G. Morrisett. (2013). “All Your IFCEXception Are Belong to Us”. In: *Proc. of the 2013 IEEE Symposium on Security and Privacy. SP '13*. Washington, DC, USA: IEEE Computer Society. 3–17. DOI: [10.1109/SP.2013.10](https://doi.org/10.1109/SP.2013.10).
- Hunt, S. and D. Sands. (2006). “On flow-sensitive security types”. In: *Conference record of the 33rd ACM SIGPLAN-SIGACT Symp. on Principles of programming languages. POPL '06*. Charleston, South Carolina, USA: ACM. 79–90.
- Jaskelioff, M. and A. Russo. (2011). “Secure Multi-execution in Haskell”. In: *Proc. Andrei Ershov International Conference on Perspectives of System Informatics. LNCS*. Springer-Verlag.

- Jia, L., J. Aljuraidan, E. Fragkaki, L. Bauer, M. Stroucken, K. Fukushima, S. Kiyomoto, and Y. Miyake. (2013). “Run-Time Enforcement of Information-Flow Properties on Android”. In: *Proc. of the 18th European Symposium on Research in Computer Security (ESORICS '13)*. Springer.
- Jung, R., R. Krebbers, J.-H. Jourdan, A. Bizjak, L. Birkedal, and D. Dreyer. (2018). “Iris from the ground up: A modular foundation for higher-order concurrent separation logic”. *Journal of Functional Programming*. 28: e20. DOI: [10.1017/S0956796818000151](https://doi.org/10.1017/S0956796818000151).
- Kozyri, E., F. B. Schneider, A. Bedford, J. Desharnais, and N. Tawbi. (2019). “Beyond Labels: Permissiveness for Dynamic Information Flow Enforcement”. In: *2019 IEEE 32nd Computer Security Foundations Symposium (CSF)*. 351–35115. DOI: [10.1109/CSF.2019.00031](https://doi.org/10.1109/CSF.2019.00031).
- Krohn, M., A. Yip, M. Brodsky, N. Cliffer, M. F. Kaashoek, E. Kohler, and R. Morris. (2007). *Information Flow Control for Standard OS Abstractions*. Stevenson, Washington, USA. DOI: [10.1145/1294261.1294293](https://doi.org/10.1145/1294261.1294293).
- Myers, A. C., L. Zheng, S. Zdancewic, S. Chong, and N. Nystrom. (2006). *Jif 3.0: Java information flow*. URL: <http://www.cs.cornell.edu/jif>.
- Nadkarni, A., B. Andow, W. Enck, and S. Jha. (2016). “Practical DIFC Enforcement on Android.” In: *USENIX Security Symposium*. 1119–1136.
- Nikiforakis, N., L. Invernizzi, A. Kapravelos, S. Van Acker, W. Joosen, C. Kruegel, F. Piessens, and G. Vigna. (2012). “You Are What You Include: Large-Scale Evaluation of Remote Javascript Inclusions”. In: *Proceedings of the 2012 ACM Conference on Computer and Communications Security. CCS '12*. Raleigh, North Carolina, USA: Association for Computing Machinery. 736–747. DOI: [10.1145/2382196.2382274](https://doi.org/10.1145/2382196.2382274).
- Norell, U. (2009). “Dependently Typed Programming in Agda”. In: *Advanced Functional Programming: 6th International School, AFP 2008, Heijen, The Netherlands, May 2008, Revised Lectures*. Ed. by P. Koopman, R. Plasmeijer, and D. Swierstra. Berlin, Heidelberg: Springer Berlin Heidelberg. 230–266. DOI: [10.1007/978-3-642-04652-0_5](https://doi.org/10.1007/978-3-642-04652-0_5).

- Pedersen, M. V. and S. Chong. (2019). “Programming with Flow-Limited Authorization: Coarser is Better”. In: *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. 63–78. DOI: [10.1109/EuroSP.2019.00015](https://doi.org/10.1109/EuroSP.2019.00015).
- Pottier, F. and V. Simonet. (2003). “Information Flow Inference for ML”. *ACM Trans. Program. Lang. Syst.* 25(1): 117–158. DOI: [10.1145/596980.596983](https://doi.org/10.1145/596980.596983).
- Rajani, V., I. Bastys, W. Rafnsson, and D. Garg. (2017). “Type Systems for Information Flow Control: The Question of Granularity”. *ACM SIGLOG News*. 4(1): 6–21. DOI: [10.1145/3051528.3051531](https://doi.org/10.1145/3051528.3051531).
- Rajani, V., A. Bichhawat, D. Garg, and C. Hammer. (2015). “Information Flow Control for Event Handling and the DOM in Web Browsers”. In: *2015 IEEE 28th Computer Security Foundations Symposium*. 366–379. DOI: [10.1109/CSF.2015.32](https://doi.org/10.1109/CSF.2015.32).
- Rajani, V. and D. Garg. (2018). “Types for Information Flow Control: Labeling Granularity and Semantic Models”. In: *Proc. of the IEEE Computer Security Foundations Symp. CSF ’18*. IEEE Computer Society.
- Roy, I., D. E. Porter, M. D. Bond, K. S. McKinley, and E. Witchel. (2009). “Laminar: Practical Fine-grained Decentralized Information Flow Control”. In: *Proc. of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation. PLDI ’09*. Dublin, Ireland: ACM. 63–74. DOI: [10.1145/1542476.1542484](https://doi.org/10.1145/1542476.1542484).
- Russo, A. (2015). “Functional Pearl: Two Can Keep a Secret, if One of Them Uses Haskell”. In: *Proc. of the 20th ACM SIGPLAN International Conference on Functional Programming. ICFP 2015*. ACM.
- Russo, A., K. Claessen, and J. Hughes. (2009). “A library for light-weight Information-Flow Security in Haskell”. *ACM SIGPLAN Notices (HASKELL ’08)*. 44(Jan.): 13. DOI: [10.1145/1543134.1411289](https://doi.org/10.1145/1543134.1411289).
- Russo, A. and A. Sabelfeld. (2010). “Dynamic vs. Static Flow-Sensitive Security Analysis”. In: *Proc. of the 2010 23rd IEEE Computer Security Foundations Symp. CSF ’10*. IEEE Computer Society. 186–199.

- Sabelfeld, A. and A. Russo. (2009). “From dynamic to static and back: Riding the roller coaster of information-flow control research”. In: *Proc. Andrei Ershov International Conference on Perspectives of System Informatics (PSI '09)*. LNCS. Springer-Verlag.
- Sabelfeld, A. and A. C. Myers. (2006). “Language-based Information-flow Security”. *IEEE J.Sel. A. Commun.* 21(1): 5–19. DOI: [10.1109/JSAC.2002.806121](https://doi.org/10.1109/JSAC.2002.806121).
- Sabelfeld, A. and D. Sands. (2001). “A Per Model of Secure Information Flow in Sequential Programs”. *Higher Order Symbol. Comput.* 14(1): 59–91. DOI: [10.1023/A:1011553200337](https://doi.org/10.1023/A:1011553200337).
- Schmitz, T., M. Alghed, C. Flanagan, and A. Russo. (2018). “Faceted Secure Multi Execution”. In: *Proc. of the 2018 ACM SIGSAC Conference on Computer and Communications Security. CCS '18*. Toronto, Canada: ACM. 1617–1634. DOI: [10.1145/3243734.3243806](https://doi.org/10.1145/3243734.3243806).
- Stefan, D., A. Russo, P. Buiras, A. Levy, J. C. Mitchell, and D. Mazières. (2012). “Addressing Covert Termination and Timing Channels in Concurrent Information Flow Systems”. In: *International Conference on Functional Programming (ICFP)*. ACM SIGPLAN.
- Stefan, D., A. Russo, D. Mazières, and J. C. Mitchell. (2017). “Flexible Dynamic Information Flow Control in the Presence of Exceptions”. *Journal of Functional Programming.* 27.
- Stefan, D., A. Russo, J. C. Mitchell, and D. Mazières. (2011). “Flexible Dynamic Information Flow Control in Haskell”. In: *Proc. of the 4th ACM Symposium on Haskell. Haskell '11*. Tokyo, Japan: ACM. 95–106. DOI: [10.1145/2034675.2034688](https://doi.org/10.1145/2034675.2034688).
- Stefan, D., E. Z. Yang, P. Marchenko, A. Russo, D. Herman, B. Karp, and D. Mazières. (2014). “Protecting Users by Confining JavaScript with COWL”. In: *Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation. OSDI'14*. Broomfield, CO: USENIX Association. 131–146. URL: <http://dl.acm.org/citation.cfm?id=2685048.2685060>.

- Surbatovich, M., J. Aljuraidan, L. Bauer, A. Das, and L. Jia. (2017). “Some Recipes Can Do More Than Spoil Your Appetite: Analyzing the Security and Privacy Risks of IFTTT Recipes”. In: *Proceedings of the 26th International Conference on World Wide Web. WWW '17*. Perth, Australia: International World Wide Web Conferences Steering Committee. 1501–1510. DOI: [10.1145/3038912.3052709](https://doi.org/10.1145/3038912.3052709).
- Tsai, T.-C., A. Russo, and J. Hughes. (2007). “A Library for Secure Multi-threaded Information Flow in Haskell”. In: *Proc. of the 20th IEEE Computer Security Foundations Symposium (CSF'07)*. 187–202. DOI: [10.1109/CSF.2007.6](https://doi.org/10.1109/CSF.2007.6).
- Vassena, M. and A. Russo. (2016). “On Formalizing Information-Flow Control Libraries”. In: *Proc. of the 2016 ACM Workshop on Programming Languages and Analysis for Security. PLAS '16*. Vienna, Austria: ACM. 15–28. DOI: [10.1145/2993600.2993608](https://doi.org/10.1145/2993600.2993608).
- Vassena, M., A. Russo, P. Buiras, and L. Waye. (2017). “MAC A Verified Static Information-Flow Control Library”. *Journal of Logical and Algebraic Methods in Programming*. DOI: <https://doi.org/10.1016/j.jlamp.2017.12.003>.
- Vassena, M., A. Russo, D. Garg, V. Rajani, and D. Stefan. (2019). “From Fine- to Coarse-Grained Dynamic Information Flow Control and Back”. *Proc. ACM Program. Lang.* 3(POPL). DOI: [10.1145/3290389](https://doi.org/10.1145/3290389).
- Volpano, D., G. Smith, and C. Irvine. (1996). “A Sound Type System for Secure Flow Analysis”. *J. Computer Security.* 4(3): 167–187.
- Volpano, D. and G. Smith. (1997). “Eliminating Covert Flows with Minimum Typings”. In: *Proc. of the 10th IEEE workshop on Computer Security Foundations. CSFW '97*. IEEE Computer Society.
- Xiang, J. and S. Chong. (2021). “Co-Inflow: Coarse-grained Information Flow Control for Java-like Languages”. In: *Proceedings of the 2021 IEEE Symposium on Security and Privacy*. Piscataway, NJ, USA: IEEE Press.
- Yang, J., K. Yessenov, and A. Solar-Lezama. (2012). “A Language for Automatically Enforcing Privacy Policies”. In: *Proc. of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL '12*. Philadelphia, PA, USA: ACM. 85–96. DOI: [10.1145/2103656.2103669](https://doi.org/10.1145/2103656.2103669).

- Yip, A., N. Narula, M. Krohn, and R. Morris. (2009). “Privacy-preserving Browser-side Scripting with BFlow”. In: *Proc. of the 4th ACM European Conference on Computer Systems. EuroSys '09*. ACM.
- Zdancewic, S. A. (2002). “Programming Languages for Information Security”. *PhD thesis*. Ithaca, NY, USA.
- Zeldovich, N., S. Boyd-Wickizer, E. Kohler, and D. Mazières. (2006). “Making Information Flow Explicit in HiStar”. In: *Proceedings of the 7th USENIX Symposium on Operating Systems Design and Implementation - Volume 7. OSDI '06*. Seattle, WA: USENIX Association. 19–19. URL: <http://dl.acm.org/citation.cfm?id=1267308.1267327>.
- Zeldovich, N., S. Boyd-Wickizer, and D. Mazières. (2008). “Securing Distributed Systems with Information Flow Control”. In: *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation. NSDI'08*. San Francisco, California: USENIX Association. 293–308. URL: <http://dl.acm.org/citation.cfm?id=1387589.1387610>.