

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

MARCO VASSENA, Chalmers University of Technology, Sweden

ALEJANDRO RUSSO, Chalmers University of Technology, Sweden

DEEPAK GARG, Max Planck Institute for Software Systems, Germany

VINEET RAJANI, Max Planck Institute for Software Systems, Germany

DEIAN STEFAN, University of California San Diego, USA

We show that fine-grained and coarse-grained dynamic information-flow control (IFC) systems are equally expressive. 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. This result addresses a longstanding open problem in IFC: whether coarse-grained dynamic IFC techniques are less expressive than fine-grained dynamic IFC techniques (they are not!). The translations also 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.

Additional Key Words and Phrases: Information-flow control, verified source-to-source transformations, Agda

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, where the *granularity* of the tracking system alone has important implications for the usage of IFC technology in practice. 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., [Efsthopoulos et al. 2005; Krohn et al. 2007b; Zeldovich et al. 2006]) 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; Hedin et al. 2014; Hritcu et al. 2013a; Yang et al. 2012; Zdancewic 2002]) track the flow of information in a more *fine-grained* fashion, e.g., at the granularity of program variables and references.

Authors' addresses: Marco Vassena, Chalmers University of Technology, Sweden; Alejandro Russo, Chalmers University of Technology, Sweden; Deepak Garg, Max Planck Institute for Software Systems, Germany; Vineet Rajani, Max Planck Institute for Software Systems, Germany; Deian Stefan, University of California San Diego, USA.

2019. 2475-1421/2019/1-ART1 \$15.00
<https://doi.org/>

50 Dynamic coarse-grained IFC systems in the style of LIO [Buiras et al. 2015; Heule et al.
51 2015; Stefan et al. 2012, 2017, 2011; Vassena et al. 2017] have several advantages over
52 dynamic fine-grained IFC systems. Such coarse-grained systems are often easier to design
53 and implement—they inherently track less information. For example, LIO protects against
54 control-flow-based *implicit flows* by tracking information at a coarse-grained level—to
55 branch on secrets, LIO programs must first taint the context where secrets are going to
56 be observed. Finally, coarse-grained systems often require considerably fewer programmer
57 annotations—unlike fine-grained ones. More specifically, developers often only need to
58 annotate a single-label to protect everything in the scope of a thread or process responsible
59 to handle sensitive data.

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

73 This paper removes the division between fine- and coarse-grained dynamic IFC systems
74 and the belief that they are fundamentally different. In particular, we show that *dynamic*
75 fine-grained and coarse-grained IFC are equally expressive. Our work is inspired by the
76 recent work of Rajani et al. [2017]; Rajani and Garg [2018], who prove similar results for
77 *static* fine-grained and coarse-grained IFC systems. Specifically, they establish a semantics-
78 and type-preserving translation from a coarse-grained IFC type system to a fine-grained
79 one and vice-versa. We complete the picture by showing a similar result for dynamic IFC
80 systems that additionally allow *introspection on labels* at run-time. While label introspection
81 is meaningless in a static IFC system, in a dynamic IFC system this feature is key to both
82 writing practical applications and mitigating the label creep problem [Stefan et al. 2017].

83 Using Agda, we formalize a traditional fine-grained system (in the style of [Austin and
84 Flanagan 2009]) extended with label introspection primitives, as well as a coarse-grained
85 system (in the style of [Stefan et al. 2017]). We then define and formalize modular semantics-
86 preserving translations between them. Our translations are macro-expressible in the sense
87 of Felleisen [1991].

88 We show that a translation from fine- to coarse-grained is possible when the coarse-grained
89 system is equipped with a primitive that limits the scope of tainting (e.g., when reading
90 sensitive data). In practice, this is not an imposing requirement since most coarse-grained
91 systems rely on such primitives for compartmentalization. For example, Stefan et al. [2012,
92 2017], provide `toLabeled` blocks and threads for precisely this purpose. Dually, we show
93 that the translation from coarse- to fine-grained is possible when the fine-grained system has
94 a primitive `taint(·)` that relaxes precision to keep the *program counter label* synchronized
95 when translating a program to the coarse-grained language. While this primitive is largely
96 necessary for us to establish the coarse- to fine-grained translation, extending existing
97 fine-grained systems with it is both secure and trivial.

Types:	$\tau ::=$	unit $\tau_1 \rightarrow \tau_2$ $\tau_1 + \tau_2$ $\tau_1 \times \tau_2$ \mathcal{L} Ref τ
Labels:	$\ell, pc \in$	\mathcal{L}
Addresses:	$n \in$	\mathbb{N}
Environment:	$\theta \in$	$Var \rightarrow Value$
Raw Value:	$r ::=$	$()$ $(x.e, \theta)$ inl (v) inr (v) (v_1, v_2) ℓ n_ℓ
Value	$v ::=$	r^ℓ
Expressions:	$e ::=$	x $\lambda x.e$ $e_1 e_2$ $()$ ℓ inl (e) inr (e) case ($e, x.e_1, x.e_2$) (e_1, e_2) fst (e) snd (e) getLabel labelOf (e) taint (e_1, e_2) new (e) $!e$ $e_1 := e_2$ labelOfRef (e) $e_1 \sqsubseteq^? e_2$
Configuration:	$c ::=$	$\langle \Sigma, e \rangle$
Store:	$\Sigma \in$	$(\ell : Label) \rightarrow Memory \ell$
Memory ℓ :	$M ::=$	$[]$ $r : M$

Fig. 1. Syntax of λ^{dFG} .

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 paper are:

- A pair of semantics-preserving translations between traditional dynamic fine-grained and coarse-grained IFC systems equipped with label introspection (Theorems 3 and 5).
- Two different proofs of *termination-insensitive* non-interference (TINI) for each calculus: one is derived directly in the usual way (Theorems 1 and 2), while the other is recovered via our verified translation (Theorems 4 and 6).
- Mechanized Agda proofs of our results ($\sim 4,000$ LOC)¹.

The rest of this paper is organized as follows. Our dynamic fine- and coarse-grained IFC calculi are introduced in Sections 2 and 3, respectively. We also prove their soundness guarantees (i.e., termination-insensitive non-interference). Section 4 presents the translation from the fine- to the coarse-grained calculus and recovers the non-interference of the former from the non-interference theorem of the latter. Section 5 has similar results in the other direction. Related work is described in Section 6 and Section 7 concludes the paper.

2 FINE-GRAINED 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 1 shows the syntax of the dynamic fine-grained IFC calculus λ^{dFG} , which is inspired by Austin and Flanagan [Austin and Flanagan 2009] and extended with a standard (security unaware) type system $\Gamma \vdash e : \tau$ (omitted), sum and

¹Artifact available at <https://hub.docker.com/r/marcovassena/granularity/>

148 product data types and security labels $\ell \in \mathcal{L}$ that form a lattice $(\mathcal{L}, \sqsubseteq)$.² In order to
 149 capture flows of information precisely at run-time, λ^{dFG} features *intrinsically labeled* values,
 150 written, r^ℓ , meaning that raw value r has security level ℓ . Compound values, e.g., pairs and
 151 sums, carry labels to tag the security level of each component, for example a pair containing
 152 a secret and a public boolean would be written $(\mathbf{true}^H, \mathbf{false}^L)$.³ Functional values are
 153 closures $(x.e, \theta)$, where x is the variable that binds the argument in the body of the function
 154 e and all other free variables are mapped to some labeled value in the environment θ . λ^{dFG}
 155 features a labeled partitioned store, i.e., $\Sigma \in (\ell : \mathcal{L}) \rightarrow \text{Memory } \ell$, where *Memory* ℓ is the
 156 memory that contains values at security level ℓ . Each reference carries an additional label
 157 annotation that records the label of the memory it refers to—reference n_ℓ points to the n -th
 158 cell of the ℓ -labeled memory, i.e., $\Sigma(\ell)$. Notice that this label has nothing to do with the
 159 *intrinsic* label that decorates the reference itself. For example, a reference $(n_H)^L$ represents
 160 a secret reference in a public context, whereas $(n_L)^H$ represents a public reference in a secret
 161 context. Notice that there is no order invariant between those labels—in the latter case, the
 162 runtime monitor enforcing IFC prevents writing data to the reference to avoid *implicit flows*.
 163 A program can create, read and write a labeled reference via constructs $\mathbf{new}(e)$, $!e$ and
 164 $e_1 := e_2$ and inspect its subscripted label with the primitive $\mathbf{labelOfRef}(\cdot)$.
 165

166 2.1 Dynamics

167 The operational semantics of λ^{dFG} includes a security monitor that propagates the label
 168 annotations of input values during program execution and assigns security labels to the
 169 result accordingly. The monitor prevents information leakage by stopping the execution of
 170 potentially leaky programs, which is reflected in the semantics by not providing reduction
 171 rules for the cases that may cause insecure information flow.⁴ The relation $\langle \Sigma, e \rangle \Downarrow_{pc}^\theta \langle \Sigma', v \rangle$
 172 denotes the evaluation of program e with initial store Σ that terminates with labeled value v
 173 and final store Σ' . The environment θ stores the input values of the program and is extended
 174 with intermediate results during function application and case analysis. The subscript pc
 175 is the *program counter* label [Sabelfeld and Myers 2003]—it is a label that represents the
 176 security level of the context in which the expression is evaluated. The semantics employs the
 177 program counter label to 1) propagate and assign labels to values computed by a program
 178 and 2) prevent implicit flow leaks that exploit the control flow and the store (explained
 179 below).

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

186 When the control flow of a program *depends* on some intermediate value, the program
 187 counter label is joined with the value's label so that the label of the final result will be
 188 tainted with the result of the intermediate value. For instance, consider case analysis, i.e.,
 189 $\mathbf{case } e \ x.e_1 \ x.e_2$. Rules $[\mathbf{CASE}_1]$ and $[\mathbf{CASE}_2]$ evaluate the scrutinee e to a value (either
 190

191 ² The lattice is arbitrary and fixed. In examples we will often use the two point lattice $\{L, H\}$, which only
 192 disallows secret to public flow of information, i.e., $H \not\sqsubseteq L$.

193 ³ We define the boolean type $\mathbf{bool} = \mathbf{unit} + \mathbf{unit}$, boolean values as raw values, i.e., $\mathbf{true} = \mathbf{inl}(()^L)$,
 194 $\mathbf{false} = \mathbf{inr}(()^L)$ and $\mathbf{if } e \ \mathbf{then } e_1 \ \mathbf{else } e_2 = \mathbf{case } e \ _ .e_1 \ _ .e_2$.

195 ⁴In this work we ignore leaks that exploit program termination. This is accounted for in the *termination*
 196 *insensitive* security condition satisfied by λ^{dFG} (Theorem 1).

$$\begin{array}{c}
\text{(VAR)} \qquad \qquad \qquad \text{(UNIT)} \qquad \qquad \qquad \text{(LABEL)} \\
\langle \Sigma, x \rangle \Downarrow_{pc}^{\theta} \langle \Sigma, \theta(x) \sqcup pc \rangle \qquad \langle \Sigma, () \rangle \Downarrow_{pc}^{\theta} \langle \Sigma, ()^{pc} \rangle \qquad \langle \Sigma, \ell \rangle \Downarrow_{pc}^{\theta} \langle \Sigma, \ell^{pc} \rangle \\
\text{(FUN)} \\
\langle \Sigma, \lambda x. e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma, (x. e, \theta)^{pc} \rangle \\
\text{(APP)} \\
\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 \qquad \qquad \text{(INR)} \\
\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 \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{)} \\
\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{)} \\
\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 \qquad \qquad \text{(FST)} \\
\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} \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} \\
\text{(SND)} \qquad \qquad \qquad \text{(TAINT)} \\
\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} \qquad \frac{\langle \Sigma, e_1 \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', \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}$$

Fig. 2. Big-step semantics for λ^{dFG} (part I).

$\mathbf{inl}(v)^{\ell}$ or $\mathbf{inr}(v)^{\ell}$), add the value to the environment, i.e., $\theta[x \mapsto v]$, and then execute 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 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

$$\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} \\
\\
\text{(GETLABEL)} \\
\langle \Sigma, \mathbf{getLabel} \rangle \Downarrow_{pc}^{\theta} \langle \Sigma'', pc^{pc} \rangle \\
\\
\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}(\cdot)^{pc} \rangle^{\ell_1 \sqcup \ell_2}} \\
\\
\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}(\cdot)^{pc} \rangle^{\ell_1 \sqcup \ell_2}}
\end{array}$$

Fig. 3. Big-step semantics for λ^{dFG} (part II).

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., $(\lambda x.()) y \Downarrow_{pc}^{y \rightarrow 42^H} ()^L$. 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 $x \Downarrow_{pc}^{x \rightarrow ()^L} ()^H$.

Label Introspection. The λ^{dFG} calculus features primitives for label introspection, namely **getLabel**, **labelOf**(\cdot) and $\sqsubseteq^?$ —see Figure 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. 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 4 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

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

Fig. 4. Big-step semantics for λ^{dFG} (references).

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 is 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 $pc \sqsubseteq \ell$.

Proof. By induction on the given evaluation derivation.

References. We now extend the semantics presented earlier with primitives that inspect, access and modify the labeled store via labeled references. See Figure 4. Rule [NEW] creates a reference n_{ℓ} , labeled with the security level of the initial content, i.e., ℓ in the ℓ -labeled memory $\Sigma(\ell)$ and updates the memory store accordingly.⁵ 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 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*)⁶, 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 *explicitly*). 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

⁵ $|M|$ denotes the length of memory M —memory indices start at 0.

⁶ Notice 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** $x = \mathbf{new}(0)$ **in if** *secret* **then** $x := 1$ **else** $()$.

$$\begin{array}{c}
\text{(VALUE}_L\text{)} \\
\frac{\ell \sqsubseteq L \quad r_1 \approx_L r_2}{r_1^\ell \approx_L r_2^\ell} \\
\text{(VALUE}_H\text{)} \\
\frac{\ell_1 \not\sqsubseteq L \quad \ell_2 \not\sqsubseteq L}{r_1^{\ell_1} \approx_L r_2^{\ell_2}} \\
\text{(UNIT)} \quad \text{(LABEL)} \\
() \approx_L () \quad \ell \approx_L \ell \\
\text{(CLOSURE)} \\
\frac{e_1 \equiv_\alpha e_2 \quad \theta_1 \approx_L \theta_2}{(e_1, \theta_1) \approx_L (e_2, \theta_2)} \\
\text{(INL)} \\
\frac{v_1 \approx_L v_2}{\mathbf{inl}(v_1) \approx_L \mathbf{inl}(v_2)} \\
\text{INR} \\
\frac{v_1 \approx_L v_2}{\mathbf{inr}(v_1) \approx_L \mathbf{inr}(v_2)} \\
\text{(PAIR)} \\
\frac{v_1 \approx_L v'_1 \quad v_2 \approx_L v'_2}{(v_1, v_2) \approx_L (v'_1, v'_2)} \\
\text{(REF}_L\text{)} \\
\frac{\ell \sqsubseteq L}{n_\ell \approx_L n_\ell} \\
\text{(REF}_H\text{)} \\
\frac{\ell_1 \not\sqsubseteq L \quad \ell_2 \not\sqsubseteq L}{n_{\ell_1} \approx_L n_{\ell_2}}
\end{array}$$

Fig. 5. L -equivalence for λ^{dFG} values and raw values.

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 4 and 5. 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., [Efsthathopoulos et al. 2005; Krohn et al. 2007b; Russo et al. 2008; Stefan et al. 2011]) enforcements are likely to assign a single *flow-insensitive* 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., [Myers et al. 2006; Pottier and Simonet 2002; Stefan et al. 2011; Vassena and Russo 2016]), in a manner similar to what we have just shown for references in λ^{dFG} .

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. L -equivalence for values and raw-values is defined formally by mutual induction in Figure 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 ([UNIT, 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 ($[\text{INL}, \text{INR}]$). In particular, closures are related if they contain the *same* function (up to α -renaming)⁷ 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 $\forall x. \theta_1(x) \approx_L \theta_2(x)$.

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 (rule $[\text{REF}_H]$) but requires the same label and address for public references (rule $[\text{REF}_L]$).

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 $v_1 \approx_L v_2$.

We now formally state and prove that λ^{dFG} semantics preserve L -equivalence under L -equivalent environments, i.e., *termination-insensitive non-interference* (TINI).

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 induction on the derivations.

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 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.

Figure 6 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 calculus λ^{dCG} features both

⁷Symbol \equiv_α denotes α -equivalence. In our mechanized proofs we use De Bruijn indexes and syntactic equivalence.

442	Type:	$\tau ::=$	unit $\tau_1 \rightarrow \tau_2$ $\tau_1 + \tau_2$ $\tau_1 \times \tau_2$ \mathcal{L} LIO τ Labeled τ Ref τ
443	Labels:	$\ell, pc \in$	\mathcal{L}
444	Addresses:	$n \in$	\mathbb{N}
445	Environment:	$\theta \in$	$Var \rightarrow Value$
446	Value:	$v ::=$	$()$ $(x.e, \theta)$ inl (v) inr (v) (v_1, v_2) ℓ Labeled ℓv (t, θ) n_ℓ
447	Expression:	$e ::=$	x $\lambda x.e$ $e_1 e_2$ $()$ ℓ inl (e_1) inr (e_2) case ($e, x.e_1, x.e_2$)
448			(e_1, e_2) fst (e) snd (e) $e_1 \sqsubseteq^? e_2$ t
449	Thunk	$t ::=$	return (e) bind ($e, x.e$) unlabel (e) toLabeled (e) labelOf (e)
450			getLabel taint (e) new (e) $!e$ $e_1 := e_2$ labelOfRef (e)
451	Type System:	$\Gamma \vdash e : \tau$	
452	Configuration:	$c ::=$	$\langle \Sigma, pc, e \rangle$
453	Store:	$\Sigma \in$	$(\ell : Label) \rightarrow Memory \ell$
454	Memory ℓ :	$M ::=$	$[]$ $v : M$
455			

Fig. 6. Syntax of λ^{dCG} .

standard (unlabeled) values and *explicitly labeled* values. For example, **Labeled** H **true** represents a secret boolean value of type **Labeled** **bool**.⁸ 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 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. It also features an operator to query if a label flows to another, written $\ell_1 \sqsubseteq^? \ell_2$. The primitives of the **LIO** monad are listed in a separate sub-category of expressions called *thunk*. Intuitively, a thunk is just a description of a stateful computation, which only the top-level security monitor can execute—a *thunk closure*, i.e., (t, θ) , provides a way to suspend 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, 2008]) and that was later applied in a language-based setting. In this technique, throughout a program's execution, the program counter *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 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 *flow-sensitive*, i.e., it may assume different values in the final configuration depending on the control flow of the program).⁹

⁸As 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**(\cdot).

⁹In 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.

491 Like λ^{dFG} , the operational semantics of λ^{dCG} consists of a security monitor that fully
 492 evaluates secure programs but prevents the execution of insecure programs and similarly
 493 enforces *termination-insensitive* non-interference (Theorem 2). Figure 7 shows the big-step
 494 operational semantics of λ^{dCG} in two parts: (i) a top-level security monitor for monadic
 495 programs and (ii) a straightforward call-by-value side-effect-free semantics for pure expres-
 496 sions. The semantics of the security monitor is further split into two mutually recursive
 497 reduction relations, one for arbitrary expressions (Fig. 7a) and one specific to thunks (Fig. 7c).
 498 These constitute the *forcing* semantics of the monad, which reduce a thunk to a pure value
 499 and perform side-effects. In particular, given the initial store Σ , program counter label
 500 pc , expression e of type **LIO** τ for some type τ and input values θ (which may or may
 501 not be labeled), the monitor executes the program, i.e., $\langle \Sigma, pc, e \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle$ and gives
 502 an updated store Σ' , updated program counter pc' and a final value v of type τ , which
 503 also might not be labeled. The execution starts with rule [FORCE], which reduces the pure
 504 expression to a thunk closure, i.e., (t, θ') and then forces the thunk t in its environment θ'
 505 with the thunk semantics. The pure semantics is fairly standard—we report some selected
 506 rules in Fig. 7b for comparison with λ^{dFG} . A pure reduction, written $e \Downarrow^\theta v$, evaluates
 507 an expression e with an appropriate environment θ to a pure value v . Notice that, unlike
 508 λ^{dFG} , all reduction rules of the pure semantics ignore security, even those that affect the
 509 control flow of the program, e.g., rule [APP]: they do not feature the program counter label
 510 or label annotations. They are also *pure*—they do not have access to the store, thus only
 511 the security monitor needs to protect against *implicit flows*.

512 If the pure evaluation reaches a side-effectful computation, i.e., thunk t , it *suspends* the
 513 computation by creating a thunk closure that captures the current environment θ (see
 514 rule [THUNK]). Notice that *thunk closures* and *function closures* are distinct values created
 515 by different rules, [THUNK] and [FUN] respectively.¹⁰ Function application succeeds only
 516 when the function evaluates to a function closure (rule [APP]). In the thunk semantics, rule
 517 [RETURN] evaluates a pure value embedded in the monad via **return**(\cdot) and leaves the state
 518 unchanged, while rule [BIND] executes the first computation with the forcing semantics,
 519 binds the result in the environment i.e., $\theta[x \mapsto v_1]$, passes it on to the second computation
 520 together with the updated state and returns the final result and state. Rule [UNLABEL] is
 521 interesting. Following the *floating-label* principle, it returns the value wrapped inside the
 522 labeled value, i.e., v , and raises the program counter with its label, i.e., $pc \sqcup \ell$, to reflect
 523 the fact that new data at security level ℓ is now in scope.

524 Floating-label based coarse-grained IFC systems like **LIO** suffer from the *label creep*
 525 problem, which occurs when the program counter gets over-tainted, e.g., because too many
 526 secrets have unlabeled, to the point that no useful further computation can be performed.
 527 Primitive **toLabeled**(\cdot) provides a mechanism to address this problem by (i) creating a
 528 separate context where some sensitive computation can take place and (ii) restoring the
 529 original program counter label afterwards. Rule [TOLABELED] formalizes this idea. Notice
 530 that the result of the nested sensitive computation, i.e., v , cannot be simply returned to
 531 the lower context—that would be a leak, so **toLabeled**(\cdot) wraps that piece of information
 532 in a labeled value protected by the final program counter of the sensitive computation,
 533 i.e., **Labeled** $pc' v$.¹¹ Furthermore, notice that pc' , the label that tags the result v , is as
 534 sensitive as the result itself because the final program counter depends on all the **unlabel**(\cdot)
 535

536 ¹⁰It would have also been possible to define thunk values in terms of function closures using explicit suspension
 537 and an opaque wrapper, e.g., **LIO** ($_t, \theta$).

538 ¹¹Stefan et al. [2017] have proposed an alternative flow-insensitive primitive, i.e., **toLabeled**(ℓ, e), which
 539 labels the result with the user-assigned label ℓ . The semantics of λ^{dFG} forced us to use **toLabeled**(e).

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} .

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 8). 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 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 2 (“the final value of the program counter is at least as sensitive as the initial value”).

PROPERTY 2.

- 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.

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. Figure 9 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 [Labeled_L] relates public labeled values that share the *same* observable label ($\ell \sqsubseteq L$) and contain related values, i.e., $v_1 \approx_L v_2$, while rule [Labeled_H] relates secret labeled values, with arbitrary sensitivity labels not below L ($\ell_1 \not\sqsubseteq L$ 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 $v_1 \approx_L v'_1$ and $\mathbf{inr}(v_2) \approx_L \mathbf{inr}(v'_2)$ iff $v_2 \approx_L v'_2$. Closures and thunks are related if the function and the monadic computations are α -equivalent and their environments are related, i.e., $\theta_1 \approx_L \theta_2$ iff $Dom(\theta_1) \equiv Dom(\theta_2)$ and $\forall x. \theta_1(x) \approx_L \theta_2(x)$. L -equivalence relates labeled references, memories and stores analogously to λ^{dFG} . Related initial configurations have related stores, equal program counters, and

589 α -equivalent programs, i.e., $c_1 \approx_L c_2$ iff $c_1 = \langle \Sigma_1, pc_1, e_1 \rangle$, $c_2 = \langle \Sigma_2, pc_2, e_2 \rangle$, $\Sigma_1 \approx_L \Sigma_2$,
 590 $pc_1 \equiv pc_2$, and $e_1 \equiv_\alpha e_2$. L -equivalence relates final configurations in which either 1) $[PC_L]$
 591 the attacker can observe the same program counter $pc \sqsubseteq L$ in *both* configurations, which
 592 then carry related stores and values, or 2) $[PC_H]$ the value of the program counter in *both*
 593 configuration is not below the attacker level, which thus contain arbitrary values and related
 594 stores.

595 We now formally state and prove that λ^{dCG} semantics preserve L -equivalence under
 596 L -equivalent environments, i.e., *termination-insensitive non-interference* (TINI).

597 **THEOREM 2** (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*
 598 *$c'_1 \approx_L c'_2$.*

600 *Proof.* By induction on the derivations.

601 At this point, we have formalized two calculi— λ^{dFG} and λ^{dCG} —that perform dynamic
 602 IFC at *fine* and *coarse* granularity, respectively. While they have some similarities, i.e., they
 603 are both functional languages that feature labeled annotated data, references and label
 604 introspection primitives, and ensure a termination-insensitive security condition, they also
 605 have striking differences. First and foremost, they differ in the number of label annotations—
 606 pervasive in λ^{dFG} and optional in λ^{dCG} —with significant implications for the programming
 607 model and usability. Second, they differ in the nature of the program counter, *flow-insensitive*
 608 in λ^{dFG} and *flow-sensitive* in λ^{dCG} . Third, they differ in the way they deal with side-effects—
 609 λ^{dCG} allows side-effectful computations exclusively inside the monad, while λ^{dFG} is *impure*,
 610 i.e., any λ^{dFG} expression can modify the state. This difference affects the effort required to
 611 implement a system that performs language-based fine- and coarse-grained dynamic IFC. In
 612 fact, several coarse-grained IFC languages [Buiras et al. 2015; Jaskelioff and Russo 2011;
 613 Russo 2015; Russo et al. 2008; Schmitz et al. 2018; Tsai et al. 2007] have been implemented
 614 as an embedded domain specific language (EDSL) in a Haskell library with little effort,
 615 exploiting the strict control that the host language provides on side-effects. Adapting an
 616 existing language to perform fine-grained IFC requires major engineering effort, because
 617 several components (all the way from the parser to the runtime system) must be adapted to
 618 be label-aware.

619 In the next two sections we show that—despite their differences—these two calculi are, in
 620 fact, equally expressive.

622 4 FINE- TO COARSE-GRAINED PROGRAM TRANSLATION

624 This section presents a provably semantics-preserving program translation from the fine-
 625 grained dynamic IFC calculus λ^{dFG} to the coarse-grained calculus λ^{dCG} . At a high level,
 626 the translation performs two tasks (i) it embeds the *intrinsic* label annotation of λ^{dFG}
 627 values into an *explicitly* labeled λ^{dCG} value via the **Labeled** type constructor and (ii) it
 628 restructures λ^{dFG} *side-effectful* expressions into *monadic operations* inside the **LIO** monad.
 629 Our type-driven approach starts by formalizing this intuition in the function $\langle\langle \cdot \rangle\rangle$, which maps
 630 the λ^{dFG} type τ to the corresponding λ^{dCG} type $\langle\langle \tau \rangle\rangle$ (see Figure 10). The function is defined
 631 by induction on types and recursively adds the **Labeled** type constructor to each existing
 632 λ^{dFG} type constructor. For the function type $\tau_1 \rightarrow \tau_2$, the result is additionally monadic, i.e.,
 633 $\langle\langle \tau_1 \rangle\rangle \rightarrow \mathbf{LIO} \langle\langle \tau_2 \rangle\rangle$. This is because the function's body in λ^{dFG} may have side-effects. The
 634 translation for values (Figure 11) is straightforward. Each λ^{dFG} label tag becomes the label
 635 annotation in a λ^{dCG} labeled value. The translation is homomorphic in constructors on raw
 636 values. The translation converts a λ^{dFG} function closure into a λ^{dCG} closure by translating
 637

the body of the function to a thunk, i.e., $\langle e \rangle$, (see below) and translating the environment pointwise, i.e., $\langle \theta \rangle = \lambda x. \langle \theta(x) \rangle$.

Expressions. We show the translation of λ^{dFG} expressions to λ^{dCG} monadic thunks in Figure 12. We use the standard **do** notation for readability.¹² 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 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 compound 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].¹³ For example, the λ^{dFG} reduction $(((), ())) \Downarrow_L^\emptyset (({}^L, ({}^L))^L$ maps to a λ^{dCG} reduction that yields **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 10 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 $\langle r^\ell \rangle = \mathbf{Labeled} \ell \langle r \rangle$ 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** $\ell \langle r \rangle$, taints the program counter with its label to make it $pc \sqcup \ell$, and returns the content, i.e., $\langle r \rangle$. Since **unlabel**(x) occurs inside a **toLabeled**(\cdot) block, the code above results in **Labeled** $(pc \sqcup \ell) \langle r \rangle$ when evaluated, matching precisely the tainting behavior of λ^{dFG} rule [VAR], i.e., $x \Downarrow_{pc}^{\theta[x \mapsto r^\ell]} r^{pc \sqcup \ell}$.

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., $lw_1 \leftarrow \langle e_1 \rangle$, then the code that computes the argument, i.e., $lw_2 \leftarrow \langle e_2 \rangle$ and then retrieves the function from the first labeled value, i.e., $v_1 \leftarrow \mathbf{unlabel}(lw_1)$.¹⁴ The function v_1 applied to the labeled argument lw_2 gives a computation that gets executed and returns a labeled value lw that gets unlabeled to expose the final result (the surrounding **toLabeled**(\cdot) wraps it again

¹²Syntax **do** $x \leftarrow e_1; e_2$ desugars to **bind**($e_1, x.e_2$) and syntax $e_1; e_2$ to **bind**($e_1, _ .e_2$).

¹³We name a variable lw if it gets bound to a labeled value, i.e., to indicate that the variable has type **Labeled** τ .

¹⁴ Notice 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.

687 right away). The translation of case analysis is analogous. The translation of pair projections
 688 first converts the λ^{dFG} pair into a computation that gives a λ^{dCG} labeled pair of labeled
 689 values, say **Labeled** ℓ (**Labeled** $\ell_1 \langle r_1 \rangle$, **Labeled** $\ell_2 \langle r_2 \rangle$) and removes the label tag on the
 690 pair via **unlabel**, thus raising the program counter label to $pc \sqcup \ell$. Then, it projects the
 691 appropriate component and unlables it, thus tainting the program counter label even further
 692 with the label of either the first or the second component. This coincides with the tainting
 693 mechanism of λ^{dFG} for projection rules, e.g., in rule [FST] where **fst**(e) $\Downarrow_{pc}^\theta r_1^{pc \sqcup \ell \sqcup \ell_1}$ if
 694 $e \Downarrow_{pc}^\theta (r_1^{\ell_1}, r_2^{\ell_2})^\ell$.

695 Lastly, translating **taint**(e_1, e_2) requires (i) translating the expression e_1 that gives the
 696 label, (ii) using **taint**(\cdot) from λ^{dCG} to explicitly taint the program counter label with the
 697 label that e_1 gives, and (iii) translating the second argument e_2 to execute in the tainted
 698 context and unlabeled the result. The construct **labelOf**(e) of λ^{dFG} uses the corresponding
 699 λ^{dCG} primitive applied on the corresponding labeled value, say **Labeled** $\ell \langle r \rangle$, obtained
 700 from the translated expression. Notice that **labelOf**(\cdot) taints the program counter label
 701 in λ^{dCG} , which rises to $pc \sqcup \ell$, so the code just described results in **Labeled** ($pc \sqcup \ell$) ℓ ,
 702 which corresponds to the translation of the result in λ^{dFG} , i.e., $\langle \ell^\ell \rangle = \mathbf{Labeled} \ell \ell$ because
 703 $pc \sqcup \ell \equiv \ell$, since $pc \sqsubseteq \ell$ from Property 1. The translation of **getLabel** follows naturally
 704 by simply wrapping λ^{dCG} 's **getLabel** inside a **toLabeled**(\cdot), which correctly returns the
 705 program counter label labeled with itself, i.e., **Labeled** pc pc .

706
 707 *Note on Environments.* λ^{dFG} and λ^{dCG} semantics feature an environment θ for input
 708 values that gets extended with intermediate values during program evaluation and that may
 709 be captured inside a closure. Unfortunately, this capturing behavior is undesirable for our
 710 program translation. The program translation defined above introduces temporary auxiliary
 711 variables that carry the value of intermediate results, e.g., the labeled value obtained from
 712 running a computation that translates some λ^{dFG} expression. When the translated program
 713 is executed, these values end up in the environment, e.g., by means of rules [APP] and
 714 [BIND], and mix with the input values of the source program and output values as well, thus
 715 complicating the correctness statement of the translation, which now has to account for
 716 those extra variables as well. In order to avoid this nuisance, we employ a special form of
 717 weakening that allows shrinking the environment at run-time and removing spurious values
 718 that are not needed in the rest of the program. In particular, expression *wken* \bar{x} e has the
 719 same type as e if variables \bar{x} are not free in e . At run-time, *wken* \bar{x} e , evaluates e in an
 720 environment from which variables \bar{x} have been dropped, so they do not get captured in
 721 closures created during the execution of e . Formally:

$$\begin{array}{c}
 722 \\
 723 \\
 724 \\
 725 \\
 726 \\
 727 \\
 728 \\
 729 \\
 730 \\
 731 \\
 732 \\
 733 \\
 734 \\
 735
 \end{array}$$

$$\frac{\Gamma \setminus \bar{x} \vdash e : \tau}{\Gamma \vdash \mathit{wken} \bar{x} e : \tau} \qquad \frac{(\text{WKEN}) \quad e \Downarrow^{\theta \setminus \bar{x}} v}{\mathit{wken} \bar{x} e \Downarrow^{\theta} v}$$

728 Rule [WKEN] is part of the pure semantics of λ^{dCG} —the semantics of λ^{dFG} includes an
 729 analogous rule (the issue of contaminated environments arises in the translations in both
 730 directions so both calculi feature *wken*). We remark that this expedient is not essential—we
 731 can avoid it by slightly complicating the correctness statement to explicitly account for those
 732 extra variables. Nor is this expedient particularly interesting. In fact, we omit *wken* from
 733 the code of the program translations to avoid clutter (our mechanization includes *wken* in
 734 the appropriate places).

736 *References.* Figure 13 shows the program translation of λ^{dFG} primitives that access the
 737 store via references. The translation of λ^{dFG} values wraps references in λ^{dCG} labeled values
 738 (Figure 11), so the translations of Figure 13 take care of boxing and unboxing references. The
 739 translation of **new**(e) has a top-level **toLabeled**(\cdot) block that simply translates the content
 740 ($lv \leftarrow \langle e \rangle$) and puts it in a new reference (**new**(lv)). λ^{dCG} assigns the label of the translated
 741 content to the new reference using the λ^{dCG} rule [NEW] (Figure 8), which also gets labeled
 742 with the original program counter label¹⁵, just as in the λ^{dFG} rule [NEW] (Figure 4). In
 743 λ^{dFG} , rule [READ] reads from a reference $n_{\ell'}^{\ell'}$ at security level ℓ' that points to the ℓ -labeled
 744 memory, and returns the content $\Sigma(\ell)[n]^{\ell \sqcup \ell'}$ at level $\ell \sqcup \ell'$. Similarly, the translation creates
 745 a **toLabeled**(\cdot) block that executes to get a labeled reference $lr = \mathbf{Labeled} \ell' n_{\ell}$, extracts
 746 the reference n_{ℓ} ($r \leftarrow \mathbf{unlabel}(lr)$) tainting the program counter label with ℓ' , and then reads
 747 the reference's content further tainting the program counter label with ℓ as well. The code
 748 that translates and updates a reference consists of two **toLabeled**(\cdot) blocks. The first block
 749 is responsible for the update: it extracts the labeled reference and the labeled new content (lr
 750 and lv resp.), extracts the reference from the labeled value ($r \leftarrow \mathbf{unlabel}(lr)$) and updates
 751 it ($r := lv$). The second block, **toLabeled**(**return**(\cdot)), returns unit at security level pc , i.e.,
 752 **Labeled** pc (\cdot), similar to the λ^{dFG} rule [WRITE]. The translation of **labelOfRef**(e) extracts
 753 the reference and projects its label via the λ^{dCG} primitive **labelOfRef**(\cdot), which additionally
 754 taints the program counter with the label itself, similar to the λ^{dFG} rule [LABELOFREF].
 755

756 4.1 Correctness

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

762 We start by showing that the program translation preserves typing. The type translation
 763 for typing contexts Γ is pointwise, i.e., $\langle \Gamma \rangle = \lambda x. \langle \Gamma(x) \rangle$.
 764

765 LEMMA 4.1 (TYPE PRESERVATION). *Given a well-typed λ^{dFG} expression, i.e., $\Gamma \vdash e : \tau$,
 766 the translated λ^{dCG} expression is also well-typed, i.e., $\langle \Gamma \rangle \vdash \langle e \rangle : \mathbf{LIO} \langle \tau \rangle$.*

767 PROOF. By induction on the given typing derivation. □
 768

769 The main correctness criterion for the translation is semantics preservation. Intuitively,
 770 proving this theorem ensures that the program translation preserves the meaning of secure
 771 λ^{dFG} programs when translated and executed with λ^{dCG} semantics (under a translated
 772 environment). In the theorem below, the translation of stores and memories is pointwise,
 773 i.e., $\langle \Sigma \rangle = \lambda \ell. \langle \Sigma(\ell) \rangle$, and $\langle [] \rangle = []$ and $\langle r : M \rangle = \langle r \rangle : \langle M \rangle$ for each ℓ -labeled memory
 774 M . Furthermore, notice that in the translation, the initial and final program counter labels
 775 are the same. This establishes that the program translation preserves the flow-insensitive
 776 program counter label of λ^{dFG} (by means of primitive **toLabeled**(\cdot)).
 777

778 THEOREM 3 (SEMANTICS PRESERVATION OF $\langle \cdot \rangle : \lambda^{dFG} \rightarrow \lambda^{dCG}$). *Given a well-typed
 779 λ^{dFG} program $\langle \Sigma, e \rangle \Downarrow_{pc}^{\theta} \langle \Sigma', v \rangle$, then $\langle \langle \Sigma \rangle, pc, \langle e \rangle \rangle \Downarrow^{\langle \theta \rangle} \langle \langle \Sigma' \rangle, pc, \langle v \rangle \rangle$.*

780 *Proof.* By induction on the given evaluation derivation using basic properties of the security
 781 lattice and of the translation function.
 782

783 ¹⁵ The nested block does not execute any **unlabel**(\cdot) nor **taint**(\cdot).
 784

785 *Recovery of non-interference.* We conclude this section by constructing a proof of termination-
 786 insensitive non-interference for λ^{dFG} (Theorem 1) from the corresponding theorem for λ^{dCG}
 787 (Theorem 2), using the semantics preservation of the translation, together with a property
 788 that the translation preserves L -equivalence. Doing so ensures that the meaning of labels is
 789 preserved by the translation [Barthe et al. 2007; Rajani and Garg 2018]. In the absence of
 790 such an artifact, one could devise a semantics-preserving translation that simply does not
 791 use the security features of the target language. While technically correct (i.e., semantics
 792 preserving), the translation would not be meaningful from the perspective of security.¹⁶

793 The theorem requires a helping lemma that L -equivalence is preserved by the translation.
 794 In the following, we define the translation for *initial* configuration as $\langle\langle c \rangle\rangle^{pc} = \langle\langle \Sigma \rangle\rangle, pc, \langle\langle e \rangle\rangle$
 795 if $c = \langle \Sigma, e \rangle$, and for *final* configurations $\langle\langle c \rangle\rangle^{pc} = \langle\langle \Sigma \rangle\rangle, pc, \langle\langle v \rangle\rangle$ if $c = \langle \Sigma, v \rangle$.

796 LEMMA 4.2. *For all values, raw values, environments and configurations:*

- 797 • $v_1 \approx_L v_2$ if and only if $\langle\langle v_1 \rangle\rangle \approx_L \langle\langle v_2 \rangle\rangle$.
- 798 • $r_1 \approx_L r_2$ if and only if $\langle\langle r_1 \rangle\rangle \approx_L \langle\langle r_2 \rangle\rangle$
- 799 • $\theta_1 \approx_L \theta_2$ if and only if $\langle\langle \theta_1 \rangle\rangle \approx_L \langle\langle \theta_2 \rangle\rangle$
- 800 • Let c_1 and c_2 be initial configurations, then for all pc , $c_1 \approx_L c_2$ if and only if
 801 $\langle\langle c_1 \rangle\rangle^{pc} \approx_L \langle\langle c_2 \rangle\rangle^{pc}$.
- 802 • Let $c'_1 = \langle \Sigma_1, r_1^{\ell_1} \rangle$, $c'_2 = \langle \Sigma_2, r_2^{\ell_2} \rangle$, if $pc \sqsubseteq \ell_1$, $pc \sqsubseteq \ell_2$ and $\langle\langle c'_1 \rangle\rangle^{pc} \approx_L \langle\langle c'_2 \rangle\rangle^{pc}$ then
 803 $c'_1 \approx_L c'_2$.

804 *Proof.* By mutual induction and using injectivity of the translation function in the if
 805 direction.

806 THEOREM 4 (λ^{dFG} -TINI VIA $\langle\langle \cdot \rangle\rangle$). *If $c_1 \Downarrow_{pc}^{\theta_1} c'_1$, $c_2 \Downarrow_{pc}^{\theta_2} c'_2$, $\theta_1 \approx_L \theta_2$ and $c_1 \approx_L c_2$
 807 then $c'_1 \approx_L c'_2$.*

808 *Proof.* We start by applying the fine to coarse grained program translation to the initial
 809 configurations and input values. By Theorem 3 (semantics preservation), we derive the
 810 corresponding λ^{dCG} reductions, i.e., $\langle\langle c_1 \rangle\rangle^{pc} \Downarrow^{\langle\theta\rangle} \langle\langle c'_1 \rangle\rangle^{pc}$ and $\langle\langle c_2 \rangle\rangle^{pc} \Downarrow^{\langle\theta\rangle} \langle\langle c'_2 \rangle\rangle^{pc}$. Then, we
 811 lift L -equivalence for initial configurations and input values to their translation (Lemma 4.2),
 812 i.e., $\langle\langle c_1 \rangle\rangle^{pc} \approx_L \langle\langle c_2 \rangle\rangle^{pc}$ and $\langle\langle \theta_1 \rangle\rangle \approx_L \langle\langle \theta_2 \rangle\rangle$ and obtain $\langle\langle c'_1 \rangle\rangle^{pc} \approx_L \langle\langle c'_2 \rangle\rangle^{pc}$ by Theorem 2 (λ^{dCG} -
 813 TINI). Finally, we deduce L -equivalence of the source final configurations again by the last
 814 point of Lemma 4.2, where $c'_1 = \langle \Sigma_1, r_1^{\ell_1} \rangle$, $c'_2 = \langle \Sigma_2, r_2^{\ell_2} \rangle$ and $pc \sqsubseteq \ell_1$ (resp. $pc \sqsubseteq \ell_2$) by
 815 Property 1 applied to the source reductions, i.e., $c_1 \Downarrow_{pc}^{\theta_1} c'_1$ (resp. $c_2 \Downarrow_{pc}^{\theta_2} c'_2$).
 816
 817
 818

819 5 COARSE- TO FINE-GRAINED PROGRAM TRANSLATION

820 We now show a verified program translation in the opposite direction—from the coarse
 821 grained calculus λ^{dCG} to the fine grained calculus λ^{dFG} . The translation in this direction is
 822 more involved—a program in λ^{dFG} contains strictly more information than its counterpart
 823 in λ^{dCG} , namely the extra *intrinsic* label annotations that tag every value. The challenge
 824 in constructing this translation is two-fold. On one hand, the translation must come up
 825 with labels for all values. However, it is not always possible to do this statically during the
 826 translation: Often, the labels depend on input values and arise at run-time with intermediate
 827 results since the λ^{dFG} calculus is designed to compute and attach labels at run-time. On
 828 the other hand, the translation cannot conservatively under-approximate the values of
 829

830 ¹⁶Note that such bogus translations are also ruled out due to the need to preserve the outcome of any label
 831 introspection. Nonetheless, building this proof artifact increases our confidence in the robustness of our
 832 translation. In contrast, if the enforcement of IFC is *static*, then there is no label introspection, and this
 833 proof artifact is extremely important, as argued in prior work [Barthe et al. 2007; Rajani and Garg 2018].

834 labels¹⁷— λ^{dCG} and λ^{dFG} have label introspection so, in order to get semantics preservation,
 835 labels must be preserved precisely. Intuitively, we solve this impasse by crafting a program
 836 translation that (i) preserves the labels that can be inspected by λ^{dCG} and (ii) lets the
 837 λ^{dFG} semantics compute the remaining label annotations automatically—we account for
 838 those labels with a *cross-language* relation that represents semantic equivalence between
 839 λ^{dFG} and λ^{dCG} modulo extra annotations (Section 5.1). The fact that the source program
 840 in λ^{dCG} cannot inspect those labels—they have no value counterpart in the source λ^{dCG}
 841 program—facilitates this aspect of the translation. We elaborate more on the technical
 842 details later.

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

853 *Types.* The λ^{dCG} -to- λ^{dFG} translation follows the same type-driven approach used in the
 854 other direction, starting from the function $\llbracket \cdot \rrbracket$ in Figure 14, that translates a λ^{dFG} type τ into
 855 the corresponding λ^{dCG} type $\llbracket \tau \rrbracket$. The translation is defined by induction on τ and preserves
 856 all the type constructors standard types. Only the cases corresponding to λ^{dCG} -specific
 857 types are interesting. In particular, it converts *explicitly* labeled types, i.e., **Labeled** τ , to a
 858 standard pair type in λ^{dFG} , i.e., $(\mathcal{L} \times \llbracket \tau \rrbracket)$, where the first component is the label and the
 859 second component the content of type τ . Type **LIO** τ becomes a *suspension* in λ^{dFG} , i.e.,
 860 the function type **unit** $\rightarrow \llbracket \tau \rrbracket$ that delays a computation and that can be forced by simply
 861 applying it to the unit value $()$.

862 *Values.* The translation of values follows the type translation, as shown in Figure 15.
 863 Notice that the translation is indexed by the program counter label (the translation is
 864 written $\llbracket v \rrbracket^{pc}$), which converts the λ^{dCG} value v in scope of a computation protected by
 865 security level pc to the corresponding fully label-annotated λ^{dFG} value. The translation is
 866 pretty straightforward and uses the program counter label to tag each value, following the
 867 λ^{dCG} principle that the program counter label protects every value in scope that is not
 868 explicitly labeled. The translation converts a λ^{dCG} function closure into a corresponding
 869 λ^{dFG} function closure by translating the body of the function to a λ^{dFG} expression (see
 870 below) and translating the environment pointwise, i.e., $\llbracket \theta \rrbracket^{pc} = \lambda x. \llbracket \theta(x) \rrbracket^{pc}$. A thunk value or
 871 a *thunk closure*, which denotes a suspended side-effectful computation, is also converted into
 872 a λ^{dFG} function closure. Technically, the translation would need to introduce a *fresh variable*
 873 that would get bound to unit when the suspension gets forced. However, the argument to
 874 the suspension does not have any purpose, so we do not bother with giving a name to it and
 875 write $_ \llbracket t \rrbracket$ instead. (In our mechanized proofs we employ unnamed De Bruijn indexes and
 876 this issue does not arise.) The translation converts an explicitly labeled value **Labeled** ℓv ,
 877
 878

879 ¹⁷In contrast, previous work on *static* type-based fine-to-coarse grained translation safely under-approximates
 880 the label annotations in types with \perp [Rajani and Garg 2018]. The proof of type preservation of the translation
 881 recovers the actual labels via *subtyping*.

883 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
 884 with itself, and the value translated at a security level equal to the label annotation, i.e., $\llbracket v \rrbracket^\ell$.
 885 Notice that tagging the label with itself allows us to translate the λ^{dCG} (label introspection)
 886 primitive **labelOf**(\cdot) by simply projecting the first component, thus preserving the label
 887 and its security level across the translation.
 888

889 *Expressions and Thunks.* The translation of pure expressions (Figure 16) is trivial: It is
 890 homomorphic in all constructs, mirroring the type translation. The translation of a thunk
 891 expression t builds a suspension explicitly with a λ -abstraction (the name of the variable is
 892 again irrelevant, thus we omit it as explained above), and carries on by translating the thunk
 893 itself according to the definition in Figure 17. The thunk **return**(e) becomes $\llbracket e \rrbracket$, since
 894 **return**(\cdot) does not have any side-effect. When two monadic computations are combined
 895 via **bind**($e_1, x.e_2$), the translation (i) converts the first computation to a suspension and
 896 forces it by applying **unit**($\llbracket e_1 \rrbracket()$), (ii) binds the result to x and passes it to the second
 897 computation¹⁸, which is also converted, forced, and, *importantly*, (iii) executed with a
 898 program counter label tainted with the security level of the result of the first computation
 899 (**taint**(**labelOf**(x), $\llbracket e_2 \rrbracket()$). Notice that **taint**(\cdot) is essential to ensure that the second
 900 computation executes with the program counter label set to the correct value—the precision
 901 of the fine-grained system would otherwise retain the initial lower program counter label
 902 according to rule [APP] and the value of the program counter labels in the source and target
 903 programs would differ in the remaining execution.

904 Similarly, the translation of **unlabel**(e) first translates the labeled expression e (the
 905 translated expression does not need to be forced because it is not of a monadic type),
 906 binds its result to x and then projects the content in a context tainted with its label, as
 907 in **taint**(**fst**(x), **snd**(x)). This closely follows λ^{dCG} 's [UNLABEL] rule. The translation of
 908 **toLabeled**(e) forces the nested computation with $\llbracket e \rrbracket()$, binds its result to x and creates
 909 the pair (**labelOf**(x), x), which corresponds to the labeled value obtained in λ^{dCG} via rule
 910 [TOLABELED]. Intuitively, the translation guarantees that the value of the final program
 911 counter label in the nested computation coincides with the security level of the translated
 912 result (bound to x). Therefore, the first component contains the correct label and it is
 913 furthermore at the right security level, because **labelOf**(\cdot) protects the projected label with
 914 the label itself in λ^{dFG} . Primitive **labelOf**(e) simply projects the first component of the pair
 915 that encodes the labeled value in λ^{dFG} as explained above. Lastly, **getLabel** in λ^{dCG} maps
 916 directly to **getLabel** in λ^{dFG} —rule [GETLABEL] in λ^{dCG} simply returns the program counter
 917 label and does not raise its value, so it corresponds exactly to rule [GETLABEL] in λ^{dFG} ,
 918 which returns label pc at security level pc . Similarly, **taint**(e) translates to **taint**($\llbracket e \rrbracket$, $()$)
 919 since rule [TAINT] in λ^{dCG} simply taints the program counter label with the label that e
 920 evaluates to, say ℓ and results in **unit** with program counter label raised to $pc \sqcup \ell$. This
 921 corresponds to the result of the translated program, i.e., $()^{pc \sqcup \ell}$.
 922

923 *References.* Figure 18 shows the translation of primitives that access the store via references.
 924 Since λ^{dCG} 's rule [NEW] in Figure 8 creates a new reference labeled with the label of the
 925 argument (which must be a labeled value), the translation converts **new**(e) to an expression
 926 that first binds $\llbracket e \rrbracket$ to x and then creates a new reference with the same content as the source,
 927 i.e., **snd**(x), but tainted with the label in x , i.e., **fst**(x). Notice that the use of **taint**(\cdot) is
 928 essential to ensure that λ^{dFG} 's rule [NEW] in Figure 4 assigns the correct label to the new
 929 reference. Due to its *fine-grained* precision, λ^{dFG} might have labeled the content with a

930 ¹⁸Syntax **let** $x = e_1$ **in** e_2 where x is free in e_2 is a shorthand for $(\lambda x.e_2) e_1$.
 931

different label that is less sensitive than the explicit label that *coarsely* approximates the security level in λ^{dCG} . In contrast, updating a reference does not require any tainting—both λ^{dFG} and λ^{dCG} accept values less sensitive than the reference in rule [WRITE]. Thus, the translation $e_1 := e_2$ simply updates the translated reference with the content of the labeled value projected from the translated pair. Hence, $\llbracket e_1 := e_2 \rrbracket$ is $\llbracket e_1 \rrbracket := \mathbf{snd}(\llbracket e_2 \rrbracket)$. The translation of the primitives that read and query the label of a reference is trivial.

5.1 Cross-Language Semantic Equivalence up to Extra Annotations

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 $\llbracket v \rrbracket^{pc'}$.

Example. Consider the λ^{dCG} program $\langle \Sigma, L, \mathbf{taint}(H); \mathbf{return}(x) \rangle \Downarrow^{x \mapsto \mathbf{true}} \langle \Sigma, H, \mathbf{true} \rangle$, which returns $\mathbf{true} = \mathbf{inl}()$ and the store Σ unchanged, after tainting the program counter label with H . Let e be the expression obtained by applying the program translation defined above to the example program:

$$e = \lambda_{-}. \\ \mathbf{let } y = \mathbf{taint}(H, ()) \mathbf{ in} \\ \mathbf{taint}(\mathbf{labelOf}(y), x)$$

Interestingly, when we force the program e and execute it starting from program counter label equal to L , and an input environment translated according to the initial program counter label (L in this case), 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 :

$$\langle \llbracket \Sigma \rrbracket, e () \rangle \Downarrow_L^{x \mapsto \mathbf{true}^L} \langle \llbracket \Sigma \rrbracket, \mathbf{true}^H \rangle = \langle \llbracket \Sigma \rrbracket, \mathbf{inl}(()^L)^H \rangle \neq \langle \llbracket \Sigma \rrbracket, \mathbf{inl}(()^H)^H \rangle = \langle \llbracket \Sigma \rrbracket, \llbracket \mathbf{true} \rrbracket^H \rangle$$

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.¹⁹ Technically, we use this equivalence in the semantics preservation theorem in Section 5.2, 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 19 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]), requires the same injection for values of the sum type ([INL, INR]) and requires the components to be related 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 $x.[e]$ (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 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 $(r_1 : M_1) \downarrow_{\approx} (r_2 : M_2)$ iff $r_1 \downarrow_{\approx \ell} r_2$ and $M_1 \downarrow_{\approx} M_2$ for λ^{dFG} and λ^{dCG} memories $M_1, M_2 : Memory \ell$. Lastly, we lift the relation to initial and final configurations.

DEFINITION 1 (EQUIVALENCE OF CONFIGURATIONS). *For all initial and final configurations:*

- $\langle \Sigma_1, [e]() \rangle \downarrow_{\approx} \langle \Sigma_2, pc, e \rangle$ iff $\Sigma_1 \downarrow_{\approx} \Sigma_2$,
- $\langle \Sigma_1, \llbracket t \rrbracket \rangle \downarrow_{\approx} \langle \Sigma_2, pc, t \rangle$ iff $\Sigma_1 \downarrow_{\approx} \Sigma_2$,
- $\langle \Sigma_1, r^{pc} \rangle \downarrow_{\approx} \langle \Sigma_2, pc, v \rangle$ iff $\Sigma_1 \downarrow_{\approx} \Sigma_2$ and $r \downarrow_{\approx pc} v$.

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 $\llbracket \cdot \rrbracket$ defined above (similar to rules [FUN] and [THUNK] in Figure 19). Furthermore, in the first case (expressions), the relation additionally forces the translated suspension $\llbracket e \rrbracket$ 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 . The third definition relates final configurations whenever the stores are related and the security level of the final λ^{dFG} result corresponds to the program counter label pc of the final λ^{dCG} configuration, and the final λ^{dCG} result corresponds 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 equivalence 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., $\llbracket \langle \Sigma, pc, t \rangle \rrbracket = \langle \llbracket \Sigma \rrbracket, \llbracket t \rrbracket \rangle$ and forcing for suspensions, i.e.,

¹⁹This 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.

1030 $\llbracket \langle \Sigma, pc, e \rangle \rrbracket = \langle \llbracket \Sigma \rrbracket, \llbracket e \rrbracket \rangle$, pointwise for stores, i.e., $\llbracket \Sigma \rrbracket = \lambda \ell. \llbracket \Sigma(\ell) \rrbracket$, and memories, i.e.,
 1031 $\llbracket [] \rrbracket = []$ and $\llbracket v : M \rrbracket = \llbracket v \rrbracket^\ell : \llbracket M \rrbracket$ for ℓ -labeled memory M .

1032 **PROPERTY 3 (REFLEXIVITY).** *For all initial configurations c , $\llbracket c \rrbracket \downarrow \approx c$.*
 1033

1034 *Proof.* The proof is by induction and relies on analogous properties for all syntactic
 1035 categories: for stores, $\llbracket \Sigma \rrbracket \downarrow \approx \Sigma$, for memories, $\llbracket M \rrbracket \downarrow \approx M$, for values $\llbracket v \rrbracket^{pc} \downarrow \approx_{pc} v$, for
 1036 environments $\llbracket \theta \rrbracket^{pc} \downarrow \approx_{pc} \theta$, for any label pc .

1037 The next property guarantees that values and environments remain in the relation when
 1038 the program counter label rises.

1039 **PROPERTY 4 (WEAKENING).** *For all labels pc and pc' such that $pc \sqsubseteq pc'$, and for all
 1040 λ^{dFG} values v_1 and environments θ_1 , and λ^{dCG} values v_2 and environments θ_2 :*

- 1041 • *If $v_1 \downarrow \approx_{pc} v_2$ then $v_1 \downarrow \approx_{pc'} v_2$*
- 1042 • *If $\theta_1 \downarrow \approx_{pc} \theta_2$ then $\theta_1 \downarrow \approx_{pc'} \theta_2$*

1043 *Proof.* By mutual induction on the relation and using basic properties of the lattice.
 1044

1045 5.2 Correctness

1046 With the help of the cross-language relation defined above, we can now state and prove that
 1047 the λ^{dCG} -to- λ^{dFG} translation is correct, i.e., it satisfies a semantics-preservation theorem
 1048 analogous to that proved for the translation in the opposite direction. At a high level,
 1049 the theorem ensures that the translation preserves the meaning of a secure terminating
 1050 λ^{dCG} program when executed under λ^{dFG} semantics, with the same program counter label
 1051 and translated input values. Since the translated λ^{dFG} program computes strictly more
 1052 information than the original λ^{dCG} program, the theorem existentially quantify over the
 1053 λ^{dFG} result, but insists that it is semantically equivalent to the original λ^{dCG} result at a
 1054 security level equal to the final value of the program counter label, using the cross-language
 1055 relation just defined.

1056 We start by proving that the program translation preserves typing.
 1057

1058 **LEMMA 5.1 (TYPE PRESERVATION).** *If $\Gamma \vdash e : \tau$ then $\llbracket \Gamma \rrbracket \vdash \llbracket e \rrbracket : \llbracket \tau \rrbracket$.*
 1059

1060 *Proof.* By straightforward induction on the typing judgment.

1061 Next, we prove semantics preservation of λ^{dCG} pure reductions. Since these reductions do
 1062 not perform any security-relevant operation (they do not read or write state), they can be
 1063 executed with *any* program counter label in λ^{dFG} and do not change the state in λ^{dFG} .

1064 **LEMMA 5.2 ($\llbracket \cdot \rrbracket : \lambda^{dCG} \rightarrow \lambda^{dFG}$ PRESERVES PURE SEMANTICS).** *If $e \Downarrow^\theta v$ then for any
 1065 program counter label pc , λ^{dFG} store Σ , environment θ' such that $\theta' \downarrow \approx_{pc} \theta$, there exists a
 1066 raw value r , such that $\langle \Sigma, \llbracket e \rrbracket \rangle \Downarrow_{pc}^{\theta'} \langle \Sigma, r^{pc} \rangle$ and $r \downarrow \approx_{pc} v$.*
 1067

1068 *Proof.* By induction on the given evaluation derivation and using basic properties of the
 1069 lattice.

1070 Notice that the lemma holds for *any* input target environment θ' in relation with the
 1071 source environment θ at security level pc rather than just for the translated environment
 1072 $\llbracket \theta \rrbracket^{pc}$. Intuitively, we needed to generalize the lemma so that the induction principle is strong
 1073 enough to discharge cases where (i) we need to prove reductions that use an existentially
 1074 quantified environment, e.g., [APP] and (ii) when some intermediate result at a security level
 1075 other than pc gets added to the environment, so the environment is no longer homogenously
 1076 labeled with pc . While the second condition does not arise in pure reductions, it does occur
 1077 in the reduction of monadic expressions considered in the following theorem.

1078

THEOREM 5 ($\llbracket \cdot \rrbracket : \lambda^{dCG} \rightarrow \lambda^{dFG}$ PRESERVES THUNK AND FORCING SEMANTICS).

- Let $c_2 = \langle \Sigma_2, pc, t \rangle$ be an initial λ^{dCG} configuration. If $c_2 \Downarrow^{\theta_2} c'_2$, then for all λ^{dFG} environments θ_1 and initial configurations c_1 such that $\theta_1 \downarrow_{pc} \theta_2$ and $c_1 \downarrow \approx c_2$, there exists a final configuration c'_1 , such that $c_1 \Downarrow_{pc}^{\theta_1} c'_1$ and $c'_1 \downarrow \approx c'_2$.
- Let $c_2 = \langle \Sigma_2, pc, e \rangle$ be an initial λ^{dCG} configuration. If $c_2 \Downarrow^{\theta_2} c'_2$, then for all λ^{dFG} environments θ_1 and initial configurations c_1 such that $\theta_1 \downarrow_{pc} \theta_2$ and $c_1 \downarrow \approx c_2$, there exists a final configuration c'_1 , such that $c_1 \Downarrow_{pc}^{\theta_1} c'_1$ and $c'_1 \downarrow \approx c'_2$.

Proof (Sketch). By mutual induction on the given derivations, using Lemma 5.2 for pure reductions and Properties 2 and 4 in cases [BIND, TOLABELED, UNLABEL, READ], basic properties of the lattice and of the translation function (for operations on the store).

We finally instantiate the semantics-preservation theorem with the translation of the input values and the initial stores at security level pc .

COROLLARY 1 (CORRECTNESS). Let $c_1 = \langle \Sigma, pc, e \rangle$, if $c_1 \Downarrow^{\theta} c'_1$, then there exists a final λ^{dFG} configuration c'_2 such that $\llbracket c_1 \rrbracket \Downarrow_{pc}^{\llbracket \theta \rrbracket pc} c'_2$ and $c'_1 \downarrow \approx c'_2$.

Proof. By Property 3 and Theorem 5.

Notice that 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 $\langle \Sigma_2, pc, e \rangle \Downarrow^{\theta} \langle \Sigma'_2, pc', v \rangle$ then, by Corollary 1 and unrolling Definition 1, there exists a raw value r at security level pc' and a store Σ'_1 , such that $\langle \llbracket \Sigma_2 \rrbracket, \llbracket e \rrbracket \rangle \Downarrow_{pc}^{\llbracket \theta \rrbracket pc} \langle \Sigma'_1, r^{pc'} \rangle$, $r \downarrow_{pc'} \approx v$ and $\Sigma'_1 \downarrow \approx \Sigma'_2$.

Recovery of non-interference. Similarly to our presentation of Theorem 4 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 and the non-interference of λ^{dFG} . By reproving non-interference of the source language from the target language, we show that our program translation is authentic.

The following lemma ensures that the translation of initial configurations preserves L -equivalence.

LEMMA 5.3. If $c_1 \approx_L c_2$, then $\llbracket c_1 \rrbracket \approx_L \llbracket c_2 \rrbracket$.

Proof. By induction on the L -equivalence judgment and proving similar lemmas for values, i.e., if $v_1 \approx_L v_2$ then $\llbracket v_1 \rrbracket^{pc} \approx_L \llbracket v_2 \rrbracket^{pc}$, for environments, i.e., if $\theta_1 \approx_L \theta_2$ then $\llbracket \theta_1 \rrbracket^{pc} \approx_L \llbracket \theta_2 \rrbracket^{pc}$, for any label pc , for memories, i.e., if $M_1 \approx_L M_2$ then $\llbracket M_1 \rrbracket \approx_L \llbracket M_2 \rrbracket$, and for stores, i.e., if $\Sigma_1 \approx_L \Sigma_2$ then $\llbracket \Sigma_1 \rrbracket \approx_L \llbracket \Sigma_2 \rrbracket$.

The following lemmas recovers λ^{dCG} L -equivalence from λ^{dFG} L -equivalence using the cross-language equivalence relation for all the syntactic categories.

LEMMA 5.4. For all 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 , memories M_1, M_2 , stores Σ_1, Σ_2 , and corresponding λ^{dCG} values v'_1, v'_2 and environments θ'_1, θ'_2 , memories M'_1, M'_2 , stores Σ'_1, Σ'_2 :

- If $v_1 \approx_L v_2$, $v_1 \downarrow_{pc} \approx v'_1$ and $v_2 \downarrow_{pc} \approx v'_2$, then $v'_1 \approx_L v'_2$,
- If $r_1 \approx_L r_2$, $r_1 \downarrow_{pc} \approx v'_1$ and $r_2 \downarrow_{pc} \approx v'_2$, then $v'_1 \approx_L v'_2$,
- If $\theta_1 \approx_L \theta_2$, $\theta_1 \downarrow_{pc} \approx \theta'_1$ and $\theta_2 \downarrow_{pc} \approx \theta'_2$, then $\theta'_1 \approx_L \theta'_2$,
- If $M_1 \approx_L M_2$, $M_1 \downarrow \approx M'_1$ and $M_2 \downarrow \approx M'_2$, then $M'_1 \approx_L M'_2$,
- If $\Sigma_1 \approx_L \Sigma_2$, $\Sigma_1 \downarrow \approx \Sigma'_1$ and $\Sigma_2 \downarrow \approx \Sigma'_2$, then $\Sigma'_1 \approx_L \Sigma'_2$.

1128 *Proof.* The lemmas are proved mutually, by induction on the L -equivalence relation and
 1129 the cross-language equivalence relations and using injectivity of the translation function $\llbracket \cdot \rrbracket$
 1130 for closure values.²⁰

1131 The next lemma lifts the previous lemma final configurations.

1132 LEMMA 5.5. *Let c_1 and c_2 be λ^{dFG} final configurations, let c'_1 and c'_2 be λ^{dCG} final*
 1133 *configurations. If $c_1 \approx_L c_2$, $c_1 \downarrow \approx c'_1$ and $c_2 \downarrow \approx c'_2$, then $c'_1 \approx_L c'_2$.*

1135 *Proof.* Let $c_1 = \langle \Sigma_1, v_1 \rangle$, $c_2 = \langle \Sigma_2, v_2 \rangle$, $c'_1 = \langle \Sigma'_1, pc_1, v'_1 \rangle$, $c'_2 = \langle \Sigma'_2, pc_2, v'_2 \rangle$. From L -
 1136 equivalence of λ^{dFG} final configurations, it follows L -equivalence for the stores and the values,
 1137 i.e., $\Sigma_1 \approx_L \Sigma_2$ and $v_1 \approx_L v_2$ from $c_1 \approx_L c_2$ (Section 2.2). Similarly, from cross-language
 1138 equivalence of final λ^{dFG} and λ^{dCG} configurations, it follows cross-language equivalence of
 1139 their components, i.e., respectively $\Sigma_1 \downarrow \approx \Sigma'_1$ and $v_1 \downarrow \approx_{pc_1} v'_1$ from $c_1 \downarrow \approx c'_1$, and $\Sigma_2 \downarrow \approx \Sigma'_2$
 1140 and $v_2 \downarrow \approx_{pc_2} v'_2$ from $c_2 \downarrow \approx c'_2$ (Definition 1). First, we show that the λ^{dCG} stores are
 1141 L -equivalent, i.e., $\Sigma'_1 \approx_L \Sigma'_2$ by Lemma 5.4 for stores, then two cases follow by case split on
 1142 $v_1 \approx_L v_2$. Either (i) both label annotations on the values are not observable ($\llbracket \text{VALUE}_H \rrbracket$),
 1143 then the program counter labels are also not observable ($pc_1 \not\sqsubseteq L$ and $pc_2 \not\sqsubseteq L$ from
 1144 $v_1 \downarrow \approx_{pc_1} v'_1$ and $v_2 \downarrow \approx_{pc_2} v'_2$) and $c'_1 \approx_L c'_2$ by rule $\llbracket \text{PC}_H \rrbracket$ or (ii) the label annotations are
 1145 equal and observable by the attacker ($\llbracket \text{VALUE}_L \rrbracket$), i.e., $pc_1 = pc_2 \sqsubseteq L$, then $v'_1 \approx_L v'_2$ by
 1146 Lemma 5.4 for values and $c'_1 \approx_L c'_2$ by rule $\llbracket \text{PC}_L \rrbracket$.

1147 THEOREM 6 (λ^{dCG} -TINI VIA $\llbracket \cdot \rrbracket$). *If $c_1 \downarrow \theta_1$, c'_1 , $c_2 \downarrow \theta_2$, c'_2 , $\theta_1 \approx_L \theta_2$ and $c_1 \approx_L c_2$,*
 1148 *then $c'_1 \approx_L c'_2$.*

1150 *Proof.* First, we apply the translation $\llbracket \cdot \rrbracket : \lambda^{dCG} \rightarrow \lambda^{dFG}$ to the initial configurations c_1
 1151 and c_2 and the respective environments θ_1 and θ_2 . Let pc be the initial program counter
 1152 label common to configurations c_1 and c_2 (it is the same because $c_1 \approx_L c_2$). Corollary 1
 1153 (Correctness) then ensures that there exist two λ^{dFG} configurations c''_1 and c''_2 , such that
 1154 $\llbracket c_1 \rrbracket \downarrow_{pc}^{\llbracket \theta_1 \rrbracket^{pc}} c''_1$ 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
 1155 source configurations and environments to L -equivalence in the target language via Lemma
 1156 5.3, i.e., $\llbracket \theta_1 \rrbracket^{pc} \approx_L \llbracket \theta_2 \rrbracket^{pc}$ and $\llbracket c_1 \rrbracket \approx_L \llbracket c_2 \rrbracket$, and apply Theorem 1 (λ^{dFG} -TINI) to the
 1157 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
 1158 configurations, i.e., $c''_1 \approx_L c''_2$. Then, we apply Lemma 5.5 to $c''_1 \approx_L c''_2$, $c''_1 \downarrow \approx c'_1$, and
 1159 $c''_2 \downarrow \approx c'_2$, and recover L -equivalence for the source configurations, i.e., $c'_1 \approx_L c'_2$.

1161 6 RELATED WORK

1162 Systematic study of the relative expressiveness of fine- and coarse-grained information flow
 1163 control (IFC) systems has started only recently. Rajani et al. [2017] initiated this study
 1164 in the context of *static* coarse- and fine-grained IFC, enforced via type systems. In more
 1165 recent work, Rajani and Garg [2018] show that a fine-grained IFC type system, which they
 1166 call FG, and two variants of a coarse-grained IFC type system, which they call CG, are
 1167 equally expressive. Their approach is based on type-directed translations, which are type-
 1168 and semantics-preserving. For proofs, they develop logical relations models of FG and the

1170 ²⁰Technically, the function $\llbracket \cdot \rrbracket$ presented in Section 5 is not injective. For example, consider the type translation
 1171 function from Figure 14: $\llbracket \text{Labeled unit} \rrbracket = \mathcal{L} \times \text{unit} = \llbracket \mathcal{L} \times \text{unit} \rrbracket$ but $\text{Labeled unit} \neq \mathcal{L} \times \text{unit}$,
 1172 and $\llbracket \text{LIO unit} \rrbracket = \text{unit} \rightarrow \text{unit} = \llbracket \text{unit} \rightarrow \text{unit} \rrbracket$ but $\text{LIO unit} \neq \text{unit} \rightarrow \text{unit}$. We make the translation
 1173 injective by (i) adding a wrapper type $\text{Id } \tau$ to λ^{dFG} , together with constructor $\text{Id}(e)$, a deconstructor
 1174 $\text{unId}(e)$ and raw value $\text{Id}(v)$, and (ii) tagging security-relevant types and terms with the wrapper, i.e.,
 1175 $\llbracket \text{Labeled } \tau \rrbracket = \text{Id } (\mathcal{L} \times \llbracket \tau \rrbracket)$ and $\llbracket \text{LIO } \tau \rrbracket = \text{Id } (\text{unit} \rightarrow \llbracket \tau \rrbracket)$. Adapting the translations in both directions
 1176 is tedious but straightforward; we refer the interested reader to our mechanized proofs for details.

1177 two variants of CG, as well as cross-language logical relations. Our work and some of our
1178 techniques are directly inspired by their work, but we examine *dynamic* IFC systems based
1179 on runtime monitors. As a result, our technical development is completely different. In
1180 particular, in our work we handle label introspection, which has no counterpart in the earlier
1181 work on static IFC systems, and which also requires significant care in translations. Our
1182 dynamic setting also necessitated the use of tainting operators in both the fine-grained and
1183 the coarse-grained systems.

1184 Our coarse-grained system λ^{dCG} is the dynamic analogue of the second variant of [Rajani](#)
1185 [and Garg \[2018\]](#)'s CG type system. This variant is described only briefly in their paper (in
1186 Section 4, paragraph "Original HLIO") but covered extensively in Part-II of the paper's
1187 appendix. [Rajani and Garg \[2018\]](#) argue that translating their fine-grained system FG to this
1188 variant of CG is very difficult and requires significant use of parametric label polymorphism.
1189 The astute reader may wonder why we do not encounter the same difficulty in translating
1190 our fine-grained system λ^{dFG} to λ^{dCG} . The reason for this is that our fine-grained system
1191 λ^{dFG} is not a direct dynamic analogue of [Rajani and Garg \[2018\]](#)'s FG. In λ^{dFG} , a value
1192 constructed in a context with program counter label pc automatically receives the security
1193 label pc . In contrast, in [Rajani and Garg \[2018\]](#)'s FG, all introduction rules create values
1194 (statically) labeled \perp . Hence, leaving aside the static-vs-dynamic difference, FG's labels are
1195 more precise than λ^{dFG} 's, and this makes [Rajani and Garg \[2018\]](#)'s FG to CG translation
1196 more difficult than our λ^{dFG} to λ^{dCG} translation. In fact, in earlier work, [Rajani et al. \[2017\]](#)
1197 introduced a different type system called FG^- , a static analogue of λ^{dFG} that labels all
1198 constructed values with pc (statically), and noted that translating it to the second variant
1199 of CG is much easier (in the static setting).

1200 Coarse-grained dynamic IFC systems are prevalent in security research in operating
1201 systems [[Efstathopoulos et al. 2005](#); [Krohn et al. 2007a](#); [Zeldovich et al. 2006](#)]. These ideas
1202 have also been successfully applied to other domains, e.g., the web [[Bauer et al. 2015](#); [Giffin](#)
1203 [et al. 2012](#); [Stefan et al. 2014](#); [Yip et al. 2009](#)], mobile applications [[Jia et al. 2013](#); [Nadkarni](#)
1204 [et al. 2016](#)], and IoT [[Fernandes et al. 2016](#)]. LIO is a domain-specific language embedded in
1205 Haskell that rephrases OS-like IFC enforcement into a language-based setting [[Stefan et al.](#)
1206 [2012, 2011](#)]. [Heule et al. \[2015\]](#) introduce a general framework for coarse-grained IFC in any
1207 programming language in which external effects can be controlled. [Laminar \[Roy et al. 2009\]](#)
1208 unifies mechanisms for IFC in programming languages and operating systems, resulting in a
1209 mix of dynamic fine- and coarse-grained enforcement.

1210 In general, dynamic fine-grained IFC systems often do not support label introspection.
1211 LIO [[Stefan et al. 2017, 2011](#)] and [Breeze \[Hritcu et al. 2013b\]](#) are notable exceptions.
1212 [Breeze](#) is conceptually similar to our λ^{dFG} except for the `taint`(\cdot) primitive. Different from
1213 our λ^{dFG} , there are dynamic fine-grained IFC systems in which labels of references are
1214 flow-sensitive [[Austin and Flanagan 2009, 2010](#); [Bichhawat et al. 2014](#); [Hedin et al. 2014](#)].
1215 This design choice, however, allows label changes to be exploited as a covert channel for
1216 information leaks [[Austin and Flanagan 2009, 2010](#); [Russo and Sabelfeld 2010](#)]. There are
1217 many approaches to preventing such leaks—from using static analysis techniques [[Sabelfeld](#)
1218 [and Myers 2003](#)], to disallowing label upgrades depending on sensitive data (i.e., no-sensitive-
1219 upgrades [[Austin and Flanagan 2009](#); [Zdancewic 2002](#)]), to avoiding branching on data
1220 whose labels have been upgraded (i.e., permissive-upgrades [[Austin and Flanagan 2010](#)]).
1221 Extending our results to a fine-grained dynamic IFC system with flow-sensitive references is
1222 an interesting direction for future work.

1223
1224
1225

1226 7 CONCLUSION

1227 We formally established a connection between dynamic fine- and coarse-grained enforcement
1228 for IFC, showing that both are equally expressive under reasonable assumptions. Indeed,
1229 this work provides a systematic way to bridging the gap between a wide range of dynamic
1230 IFC techniques often proposed by the programming languages (fine-grained) and operating
1231 systems (coarse-grained) communities. As consequence, this allows future designs of dynamic
1232 IFC to choose a coarse-grained approach, which is easier to implement and use, without
1233 giving up on the precision of fine-grained IFC.
1234

1235 ACKNOWLEDGMENTS

1236 We thanks the anonymous POPL and POPL AEC reviewers for the insightful comments.
1237 This work was funded by the Swedish Foundation for Strategic Research (SSF) under the
1238 project Octopi (Ref. RIT17-0023) and WebSec (Ref. RIT17-0011) as well as the Swedish
1239 research agency Vetenskapsrådet. Vineet Rajani was partly funded through the Collaborative
1240 Research Center “Methods and Tools for Understanding and Controlling Privacy” (SFB
1241 1223) of the DFG, project “Programming Principles and Abstractions for Privacy.” This
1242 material is based upon work supported by the National Science Foundation under Grant
1243 No. nnnnnnn and Grant No. mmmmmmm. Any opinions, findings, and conclusions or
1244 recommendations expressed in this material are those of the author and do not necessarily
1245 reflect the views of the National Science Foundation. This work was supported in part
1246 by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research
1247 Corporation (SRC) program sponsored by DARPA.
1248

1249 REFERENCES

- 1250 T. H. Austin and C. Flanagan. 2009. Efficient Purely-Dynamic Information Flow Analysis. In *Proc. ACM*
1251 *Workshop on Programming Languages and Analysis for Security (PLAS)*.
- 1252 Thomas H. Austin and Cormac Flanagan. 2010. Permissive dynamic information flow analysis. In *Proceedings*
1253 *of the 5th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security (PLAS '10)*.
1254 ACM.
- 1255 Gilles Barthe, Tamara Rezk, and Amitabh Basu. 2007. Security types preserving compilation. *Computer*
1256 *Languages, Systems & Structures* 33, 2 (2007), 35–59. <https://doi.org/10.1016/j.cl.2005.05.002>
- 1257 Lujo Bauer, Shaoying Cai, Limin Jia, Timothy Passaro, Michael Stroucken, and Yuan Tian. 2015. Run-time
1258 monitoring and formal analysis of information flows in Chromium. In *Proceedings of the 22nd Annual*
Network & Distributed System Security Symposium. Internet Society.
- 1259 Abhishek Bichhawat, Vineet Rajani, Deepak Garg, and Christian Hammer. 2014. Information Flow Control
1260 in WebKit’s JavaScript Bytecode. In *International Conference on Principles of Security and Trust*
(POST). 159–178.
- 1261 Niklas Broberg, Bart van Delft, and David Sands. 2013. Paragon for Practical Programming with Information-
1262 Flow Control. In *Programming Languages and Systems - 11th Asian Symposium, APLAS 2013, Melbourne,*
1263 *VIC, Australia, December 9-11, 2013. Proceedings*. 217–232.
- 1264 Pablo Buiras, Deian Stefan, and Alejandro Russo. 2014. On Dynamic Flow-Sensitive Floating-Label Systems.
1265 In *Proceedings of the 2014 IEEE 27th Computer Security Foundations Symposium (CSF '14)*. IEEE
1266 Computer Society, Washington, DC, USA, 65–79. <https://doi.org/10.1109/CSF.2014.13>
- 1267 Pablo Buiras, Dimitrios Vytiniotis, and Alejandro Russo. 2015. HLIO: Mixing static and dynamic typing for
1268 information-flow control in Haskell. In *ACM SIGPLAN Notices*, Vol. 50. ACM, 280–288.
- 1269 Petros Efstathopoulos, Maxwell Krohn, Steve VanDeBogart, Cliff Frey, David Ziegler, Eddie Kohler, David
1270 Mazières, Frans Kaashoek, and Robert Morris. 2005. Labels and event processes in the asbestos operating
1271 system. In *Proc. of the twentieth ACM symp. on Operating systems principles (SOSP '05)*. ACM.
- 1272 Matthias Felleisen. 1991. On the Expressive Power of Programming Languages. *Sci. Comput. Program.* 17,
1273 1-3 (Dec. 1991), 35–75. [https://doi.org/10.1016/0167-6423\(91\)90036-W](https://doi.org/10.1016/0167-6423(91)90036-W)
- 1274 Earlene Fernandes, Justin Paupore, Amir Rahmati, Daniel Simionato, Mauro Conti, and Atul Prakash.
2016. FlowFence: Practical Data Protection for Emerging IoT Application Frameworks.. In *USENIX*

- 1275 *Security Symposium*. 531–548.
- 1276 Daniel B. Giffin, Amit Levy, Deian Stefan, David Terei, David Mazières, John C. Mitchell, and Alejandro
1277 Russo. 2012. Hails: Protecting Data Privacy in Untrusted Web Applications. In *10th USENIX Symposium*
1278 *on Operating Systems Design and Implementation, OSDI*.
- 1279 J.A. Goguen and J. Meseguer. 1982. Security policies and security models. In *Proc of IEEE Symposium on*
1280 *Security and Privacy*. IEEE Computer Society.
- 1281 D. Hedin, A. Birgisson, L. Bello, and A. Sabelfeld. 2014. JSFlow: Tracking information flow in JavaScript
1282 and its APIs. In *Proc. of the ACM Symposium on Applied Computing (SAC '14)*. ACM.
- 1283 Stefan Heule, Deian Stefan, Edward Z. Yang, John C. Mitchell, and Alejandro Russo. 2015. IFC Inside:
1284 Retrofitting Languages with Dynamic Information Flow Control. In *Proceedings of the Conference on*
1285 *Principles of Security and Trust*. Springer.
- 1286 C. Hritcu, M. Greenberg, B. Karel, B. C. Peirce, and G. Morrisett. 2013a. All Your IFCException Are
1287 Belong to Us. In *Proc. of the IEEE Symposium on Security and Privacy*. IEEE Computer Society.
- 1288 Catalin Hritcu, Michael Greenberg, Ben Karel, Benjamin C Pierce, and Greg Morrisett. 2013b. All your
1289 IFCException are belong to us. In *Security and Privacy (SP), 2013 IEEE Symposium on*. IEEE, 3–17.
- 1290 M. Jaskelioff and A. Russo. 2011. Secure multi-execution in Haskell. In *Proc. Andrei Ershov International*
1291 *Conference on Perspectives of System Informatics (LNCS)*. Springer-Verlag.
- 1292 Limin Jia, Jassim Aljuraidan, Elli Fragkaki, Lujo Bauer, Michael Stroucken, Kazuhide Fukushima, Shinsaku
1293 Kiyomoto, and Yutaka Miyake. 2013. Run-Time Enforcement of Information-Flow Properties on Android
1294 (Extended Abstract). In *Computer Security—ESORICS 2013: 18th European Symposium on Research in*
1295 *Computer Security*. Springer.
- 1296 Maxwell Krohn, Alexander Yip, Micah Brodsky, Natan Cliffer, M. Frans Kaashoek, Eddie Kohler, and
1297 Robert Morris. 2007a. Information Flow Control for Standard OS Abstractions. In *Proc. of the 21st*
1298 *Symp. on Operating Systems Principles*.
- 1299 Maxwell N. Krohn, Alexander Yip, Micah Z. Brodsky, Natan Cliffer, M. Frans Kaashoek, Eddie Kohler, and
1300 Robert Tappan Morris. 2007b. Information flow control for standard OS abstractions. In *Proc. of the*
1301 *21st ACM Symposium on Operating Systems Principles*. 321–334.
- 1302 Andrew C. Myers, Lantian Zheng, Steve Zdancewic, Stephen Chong, and Nathaniel Nystrom. 2006. Jif 3.0:
1303 Java information flow. (July 2006). <http://www.cs.cornell.edu/jif>
- 1304 Adwait Nadkarni, Benjamin Andow, William Enck, and Somesh Jha. 2016. Practical DIFC Enforcement on
1305 Android.. In *USENIX Security Symposium*. 1119–1136.
- 1306 F. Pottier and V. Simonet. 2002. Information Flow Inference for ML. In *Proc. ACM Symp. on Principles of*
1307 *Programming Languages*. 319–330.
- 1308 Vineet Rajani, Iulia Bastys, Willard Rafnsson, and Deepak Garg. 2017. Type systems for information flow
1309 control: the question of granularity. *SIGLOG News* 4, 1 (2017), 6–21.
- 1310 Vineet Rajani and Deepak Garg. 2018. Types for Information Flow Control: Labeling Granularity and
1311 Semantic Models. In *Proc. of the IEEE Computer Security Foundations Symp. (CSF '18)*. IEEE Computer
1312 Society.
- 1313 Indrajit Roy, Donald E Porter, Michael D Bond, Kathryn S McKinley, and Emmett Witchel. 2009. *Laminar:*
1314 *Practical fine-grained decentralized information flow control*. Vol. 44. ACM.
- 1315 Alejandro Russo. 2015. Functional Pearl: Two Can Keep a Secret, if One of Them Uses Haskell. In *Proc. of*
1316 *the 20th ACM SIGPLAN International Conference on Functional Programming (ICFP 2015)*. ACM.
- 1317 A. Russo, K. Claessen, and J. Hughes. 2008. A library for light-weight information-flow security in Haskell.
1318 In *Proc. ACM SIGPLAN symposium on Haskell (HASKELL '08)*. ACM.
- 1319 Alejandro Russo and Andrei Sabelfeld. 2010. Dynamic vs. Static Flow-Sensitive Security Analysis. In *Proc.*
1320 *of the 2010 23rd IEEE Computer Security Foundations Symp. (CSF '10)*. IEEE Computer Society, 186–199.
- 1321 A. Sabelfeld and A. C. Myers. 2003. Language-Based Information-Flow Security. *IEEE J. Selected Areas in*
1322 *Communications* 21, 1 (Jan. 2003), 5–19.
- 1323 Thomas Schmitz, Maximilian Algehed, Cormac Flanagan, and Alejandro Russo. 2018. Faceted Secure Multi
1324 Execution. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications*
1325 *Security (CCS '18)*. ACM, New York, NY, USA, 1617–1634. <https://doi.org/10.1145/3243734.3243806>
- 1326 Deian Stefan, Alejandro Russo, Pablo Buiras, Amit Levy, John C. Mitchell, and David Mazières. 2012.
1327 Addressing Covert Termination and Timing Channels in Concurrent Information Flow Systems. In
1328 *International Conference on Functional Programming (ICFP)*. ACM SIGPLAN.
- 1329 Deian Stefan, Alejandro Russo, David Mazières, and John C. Mitchell. 2017. Flexible Dynamic Information
1330 Flow Control in the Presence of Exceptions. *Journal of Functional Programming* 27 (2017).

- 1324 D. Stefan, A. Russo, J. C. Mitchell, and D. Mazières. 2011. Flexible Dynamic Information Flow Control in
1325 Haskell. In *Proc. of the ACM SIGPLAN Haskell symposium (HASKELL '11)*.
- 1326 Deian Stefan, Edward Z. Yang, Petr Marchenko, Alejandro Russo, Dave Herman, Brad Karp, and David
1327 Mazières. 2014. Protecting Users by Confining JavaScript with COWL. In *Symposium on Operating
1328 Systems Design and Implementation (OSDI)*. USENIX.
- 1329 T. C. Tsai, A. Russo, and J. Hughes. 2007. A Library for Secure Multi-threaded Information Flow in Haskell.
1330 In *Proc. IEEE Computer Security Foundations Symposium (CSF '07)*.
- 1331 Marco Vassena and Alejandro Russo. 2016. On Formalizing Information-Flow Control Libraries. In *Proceedings
1332 of the 2016 ACM Workshop on Programming Languages and Analysis for Security (PLAS '16)*. ACM,
1333 New York, NY, USA, 15–28. <https://doi.org/10.1145/2993600.2993608>
- 1334 Marco Vassena, Alejandro Russo, Pablo Buiras, and Lucas Waye. 2017. MAC A Verified Static Information-
1335 Flow Control Library. *Journal of Logical and Algebraic Methods in Programming* (2017). <https://doi.org/10.1016/j.jlamp.2017.12.003>
- 1336 Dennis Volpano and Geoffrey Smith. 1997. Eliminating Covert Flows with Minimum Typings. In *Proc. of
1337 the 10th IEEE workshop on Computer Security Foundations (CSFW '97)*. IEEE Computer Society.
- 1338 Jean Yang, Kuat Yessenov, and Armando Solar-Lezama. 2012. A language for automatically enforcing
1339 privacy policies. In *ACM SIGPLAN Notices*, Vol. 47. ACM, 137–148.
- 1340 Alexander Yip, Neha Narula, Maxwell Krohn, and Robert Morris. 2009. Privacy-preserving Browser-side
1341 Scripting with BFlow. In *Proceedings of the 4th ACM European Conference on Computer Systems
1342 (EuroSys '09)*. ACM.
- 1343 Stephan Arthur Zdancewic. 2002. *Programming languages for information security*. Ph.D. Dissertation.
1344 Cornell University.
- 1345 Nickolai Zeldovich, Silas Boyd-Wickizer, Eddie Kohler, and David Mazières. 2006. Making information flow
1346 explicit in HiStar. In *Proc. of the 7th USENIX Symp. on Operating Systems Design and Implementation*.
1347 USENIX.
- 1348 Nickolai Zeldovich, Silas Boyd-Wickizer, and David Mazières. 2008. Securing Distributed Systems with
1349 Information Flow Control. (April 2008), 293–308 pages.
- 1350
- 1351
- 1352
- 1353
- 1354
- 1355
- 1356
- 1357
- 1358
- 1359
- 1360
- 1361
- 1362
- 1363
- 1364
- 1365
- 1366
- 1367
- 1368
- 1369
- 1370
- 1371
- 1372

(FORCE)

$$\frac{e \Downarrow^\theta (t, \theta') \quad \langle \Sigma, pc, t \rangle \Downarrow^{\theta'} \langle \Sigma', pc', v \rangle}{\langle \Sigma, pc, e \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle}$$

(a) Forcing semantics: $\langle \Sigma, pc, e \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle$.

(THUNK)

$$t \Downarrow^\theta (t, \theta)$$

(FUN)

$$\lambda x.e \Downarrow^\theta (x.e, \theta)$$

(VAR)

$$x \Downarrow^\theta \theta(x)$$

(APP)

$$\frac{e_1 \Downarrow^\theta (x.e, \theta') \quad e_2 \Downarrow^\theta v_2 \quad e \Downarrow^{\theta'[x \mapsto v_2]} v}{e_1 e_2 \Downarrow^\theta v}$$

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

(RETURN)

$$\frac{e \Downarrow^\theta v}{\langle \Sigma, pc, \mathbf{return}(e) \rangle \Downarrow^\theta \langle \Sigma, pc, v \rangle}$$

(BIND)

$$\frac{\langle \Sigma, pc, e_1 \rangle \Downarrow^\theta \langle \Sigma', pc', v_1 \rangle \quad \langle \Sigma', pc', e_2 \rangle \Downarrow^{\theta[x \mapsto v_1]} \langle \Sigma'', pc'', v \rangle}{\langle \Sigma, pc, \mathbf{bind}(e_1, x.e_2) \rangle \Downarrow^\theta \langle \Sigma'', pc'', v \rangle}$$

(TOLABELED)

$$\frac{\langle \Sigma, pc, e \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle}{\langle \Sigma, pc, \mathbf{toLabeled}(e) \rangle \Downarrow^\theta \langle \Sigma', pc, \mathbf{Labeled} \ pc' \ v \rangle}$$

(UNLABEL)

$$\frac{e \Downarrow^\theta \mathbf{Labeled} \ \ell \ v}{\langle \Sigma, pc, \mathbf{unlabel}(e) \rangle \Downarrow^\theta \langle \Sigma', pc \sqcup \ell, v \rangle}$$

(LABELOF)

$$\frac{e \Downarrow^\theta \mathbf{Labeled} \ \ell \ v}{\langle \Sigma, pc, \mathbf{labelOf}(e) \rangle \Downarrow^\theta \langle \Sigma, pc \sqcup \ell, \ell \rangle}$$

(GETLABEL)

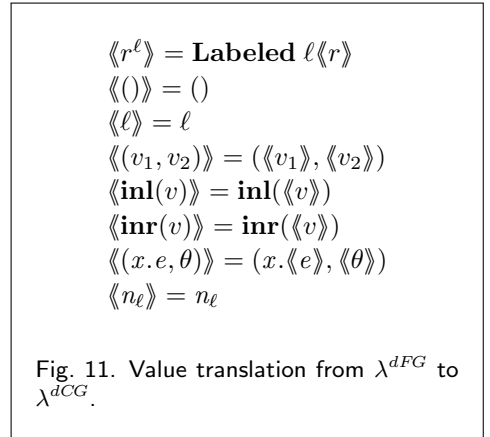
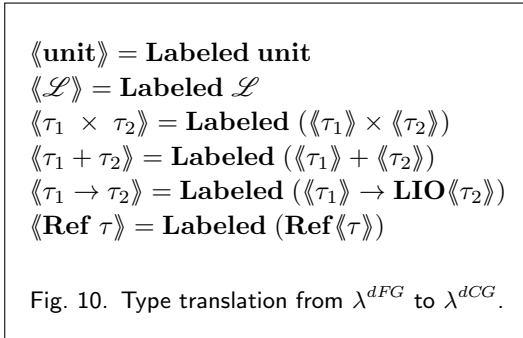
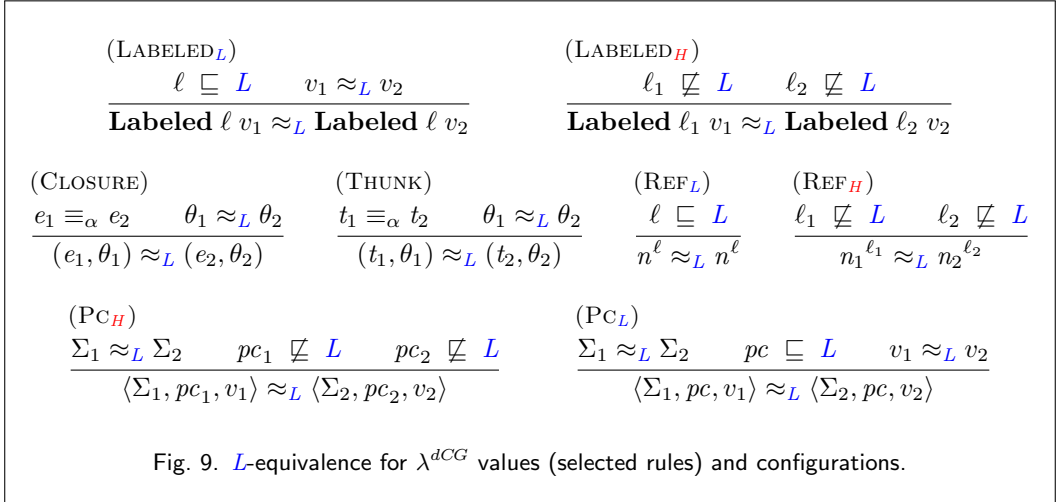
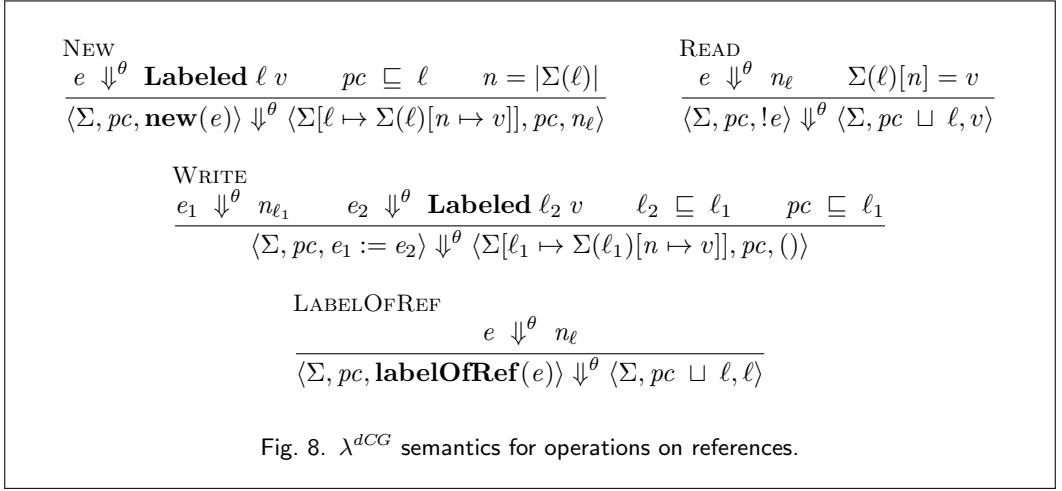
$$\langle \Sigma, pc, \mathbf{getLabel} \rangle \Downarrow^\theta \langle \Sigma, pc, pc \rangle$$

(TAINT)

$$\frac{e \Downarrow^\theta \ell}{\langle \Sigma, pc, \mathbf{taint}(e) \rangle \Downarrow^\theta \langle \Sigma, pc \sqcup \ell, () \rangle}$$

(c) Think semantics: $\langle \Sigma, pc, t \rangle \Downarrow^\theta \langle \Sigma', pc', v \rangle$.Fig. 7. Semantics of λ^{dCG} .

1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470



1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519

$\langle\langle () \rangle\rangle = \text{toLabeled}(\text{return}(()))$	$\langle\langle \text{case}(e, x.e_1, x.e_2) \rangle\rangle = \text{toLabeled}(\text{ do}$
$\langle\langle \ell \rangle\rangle = \text{toLabeled}(\text{return}(\ell))$	$lv \leftarrow \langle\langle e \rangle\rangle$
$\langle\langle \lambda x.e \rangle\rangle = \text{toLabeled}(\text{return}(\lambda x.\langle\langle e \rangle\rangle))$	$v \leftarrow \text{unlabel}(lv)$
$\langle\langle \text{inl}(e) \rangle\rangle = \text{toLabeled}(\text{ do}$	$lv' \leftarrow \text{case}(v, x.\langle\langle e_1 \rangle\rangle, x.\langle\langle e_2 \rangle\rangle)$
$le \leftarrow \langle\langle e \rangle\rangle$	$\text{unlabel}(lv')$
$\text{return}(\text{inl}(lv))$	$\langle\langle \text{fst}(e) \rangle\rangle = \text{toLabeled}(\text{ do}$
$\langle\langle \text{inr}(e) \rangle\rangle = \text{toLabeled}(\text{ do}$	$lv \leftarrow \langle\langle e \rangle\rangle$
$le \leftarrow \langle\langle e \rangle\rangle$	$v \leftarrow \text{unlabel}(lv)$
$\text{return}(\text{inr}(lv))$	$\text{unlabel}(\text{fst}(v))$
$\langle\langle (e_1, e_2) \rangle\rangle = \text{toLabeled}(\text{ do}$	$\langle\langle \text{snd}(e) \rangle\rangle = \text{toLabeled}(\text{ do}$
$lv_1 \leftarrow \langle\langle e_1 \rangle\rangle$	$lv \leftarrow \langle\langle e \rangle\rangle$
$lv_2 \leftarrow \langle\langle e_2 \rangle\rangle$	$v \leftarrow \text{unlabel}(lv)$
$\text{return}(lv_1, lv_2)$	$\text{unlabel}(\text{snd}(v))$
$\langle\langle x \rangle\rangle = \text{toLabeled}(\text{unlabel}(x))$	$\langle\langle \text{taint}(e_1, e_2) \rangle\rangle = \text{toLabeled}(\text{ do}$
$\langle\langle e_1 e_2 \rangle\rangle = \text{toLabeled}(\text{ do}$	$lv_1 \leftarrow \langle\langle e_1 \rangle\rangle$
$lv_1 \leftarrow \langle\langle e_1 \rangle\rangle$	$v_1 \leftarrow \text{unlabel}(lv_1)$
$lv_2 \leftarrow \langle\langle e_2 \rangle\rangle$	$\text{taint}(v_1)$
$v_1 \leftarrow \text{unlabel}(lv_1)$	$lv_2 \leftarrow \langle\langle e_2 \rangle\rangle$
$lv \leftarrow v_1 lv_2$	$\text{unlabel}(lv_2)$
$\text{unlabel}(lv)$	$\langle\langle \text{labelOf}(e) \rangle\rangle = \text{toLabeled}(\text{ do}$
	$lv \leftarrow \langle\langle e \rangle\rangle$
	$\text{labelOf}(lv)$
	$\langle\langle \text{getLabel} \rangle\rangle = \text{toLabeled}(\text{getLabel})$

Fig. 12. Expression translation from λ^{dFG} to λ^{dCG} .

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

Fig. 13. λ^{dFG} to λ^{dCG} translation of memory operations.

1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568

$$\begin{aligned} \llbracket \mathcal{L} \rrbracket &= \mathcal{L} \\ \llbracket \mathbf{unit} \rrbracket &= \mathbf{unit} \\ \llbracket \tau_1 \rightarrow \tau_2 \rrbracket &= \llbracket \tau_1 \rrbracket \rightarrow \llbracket \tau_2 \rrbracket \\ \llbracket \tau_1 + \tau_2 \rrbracket &= \llbracket \tau_1 \rrbracket + \llbracket \tau_2 \rrbracket \\ \llbracket \tau_1 \times \tau_2 \rrbracket &= \llbracket \tau_1 \rrbracket \times \llbracket \tau_2 \rrbracket \\ \llbracket \mathbf{Ref} \tau \rrbracket &= \mathbf{Ref} \llbracket \tau \rrbracket \\ \llbracket \mathbf{Labeled} \tau \rrbracket &= \mathcal{L} \times \llbracket \tau \rrbracket \\ \llbracket \mathbf{LIO} \tau \rrbracket &= \mathbf{unit} \rightarrow \llbracket \tau \rrbracket \end{aligned}$$

Fig. 14. Type translation from λ^{dCG} to λ^{dFG} .

$$\begin{aligned} \llbracket () \rrbracket^{pc} &= ()^{pc} \\ \llbracket \ell \rrbracket^{pc} &= \ell^{pc} \\ \llbracket \mathbf{inl}(v) \rrbracket^{pc} &= \mathbf{inl}(\llbracket v \rrbracket^{pc})^{pc} \\ \llbracket \mathbf{inr}(v) \rrbracket^{pc} &= \mathbf{inr}(\llbracket v \rrbracket^{pc})^{pc} \\ \llbracket (v_1, v_2) \rrbracket^{pc} &= (\llbracket v_1 \rrbracket^{pc}, \llbracket v_2 \rrbracket^{pc})^{pc} \\ \llbracket (x.e, \theta) \rrbracket^{pc} &= (x.\llbracket e \rrbracket, \llbracket \theta \rrbracket^{pc})^{pc} \\ \llbracket (t, \theta) \rrbracket^{pc} &= (_.\llbracket t \rrbracket, \llbracket \theta \rrbracket^{pc})^{pc} \\ \llbracket \mathbf{Labeled} \ell v \rrbracket^{pc} &= (\ell^\ell, \llbracket v \rrbracket^\ell)^{pc} \\ \llbracket n_\ell \rrbracket^{pc} &= (n_\ell)^{pc} \end{aligned}$$

Fig. 15. Value translation from λ^{dCG} to λ^{dFG} .

$$\begin{aligned} \llbracket () \rrbracket &= () \\ \llbracket \ell \rrbracket &= \ell \\ \llbracket x \rrbracket &= x \\ \llbracket \lambda x.e \rrbracket &= \lambda x.\llbracket e \rrbracket \\ \llbracket e_1 e_2 \rrbracket &= \llbracket e_1 \rrbracket \llbracket e_2 \rrbracket \\ \llbracket (e_1, e_2) \rrbracket &= (\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket) \\ \llbracket \mathbf{fst}(e) \rrbracket &= \mathbf{fst}(\llbracket e \rrbracket) \\ \llbracket \mathbf{snd}(e) \rrbracket &= \mathbf{snd}(\llbracket e \rrbracket) \\ \llbracket \mathbf{inl}(e) \rrbracket &= \mathbf{inl}(\llbracket e \rrbracket) \\ \llbracket \mathbf{inr}(e) \rrbracket &= \mathbf{inr}(\llbracket e \rrbracket) \\ \llbracket \mathbf{case}(e, x.e_1, x.e_2) \rrbracket &= \mathbf{case}(\llbracket e \rrbracket, x.\llbracket e_1 \rrbracket, x.\llbracket e_2 \rrbracket) \\ \llbracket t \rrbracket &= \lambda_.\llbracket t \rrbracket \end{aligned}$$

Fig. 16. Expr. translation from λ^{dCG} to λ^{dFG} .

$$\begin{aligned} \llbracket \mathbf{return}(e) \rrbracket &= \llbracket e \rrbracket \\ \llbracket \mathbf{bind}(e_1, x.e_2) \rrbracket &= \\ &\quad \mathbf{let} \ x = \llbracket e_1 \rrbracket() \ \mathbf{in} \\ &\quad \quad \mathbf{taint}(\mathbf{labelOf}(x), \llbracket e_2 \rrbracket()) \\ \llbracket \mathbf{unlabel}(e) \rrbracket &= \\ &\quad \mathbf{let} \ x = \llbracket e \rrbracket \ \mathbf{in} \\ &\quad \quad \mathbf{taint}(\mathbf{fst}(x), \mathbf{snd}(x)) \\ \llbracket \mathbf{toLabeled}(e) \rrbracket &= \\ &\quad \mathbf{let} \ x = \llbracket e \rrbracket() \ \mathbf{in} \\ &\quad \quad (\mathbf{labelOf}(x), x) \\ \llbracket \mathbf{labelOf}(e) \rrbracket &= \mathbf{fst}(\llbracket e \rrbracket) \\ \llbracket \mathbf{getLabel} \rrbracket &= \mathbf{getLabel} \\ \llbracket \mathbf{taint}(e) \rrbracket &= \mathbf{taint}(\llbracket e \rrbracket, ()) \end{aligned}$$

Fig. 17. Thunk translation λ^{dCG} to λ^{dFG} .

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

Fig. 18. λ^{dCG} to λ^{dFG} translation of memory operations.

$\frac{\text{(VALUE)} \quad \ell_1 \sqsubseteq_{pc} r_1 \quad r_1 \downarrow_{\approx_{pc}} v_2}{r_1^{\ell_1} \downarrow_{\approx_{pc}} v_2}$	$\text{(UNIT)} \quad () \downarrow_{\approx_{pc}} ()$	$\text{(LABEL)} \quad \ell \downarrow_{\approx_{pc}} \ell$	$\text{(REF)} \quad n_\ell \downarrow_{\approx_{pc}} n_\ell$
$\frac{\text{(INL)} \quad v_1 \downarrow_{\approx_{pc}} v'_1}{\mathbf{inl}(v_1) \downarrow_{\approx_{pc}} \mathbf{inl}(v'_1)}$	$\frac{\text{(INR)} \quad v_2 \downarrow_{\approx_{pc}} v'_2}{\mathbf{inr}(v_2) \downarrow_{\approx_{pc}} \mathbf{inr}(v'_2)}$	$\frac{\text{(PAIR)} \quad v_1 \downarrow_{\approx_{pc}} v'_1 \quad v_2 \downarrow_{\approx_{pc}} v'_2}{(v_1, v_2) \downarrow_{\approx_{pc}} (v'_1, v'_2)}$	
$\frac{\text{(FUN)} \quad \theta_1 \downarrow_{\approx_{pc}} \theta_2}{(x. \llbracket e \rrbracket, \theta_1) \downarrow_{\approx_{pc}} (x.e, \theta_2)}$	$\frac{\text{(THUNK)} \quad \theta_1 \downarrow_{\approx_{pc}} \theta_2}{(-. \llbracket t \rrbracket, \theta_1) \downarrow_{\approx_{pc}} (t, \theta_2)}$	$\frac{\text{(LABELED)} \quad v_1 \downarrow_{\approx_{\ell}} v_2}{(\ell^\ell, v_1) \downarrow_{\approx_{pc}} (\mathbf{Labeled} \ell v_2)}$	

Fig. 19. Cross-language value equivalence modulo label annotations.