

Type Systems for Information Flow Control: The Question of Granularity



Vineet Rajani
MPI-SWS



Iulia Bastys
MPI-SWS



Willard Rafnsson
MPI-SWS



Deepak Garg
MPI-SWS

Information flow control is central to computer security. The objective of information flow control is to prevent unauthorized flows of secret information to the public outputs of a computation. This task is often accomplished using type systems that rely on modal operators to label and track information and, hence, this style of enforcing information flow control is deeply ingrained in logic. One key choice in designing a type system for information flow control, or dependence analysis in general, is the granularity at which dependencies are tracked. This article considers two extreme design points in this vast design space and examines their relative expressiveness.

1. INTRODUCTION

Information flow control (IFC) is a basic building block of computer security. IFC prevents the flow of high-confidentiality (or, simply, high) information to low-confidentiality (low) outputs that may be visible to attackers. For instance, one would not want private data stored on a file server to flow unencrypted to network packets since such packets can be read by all machines connected to the network, even those that are untrusted. Here, the private data is the high information and all unencrypted network packets are low outputs.

Ideally, IFC demands semantic independence of low outputs from high inputs. This is often called *noninterference* [Goguen and Meseguer 1982]: low outputs of a program should not be affected by changes to the program's high inputs. In practice, this ideal property is too restrictive but it is useful in designing enforcement techniques, which often start by aiming for noninterference, and then relax the property by allowing declassification in various ways [Sabelfeld and Sands 2009].

Although IFC can be enforced through several techniques—OS kernel mediation of process I/O [Elnikety et al. 2016; Krohn et al. 2007; Zeldovich et al. 2006], static analysis and type systems [Barthe et al. 2011; Myers 1999; Pottier and Simonet 2003; Buiras et al. 2015; Hunt and Sands 2006; Abadi et al. 1999], language runtime modification [Austin and Flanagan 2009; Hedin and Sabelfeld 2012; Rajani et al. 2015], the use of dedicated libraries [Li and Zdancewic 2006; Russo et al. 2008; Stefan et al. 2011], or compilation [Chudnov and Naumann 2015; Fournet et al. 2009]—our focus in this article is the enforcement of IFC in higher-order languages using type systems. Building on the seminal work of Volpano, Smith and Irvine [Volpano et al. 1996], which was not in a higher-order setting, many type systems have been proposed to enforce IFC in many different languages, including higher-order ones [Abadi et al. 1999; Pottier and Simonet 2003; Buiras et al. 2015].

The common denominator of all these type systems is type annotations or *labels* to mark program inputs, outputs and intermediate values as high or low, and a mechanism to track dependencies between program values, including inputs and outputs,

within the type system. However, there is significant variance in how the type systems track dependencies. Broadly speaking, dependencies may be tracked at *coarse-granularity* or at *fine-granularity*.

In coarse-grained dependence analysis, the type system forces any output temporally after the analysis (elimination) of a high-labeled value to be labeled high, since there could *potentially* be a dependence from the analyzed value to the output. Obviously, this introduces a coarse approximation, since not all outputs after the analysis of a high value may actually depend on the analyzed value. In information flow terminology, this unnecessary forcing of labels to high is called a *label creep*. To prevent label creep, the language may provide a scoping mechanism that syntactically delimits the effect of the analysis of a value. Despite the problem of label creep, the main advantage of coarse-grained dependence analysis is that it significantly reduces the need to label intermediate values since, by design, their labels are known implicitly from the labels of values analyzed in the past.

In contrast to coarse-grained dependence analysis, fine-grained analysis requires annotating (or inferring) the label of every intermediate value, and then carefully tracks dependencies among values. This makes the type system more precise but increases the annotation burden for either the programmer or a type-label inference engine.

The goal of this article is to provide an introduction to coarse- and fine-grained dependence analysis for IFC and to comment on their relative expressiveness. Specifically, we describe one type system each for coarse- and fine-grained dependence analysis. For coarse-grained dependence analysis, we choose a type system that tracks dependencies using a construct similar to an indexed family of monads. This type system is a simplification of an existing hybrid (mixed static and dynamic) system for dependence analysis called HLIO [Buiras et al. 2015]. We call this type system CG (for coarse-grained). For fine-grained dependence analysis, we choose a slight variant of Flow Caml [Pottier and Simonet 2003], an extension of ML's type system with information flow types. We call this type system FG (for fine-grained). In both cases, our setting is a simply-typed call-by-value lambda-calculus with references. To keep the presentation simple, we do not delve into concurrency or other evaluation strategies like call-by-name, which have nontrivial implications for dependence analysis and IFC.

Having presented the two type systems, we examine their relative expressiveness through translations. Specifically, we show that programs typable in CG can be translated in a type-preserving manner to FG. Although this may be unsurprising given the description of coarse- and fine-grained analysis above, the translation shows how the dependence analysis in CG can be simulated using specific monads in FG. We then attempt a translation from FG to CG, relying on a scope restriction construct in CG to prevent label creep. While we fail to do this (we explain why), we show that a fragment of FG can be translated, type-preserving, to CG.¹

It is not our goal to provide a comprehensive survey of all existing work on type systems for IFC. Indeed, this area is vast. Instead, we focus on one dimension of the design space—the granularity of the dependence analysis.

2. TYPE SYSTEMS FOR INFORMATION-FLOW CONTROL

We first present two state-of-the-art information-flow security type systems, FG and CG, for higher-order, stateful functional programming languages. The two type systems differ substantially in the approaches they follow to track dependencies. This is a consequence of how FG and CG differ computationally: FG allows (side-) effects in all expressions, à la ML. Since effects can occur so freely, information flows must be

¹Due to lack of space, we omit some details of the translations, which are provided in an accompanying technical report [Rajani et al. 2016].

tracked pervasively. Hence, FG is fine-grained. In contrast, CG isolates effects to a monad, à la Haskell. As a result, flows have to be tracked only at the granularity of the monad, but not within pure expressions. This makes CG coarse-grained.

Both FG and CG use *labels* drawn from a *lattice* $(\mathcal{L}, \sqsubseteq)$ of confidentiality levels l . Labels higher in the lattice represent higher confidentiality. The goal of dependence analysis for information flow is to ensure that terms labeled l can depend only on terms labeled l or lower. In examples, we often use the two-point lattice, $\text{LH} = (\{L, H\}, \sqsubseteq)$, which contains two levels L (low) and H (high) with $L \sqsubseteq H$ and $H \not\sqsubseteq L$. We use \perp and \top to denote the least and the greatest elements of any lattice. In LH, $\perp = L$ and $\top = H$.

2.1. Fine-grained type system

The fine-grained type system we consider, FG, is shown in Figure 1. FG is a slight modification of Flow Caml [Pottier and Simonet 2003], an extension of ML’s type system for information flow control. Computationally, FG is the call-by-value simply-typed lambda calculus, extended with products, sums, references, label polymorphism, and ordering constraints on labels.

Since side-effects may appear in any sub-expression in this language, FG must, when analyzing sub-expressions, account for all information that data concerning the sub-expression can contain. To this end, FG labels *all* of the (otherwise standard) types for this language with a *structural label* ℓ , reflecting an upper bound on the information conveyed by observing the structure of the expression. For instance, say `bool` is one of the base types that the symbol `b` in Figure 1 ranges over. Then observing a value of type `boolH` may reveal H information.

When analyzing non-ground expressions, FG tracks the propagation of information through the evaluation of expressions. For instance, FG concludes that the conjunction of a `boolH` and a `boolL` value is a `boolH` value, as observing the result may convey information about each component in the conjunction.

This tracking alone, however, is insufficient; since (sub)expressions can be evaluated conditionally, observing the presence or absence of effects can convey information about the *control-flow* conditions that facilitated or prevented the effects. Structural labels do not account for this information. For instance, let $x_c : (\text{unit}^L + \text{unit}^L)^H$, $x : (\text{ref nat}^L)^L$, and consider $e = \text{case}(x_c, _., _., x := 42)$.² The result of evaluating e is invariably $()$, so no information is conveyed by observing the result. However, on evaluation, e reveals whether $x_c = \text{inl}()$ or $x_c = \text{inr}()$ through the absence or presence of the write to x . FG tracks this information by recording control flow information in a *control label* pc (aka program counter), making it a lower bound on the write effects that the (sub)expression being typed can perform. For instance, when attempting to type the previous example, FG raises the pc by the information in the control-flow condition x , which is H , and checks whether the branches only have write effects at or above this new pc . However, the right branch writes 42 to x , which stores L -labeled natural numbers. So, with these labels on the types of x and x_c , e does not type-check.

Effects in a function’s body are suspended until the function is applied. Further, since our language is higher-order, a function can take another function as a parameter and apply it. This necessitates additional type annotations on function types. For instance, let $x_c : (\text{unit}^L + \text{unit}^L)^H$ and $x : (\text{ref nat}^L)^L$. Consider $e = \lambda x_f. \text{case}(x_c, _., _., (x_f _))$. Assuming that x_f maps unit^L to unit^L , e maps such mappings to unit^H , possibly applying x_f in the process. Now consider $e' = \lambda _.(x := 42)$, a function with a suspended effect, which maps unit^L to unit^L . While $e e'$ always returns a result of type unit^H , $e e'$ conditionally applies e' , and thus, the L effect in e' leaks the

²We use the symbol $_.$ to denote a variable, label or type whose actual value is irrelevant. Here, $_.$ denotes anonymous variables. Later, we use $_.$ to denote labels and types that are irrelevant to the discussion.

Syntax, types, constraints:

Expressions	e	$::=$	$x \mid \lambda x.e \mid e e \mid (e, e) \mid \text{fst}(e) \mid \text{snd}(e) \mid \text{inl}(e) \mid \text{inr}(e) \mid \text{case}(e, x.e, x.e) \mid \text{new } e \mid !e \mid e := e \mid () \mid \Lambda e \mid e [] \mid \nu e \mid e \bullet$
Labels	ℓ, pc	$::=$	$l \mid \alpha \mid \ell \sqcup \ell \mid \ell \sqcap \ell$
Types	τ	$::=$	A^ℓ
Base types	A	$::=$	$\mathbf{b} \mid \tau \xrightarrow{\ell} \tau \mid \tau \times \tau \mid \tau + \tau \mid \text{ref } \tau \mid \text{unit} \mid \forall \alpha. (\ell, \tau) \mid c \xrightarrow{\ell} \tau$
Constraints	c	$::=$	$\ell \sqsubseteq \ell \mid (c, c)$

Type system: $\boxed{\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau}$

$\frac{}{\Sigma; \Psi; \Gamma, x : \tau \vdash_{pc} x : \tau}$	FG-var	$\frac{\Sigma; \Psi; \Gamma, x : \tau_1 \vdash_{\ell_e} e : \tau_2}{\Sigma; \Psi; \Gamma \vdash_{pc} \lambda x.e : (\tau_1 \xrightarrow{\ell_e} \tau_2)^\perp}$	FG-lam
$\frac{\Sigma; \Psi; \Gamma \vdash_{pc} e_1 : (\tau_1 \xrightarrow{\ell} \tau_2)^\ell \quad \Sigma; \Psi; \Gamma \vdash_{pc} e_2 : \tau_1 \quad \Sigma; \Psi \vdash \tau_2 \searrow \ell \quad \Sigma; \Psi \vdash pc \sqcup \ell \sqsubseteq \ell_e}{\Sigma; \Psi; \Gamma \vdash_{pc} e_1 e_2 : \tau_2}$			FG-app
$\frac{\Sigma; \Psi; \Gamma \vdash_{pc} e_1 : \tau_1 \quad \Sigma; \Psi; \Gamma \vdash_{pc} e_2 : \tau_2}{\Sigma; \Psi; \Gamma \vdash_{pc} (e_1, e_2) : (\tau_1 \times \tau_2)^\perp}$			FG-prod
$\frac{\Sigma; \Psi; \Gamma \vdash_{pc} e : (\tau_1 \times \tau_2)^\ell \quad \Sigma; \Psi \vdash \tau_1 \searrow \ell}{\Sigma; \Psi; \Gamma \vdash_{pc} \text{fst}(e) : \tau_1}$	FG-fst	$\frac{\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau_1}{\Sigma; \Psi; \Gamma \vdash_{pc} \text{inl}(e) : (\tau_1 + \tau_2)^\perp}$	FG-inl
$\frac{\Sigma; \Psi; \Gamma, x : \tau_1 \vdash_{pc \sqcup \ell} e_1 : \tau \quad \Sigma; \Psi; \Gamma, y : \tau_2 \vdash_{pc \sqcup \ell} e_2 : \tau \quad \Sigma; \Psi \vdash \tau \searrow \ell}{\Sigma; \Psi; \Gamma \vdash_{pc} \text{case}(e, x.e_1, y.e_2) : \tau}$			FG-case
$\frac{\Sigma; \Psi; \Gamma \vdash_{pc'} e : \tau' \quad \Sigma; \Psi \vdash pc \sqsubseteq pc' \quad \Sigma; \Psi \vdash \tau' <: \tau}{\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau}$			FG-sub
$\frac{\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau \quad \Sigma; \Psi \vdash \tau \searrow pc}{\Sigma; \Psi; \Gamma \vdash_{pc} \text{new } e : (\text{ref } \tau)^\perp}$	FG-ref	$\frac{\Sigma; \Psi; \Gamma \vdash_{pc} e : (\text{ref } \tau)^\ell \quad \Sigma; \Psi \vdash \tau \searrow \ell}{\Sigma; \Psi; \Gamma \vdash_{pc} !e : \tau}$	FG-deref
$\frac{\Sigma; \Psi; \Gamma \vdash_{pc} e_1 : (\text{ref } \tau)^\ell \quad \Sigma; \Psi; \Gamma \vdash_{pc} e_2 : \tau \quad \Sigma; \Psi \vdash \tau \searrow (pc \sqcup \ell)}{\Sigma; \Psi; \Gamma \vdash_{pc} e_1 := e_2 : \text{unit}}$			FG-assign
$\frac{\Sigma, \alpha; \Psi; \Gamma \vdash_{\ell} e : \tau}{\Sigma; \Psi; \Gamma \vdash_{pc} \Lambda e : (\forall \alpha. (\ell, \tau))^\perp}$	FG-FI	$\frac{\Sigma; \Psi, c; \Gamma \vdash_{\ell} e : \tau}{\Sigma; \Psi; \Gamma \vdash_{pc} \nu e : (c \xrightarrow{\ell} \tau)^\perp}$	FG-CI
$\frac{\Sigma; \Psi; \Gamma \vdash_{pc} e : (\forall \alpha. (\ell, \tau))^\ell \quad \ell'' \in \Sigma \cup \mathcal{L} \quad \Sigma; \Psi \vdash pc \sqcup \ell' \sqsubseteq \ell \quad \Sigma; \Psi \vdash \tau \searrow \ell'}{\Sigma; \Psi; \Gamma \vdash_{pc} e [] : \tau[\ell''/\alpha]}$			FG-FE
$\frac{\Sigma; \Psi; \Gamma \vdash_{pc} e : (c \xrightarrow{\ell} \tau)^\ell \quad \Sigma; \Psi \vdash c \quad \Sigma; \Psi \vdash pc \sqcup \ell' \sqsubseteq \ell \quad \Sigma; \Psi \vdash \tau \searrow \ell'}{\Sigma; \Psi; \Gamma \vdash_{pc} e \bullet : \tau}$			FG-CE

Fig. 1. Syntax and type system of FG.

$$\begin{array}{c}
\frac{\Sigma; \Psi \vdash \ell \sqsubseteq \ell'}{\Sigma; \Psi \vdash \mathbf{b}^\ell <: \mathbf{b}^{\ell'}} \text{FGsub-base} \quad \frac{\Sigma; \Psi \vdash \ell \sqsubseteq \ell'}{\Sigma; \Psi \vdash (\text{ref } \tau)^\ell <: (\text{ref } \tau)^{\ell'}} \text{FGsub-ref} \\
\\
\frac{\Sigma; \Psi \vdash \tau_1 <: \tau'_1 \quad \Sigma; \Psi \vdash \tau_2 <: \tau'_2 \quad \Sigma; \Psi \vdash \ell \sqsubseteq \ell'}{\Sigma; \Psi \vdash (\tau_1 \times \tau_2)^\ell <: (\tau'_1 \times \tau'_2)^{\ell'}} \text{FGsub-prod} \\
\\
\frac{\Sigma; \Psi \vdash \tau_1 <: \tau'_1 \quad \Sigma; \Psi \vdash \tau_2 <: \tau'_2 \quad \Sigma; \Psi \vdash \ell \sqsubseteq \ell'}{\Sigma; \Psi \vdash (\tau_1 + \tau_2)^\ell <: (\tau'_1 + \tau'_2)^{\ell'}} \text{FGsub-sum} \\
\\
\frac{\Sigma; \Psi \vdash \tau'_1 <: \tau_1 \quad \Sigma; \Psi \vdash \tau_2 <: \tau'_2 \quad \Sigma; \Psi \vdash \ell \sqsubseteq \ell' \quad \Sigma; \Psi \vdash \ell'_e \sqsubseteq \ell_e}{\Sigma; \Psi \vdash (\tau_1 \xrightarrow{\ell'_e} \tau_2)^\ell <: (\tau'_1 \xrightarrow{\ell'_e} \tau'_2)^{\ell'}} \text{FGsub-arrow} \\
\\
\frac{\Sigma, \alpha; \Psi \vdash \tau_1 <: \tau_2 \quad \Sigma; \Psi \vdash \ell \sqsubseteq \ell' \quad \Sigma; \Psi \vdash \ell_2 \sqsubseteq \ell_1}{\Sigma; \Psi \vdash (\forall \alpha. (\ell_1, \tau_1))^\ell <: (\forall \alpha. (\ell_2, \tau_2))^{\ell'}} \text{FGsub-forall} \\
\\
\frac{\Sigma; \Psi \vdash c_2 \implies c_1 \quad \Sigma; \Psi \vdash \tau_1 <: \tau_2 \quad \Sigma; \Psi \vdash \ell \sqsubseteq \ell' \quad \Sigma; \Psi \vdash \ell_2 \sqsubseteq \ell_1}{\Sigma; \Psi \vdash (c_1 \xRightarrow{\ell_1} \tau_1)^\ell <: (c_2 \xRightarrow{\ell_2} \tau_2)^{\ell'}} \text{FGsub-constraint}
\end{array}$$

Fig. 2. FG subtyping.

control condition (x_c) in e , which is H . FG resolves this by having function types carry a separate control label; in $\tau \xrightarrow{\ell_e} \tau'$, ℓ_e is a lower bound on the level of the write effects that can occur when a function of this type is applied. In the example, $e' : (\text{unit}^L \xrightarrow{L} \text{unit}^L)^L$; thus FG rejects e' since e applies a function with L effects in a H context. Finally, note that functions of type $(\tau \xrightarrow{L} \tau')^H$ can be constructed but not applied in FG. This is because such a function can leak its identity, which is labeled H , to L when it is applied. However, if the function is merely passed around, it cannot leak information.

For the same reason, the types $\forall \alpha. (\ell_e, \tau)$ and $c \xRightarrow{\ell_e} \tau$ also carry the control label ℓ_e . In FG, values of these types (Λe and νe , respectively) are also suspended computations. However, the decision to suspend the computations inside these values is not fundamental since neither labels nor constraints have a runtime representation in FG.

FG performs security checks by checking the satisfiability of flow constraints, using the judgment $\Sigma, \Psi \vdash c$. A constraint c is a conjunction of terms of the form $\ell \sqsubseteq \ell'$, where ℓ ranges over levels, label-variables, and lattice operations on these. Let Ψ range over sets of constraints, and Σ range over sets of label parameters α . The judgment $\Sigma, \Psi \vdash c$ checks whether, for all instantiations of Σ , assuming Ψ , c holds. Label ℓ covers type $A^{\ell'}$ (from below), written $\Sigma, \Psi \vdash A^{\ell'} \searrow \ell$, iff $\Sigma, \Psi \vdash \ell \sqsubseteq \ell'$.

Subtyping. FG uses subtyping to allow upwards-flows of information. Subtyping amounts to weakening a guarantee for an expression. In our case, this guarantee is the type of an expression, which specifies how the information is classified. The subtyping judgment, defined in Figure 2, has the form $\Sigma; \Psi \vdash \tau <: \tau'$. In effect, this judgment extends (\sqsubseteq) to labeled expression types. For any A^ℓ , $<:$ is covariant in ℓ . This weakening of the type amounts to up-classifying information, which is safe since it only labels less confidential information as more confidential. Subtyping is covariant everywhere else, with two exceptions: control labels, and function arguments. A control label guarantees a lower bound on effects. This guarantee is weakened if the control label is lowered. For instance, if an expression has type $(\text{nat}^H \xrightarrow{H} \text{unit}^H)^L$, the function

may produce effects at or above H . This implies the weaker statement that the function may produce effects at or above L . Hence $(\text{nat}^H \xrightarrow{H} \text{unit}^H)^L <: (\text{nat}^H \xrightarrow{L} \text{unit}^H)^L$. A function argument appears as an assumption in the function type, and strengthening an assumption amounts to weakening the guarantee. For instance, if an expression has type $(\text{nat}^H \xrightarrow{H} \text{unit}^H)^L$, the function does not leak despite receiving H input. The function still will not leak if given L input. Hence, $(\text{nat}^H \xrightarrow{H} \text{unit}^H)^L <: (\text{nat}^L \xrightarrow{H} \text{unit}^H)^L$.

Typing judgment and typing rules. FG's type system prevents illicit flows of information by ensuring that

- eliminating an expression labeled ℓ produces a result covered by ℓ .
- an expression executing under pc can only cause write effects at or above pc .

The typing judgment has the form $\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau$. It reads: for all Σ , assuming Ψ and Γ , e has type τ , and pc is a lower bound on the level of all write effects which can occur when e is evaluated. We focus on three constructs, since these involve the pc : case, abstraction, and references.

In the rule FG-case, since case deconstructs its sum, the results of the branches must be covered by the label on the sum. Also, since either one or the other branch is evaluated depending on the sum, in typing the branches, the pc label is raised by the label on the sum, thus ensuring that the branches do not have write effects below the label of the sum.

In the rule FG-lam, FG can disregard the pc when typing the body of the function, because the body will not be evaluated immediately. FG thus only needs to check that the function satisfies what the type $(\tau_1 \xrightarrow{\ell_e} \tau_2)^\perp$ says it satisfies: (1) that the body has type τ_2 given input of type τ_1 , and (2) that all of its effects are at or above ℓ_e , which is ensured by checking the body of the function with pc set to ℓ_e . The outermost label on the conclusion's type $\tau_1 \xrightarrow{\ell_e} \tau_2$ is \perp because the fact that the function is constructed at this point in the program reveals no information. In fact, the outermost label is \perp in the introduction rules of all types, not just $\tau_1 \xrightarrow{\ell_e} \tau_2$. Rule FG-app checks that the result of applying a function is covered by the label on the function type, and that the effect of running the function does not leak contextual information, or structural information about the function.

In rules FG-ref and FG-assign, pc must cover the type of the value written to the reference. This ensures that write effects of the expression being typed are lower-bounded by pc . Additionally, in FG-assign, the label of the value written must cover the label on the reference to prevent leaking which reference was written. In the rule FG-deref, reading a reference conveys information about which reference was read; the result of the read must thus be covered by the label on the reference. (We implicitly assume that in the type ref τ , the type τ is closed, i.e., it has no free label parameters. Not enforcing this can break both subject reduction and the following noninterference property.)

Noninterference. FG enforces noninterference: The result of evaluating an expression of a labeled base type cannot depend on an input whose label does not cover the label of the base type.

Theorem 2.1. [Noninterference for FG] Suppose (1) $\ell_i \not\sqsubseteq \ell$, (2) $x : A^{\ell_i} \vdash_{pc} e : b^\ell$, and (3) $v_1, v_2 : A^{\ell_i}$. If both $e[v_1/x]$ and $e[v_2/x]$ terminate, then they produce the same value (of type b).

2.2. Coarse-grained type system

Next, we describe CG, a type system for coarse-grained dependence analysis. CG is not a new type system: It is the static fragment of HLIO [Buiras et al. 2015], a hybrid type system that mixes static and dynamic analyses to track flows. One minor

Syntax, types, constraints:

Expressions	$e ::= x \mid \lambda x.e \mid e e \mid (e, e) \mid \text{fst}(e) \mid \text{snd}(e) \mid \text{inl}(e) \mid \text{inr}(e) \mid \text{case}(e, x.e, y.e) \mid \text{new } e \mid !e \mid e := e \mid () \mid \Lambda e \mid e [] \mid \nu e \mid e \bullet \mid \text{label}_\ell(e) \mid \text{unlabel}(e) \mid \text{toLabeled}(e) \mid \text{ret}(e) \mid \text{bind}(e, x.e)$
Labels	$\ell ::= l \mid \alpha \mid \ell \sqcup \ell \mid \ell \sqcap \ell$
Types	$\tau ::= \mathbf{b} \mid \tau \rightarrow \tau \mid \tau \times \tau \mid \tau + \tau \mid \text{ref } \ell \tau \mid \text{unit} \mid \forall \alpha. \tau \mid c \Rightarrow \tau \mid \text{Labeled } \ell \tau \mid \text{CG } \ell_i \ell_o \tau$
Constraints	$c ::= \ell \sqsubseteq \ell \mid (c, c)$

Type system: $\boxed{\Sigma; \Psi; \Gamma \vdash e : \tau}$

(All rules of the simply typed lambda-calculus pertaining to the types $\mathbf{b}, \tau \rightarrow \tau, \tau \times \tau, \tau + \tau, \text{unit}$ are included.)

$$\begin{array}{c}
\frac{\Sigma; \Psi; \Gamma \vdash e : \tau \quad \Sigma; \Psi \vdash \ell_i \sqsubseteq \ell}{\Sigma; \Psi; \Gamma \vdash \text{label}_\ell(e) : \text{CG } \ell_i \ell_i \text{ (Labeled } \ell \tau)} \text{CG-label} \\
\\
\frac{\Sigma; \Psi; \Gamma \vdash e : \text{Labeled } \ell \tau}{\Sigma; \Psi; \Gamma \vdash \text{unlabel}(e) : \text{CG } \ell_i (\ell_i \sqcup \ell) \tau} \text{CG-unlabel} \\
\\
\frac{\Sigma; \Psi; \Gamma \vdash e : \text{CG } \ell_i \ell_o \tau}{\Sigma; \Psi; \Gamma \vdash \text{toLabeled}(e) : \text{CG } \ell_i \ell_i \text{ (Labeled } \ell_o \tau)} \text{CG-toLabeled} \\
\\
\frac{\Sigma; \Psi; \Gamma \vdash e : \tau}{\Sigma; \Psi; \Gamma \vdash \text{ret}(e) : \text{CG } \ell_i \ell_i \tau} \text{CG-ret} \\
\\
\frac{\Sigma; \Psi; \Gamma \vdash e_1 : \text{CG } \ell_i \ell \tau \quad \Sigma; \Psi; \Gamma, x : \tau \vdash e_2 : \text{CG } \ell \ell_o \tau'}{\Sigma; \Psi; \Gamma \vdash \text{bind}(e_1, x.e_2) : \text{CG } \ell_i \ell_o \tau'} \text{CG-bind} \\
\\
\frac{\Sigma; \Psi; \Gamma \vdash e : \tau' \quad \Sigma; \Psi \vdash \tau' <: \tau}{\Sigma; \Psi; \Gamma \vdash e : \tau} \text{CG-sub} \\
\\
\frac{\Sigma; \Psi; \Gamma \vdash e : \text{Labeled } \ell' \tau \quad \Sigma; \Psi \vdash \ell \sqsubseteq \ell'}{\Sigma; \Psi; \Gamma \vdash \text{new } e : \text{CG } \ell \ell \text{ (ref } \ell' \tau)} \text{CG-new} \\
\\
\frac{\Sigma; \Psi; \Gamma \vdash e : \text{ref } \ell \tau}{\Sigma; \Psi; \Gamma \vdash !e : \text{CG } \ell' \ell' \text{ (Labeled } \ell \tau)} \text{CG-deref} \\
\\
\frac{\Sigma; \Psi; \Gamma \vdash e_1 : \text{ref } \ell' \tau \quad \Sigma; \Psi; \Gamma \vdash e_2 : \text{Labeled } \ell' \tau \quad \Sigma; \Psi \vdash \ell \sqsubseteq \ell'}{\Sigma; \Psi; \Gamma \vdash e_1 := e_2 : \text{CG } \ell \ell \text{ unit}} \text{CG-assign} \\
\\
\frac{\Sigma, \alpha; \Psi; \Gamma \vdash e : \tau}{\Sigma; \Gamma \vdash \Lambda e : \forall \alpha. \tau} \text{CG-FI} \qquad \frac{\Sigma; \Psi; \Gamma \vdash e : \forall \alpha. \tau \quad \ell \in (\mathcal{L} \cup \Sigma)}{\Sigma; \Psi; \Gamma \vdash e [] : \tau[\ell/\alpha]} \text{CG-FE} \\
\\
\frac{\Sigma; \Psi, c; \Gamma \vdash e : \tau}{\Sigma; \Gamma \vdash \nu e : c \Rightarrow \tau} \text{CG-CI} \qquad \frac{\Sigma; \Psi; \Gamma \vdash e : c \Rightarrow \tau \quad \Sigma; \Psi \vdash c}{\Sigma; \Psi; \Gamma \vdash e \bullet : \tau} \text{CG-CE}
\end{array}$$

Fig. 3. Syntax and type system of CG.

difference from HLIO is that CG has call-by-value semantics to match those of FG whereas HLIO's semantics are call-by-name. This difference has little consequence for the discussion here.

CG is designed to minimize type-label annotations. To this end, it isolates all effects in a monad-like type construct. The syntax and typing rules of CG are shown in Figure 3. Unlike FG, standard typing constructs like products, arrows and sums are not refined with labels. These types behave exactly as in the simply typed lambda calculus (which CG extends conservatively) and the corresponding expressions do not have side-effects. For labeling, CG has a dedicated *type constructor* Labeled $\ell \tau$, which means τ labeled with ℓ . This is the only way to label a type in CG. Expressions are augmented with the constructs $\text{label}_\ell(e)$ and $\text{unlabel}(e)$ to introduce and eliminate Labeled $\ell \tau$.

Effects are limited to computations that have the type CG $\ell_i \ell_o \tau$. This type is similar to a monad and has the usual monadic return and bind constructs. Importantly, the bind construct is used to track dependencies coarsely. Finally, CG adds a scoping construct $\text{toLabeled}(e)$ that limits label creep. References in CG store only labeled values. A reference of type $\text{ref } \ell \tau$ stores values of type Labeled $\ell \tau$.

The type CG $\ell_i \ell_o \tau$. The type CG $\ell_i \ell_o \tau$ ascribes (suspended) computations that have effects. We define two kinds of effects in CG. *Input effects* cause a computation to learn new information and happen when a computation unlabels a labeled value. An *output effect* causes a computation to release information. This happens when a computation either creates a labeled value or writes to a reference. (Since references store only labeled values, merely reading a reference is not an input effect—to learn the actual content, the program must unlabel the value. Strictly speaking, it is also not essential to treat writing a reference as an output effect in CG. However, in many practical scenarios, attackers can observe writes to memory through side-channels outside the language, so we treat all writes as outputs.)

The type system enforces that the output effects of a computation of type CG $\ell_i \ell_o \tau$ are lower-bounded by ℓ_i and that its input effects are upper-bounded by ℓ_o . We call ℓ_i the “initial” program counter (pc) and ℓ_o the “final” pc for the computation. For instance, when writing to a reference, it is checked that the initial pc is below the label of the written value (last premise of rule CG-assign). A similar check is made when a labeled value is created (rule CG-label). When a value of type Labeled $\ell \tau$ is unlabeled, the final pc of the computation is joined with ℓ (rule CG-unlabel).

The construct $\text{bind}(e_1, x.e_2)$ allows sequencing two computations of types CG $\ell_i \ell \tau$ and $\tau \rightarrow \text{CG } \ell \ell_o \tau'$ to obtain a computation of type CG $\ell_i \ell_o \tau'$. Importantly, the final pc ℓ of the first computation must match the initial pc of the second computation. This ensures that the second computation's output effects (which are lower-bounded by ℓ) are at labels higher than the input effects of the first computation (which are upper-bounded by ℓ) and, hence, prevents any information leak. This is the only mechanism for tracking dependencies in CG.

It is an invariant of the type system that if $e : \text{CG } \ell_i \ell_o \tau$, then $\ell_i \sqsubseteq \ell_o$.

Construct toLabeled(e). As described above, sequencing a second computation after a computation of type CG $\ell_i \ell_o \tau$ using bind requires that the second computation's output effects be labeled higher than ℓ_o . This causes a label creep when the second computation does not actually examine the result of the first computation (e.g., the second computation may write the first computation's result to memory without examining it). To work around such a label creep, CG provides the expression construct toLabeled that coerces the type CG $\ell_i \ell_o \tau$ to CG $\ell_i \ell_i$ (Labeled $\ell_o \tau$). The computation returned by toLabeled , when forced, forces the original computation and labels the result

$$\begin{array}{c}
\frac{}{\Sigma; \Psi \vdash \tau <: \tau} \\
\frac{\Sigma; \Psi \vdash \tau_1 <: \tau_1 \quad \Sigma; \Psi \vdash \tau_2 <: \tau_2'}{\Sigma; \Psi \vdash \tau_1 \rightarrow \tau_2 <: \tau_1' \rightarrow \tau_2'} \\
\frac{\Sigma; \Psi \vdash \tau_1 <: \tau_1' \quad \Sigma; \Psi \vdash \tau_2 <: \tau_2'}{\Sigma; \Psi \vdash \tau_1 \times \tau_2 <: \tau_1' \times \tau_2'} \\
\frac{\Sigma; \Psi \vdash \tau_1 <: \tau_1' \quad \Sigma; \Psi \vdash \tau_2 <: \tau_2'}{\Sigma; \Psi \vdash \tau_1 + \tau_2 <: \tau_1' + \tau_2'} \\
\frac{\Sigma; \Psi \vdash \tau <: \tau' \quad \Sigma; \Psi \vdash \ell \sqsubseteq \ell'}{\Sigma; \Psi \vdash \text{Labeled } \ell \tau <: \text{Labeled } \ell' \tau'} \\
\frac{\Sigma; \Psi \vdash \tau <: \tau' \quad \Sigma; \Psi \vdash \ell'_i \sqsubseteq \ell_i \quad \Sigma; \Psi \vdash \ell_o \sqsubseteq \ell'_o}{\Sigma; \Psi \vdash \text{CG } \ell_i \ell_o \tau <: \text{CG } \ell'_i \ell'_o \tau'} \\
\frac{\Sigma, \alpha; \Psi \vdash \tau_1 <: \tau_2}{\Sigma; \Psi \vdash \forall \alpha. \tau_1 <: \forall \alpha. \tau_2} \\
\frac{\Sigma; \Psi \vdash c_2 \implies c_1 \quad \tau_1 <: \tau_2}{\Sigma; \Psi \vdash c_1 \implies \tau_1 <: c_2 \implies \tau_2}
\end{array}$$

Fig. 4. CG subtyping.

with ℓ_o .³ A computation of the type $\text{CG } \ell_i \ell_i$ (Labeled $\ell_o \tau$) can be followed by a second computation whose output effects are at level ℓ_i or higher. The pc increases to ℓ_o only if the second computation actually unlabels the result of the first computation.

Subtyping. CG includes the usual subtyping rules of the simply typed lambda calculus. Subtyping for Labeled $\ell \tau$ is covariant in ℓ . Subtyping for $\text{CG } \ell_i \ell_o \tau$ is contravariant in ℓ_i and covariant in ℓ_o . This is natural since ℓ_i is a lower-bound (on the output effects) and ℓ_o is an upper-bound (on the input effects). The subtyping rules of CG are shown in Figure 4.

Noninterference. CG satisfies noninterference: If a computation has only low input effects and returns a value of base type, then the returned value must be independent of any high input.

Theorem 2.2. [Noninterference for CG] Suppose (1) $\ell_i \not\sqsubseteq \ell$, (2) $x : \text{Labeled } \ell_i \tau \vdash e : \text{CG } _ \ell b$, and (3) $v_1, v_2 : \text{Labeled } \ell_i \tau$. If both $e[v_1/x]$ and $e[v_2/x]$ terminate when forced, then they produce the same value (of type b).

3. TRANSLATIONS

Having described the fine- and coarse-grained dependence analysis type systems FG and CG, we now turn to understanding their relative expressiveness. We do so by presenting (attempted) type-preserving translations from CG to FG, and vice-versa. We start by showing a type-preserving translation from CG to FG in Section 3.1. We then attempt a translation in the reverse direction, show where it fails and why (Section 3.2). Based on our attempt, we identify a smaller fragment of FG which can be translated to CG, preserving types.

3.1. Translating CG to FG

In this section, we define a translation $\llbracket \cdot \rrbracket$ from CG to FG and show that it is type-preserving. The translation of types is shown below.

³The term “forcing” is used here in the sense of monads. Forcing a value of type $\text{CG } \ell_i \ell_o \tau$ runs the suspended computation, records its write effects and eventually returns whatever the computation returns.

$$\begin{array}{ll}
\llbracket \mathbf{b} \rrbracket = \mathbf{b}^\perp & \llbracket \text{unit} \rrbracket = \text{unit}^\perp \\
\llbracket \tau_1 \rightarrow \tau_2 \rrbracket = (\llbracket \tau_1 \rrbracket \xrightarrow{\top} \llbracket \tau_2 \rrbracket)^\perp & \llbracket \text{ref } \ell \tau \rrbracket = (\text{ref } (\llbracket \tau \rrbracket + \text{unit})^\ell)^\perp \\
\llbracket \tau_1 \times \tau_2 \rrbracket = (\llbracket \tau_1 \rrbracket \times \llbracket \tau_2 \rrbracket)^\perp & \llbracket \text{CG } \ell_i \ell_o \tau \rrbracket = (\text{unit} \xrightarrow{\ell_i} (\llbracket \tau \rrbracket + \text{unit})^{\ell_o})^\perp \\
\llbracket \tau_1 + \tau_2 \rrbracket = (\llbracket \tau_1 \rrbracket + \llbracket \tau_2 \rrbracket)^\perp & \llbracket c \Rightarrow \tau \rrbracket = (c \xrightarrow{\top} \llbracket \tau \rrbracket)^\perp \\
\llbracket \text{Labeled } \ell \tau \rrbracket = (\llbracket \tau \rrbracket + \text{unit})^\ell & \llbracket \forall \alpha. \tau \rrbracket = (\forall \alpha. (\top, \llbracket \tau \rrbracket))^\perp
\end{array}$$

This translation relies on three key ideas. First, in CG, labels are limited to the type construct Labeled $\ell \tau$, so the translation of all other types can simply use the outer label \perp . There are several choices for translating Labeled $\ell \tau$. A natural translation would be $A^{\ell' \sqcup \ell}$, where $A^{\ell'}$ is the translation of τ . However, this translation “flattens” nested labels of the form Labeled ℓ (Labeled $\ell' \tau$), making it impossible to simulate, in the translation, the selective unlabeled of only the outer ℓ , but not the inner ℓ' , which is allowed in CG. To keep the labels ℓ and ℓ' separate in the translation, we translate Labeled $\ell \tau$ to $(\llbracket \tau \rrbracket + \text{unit})^\ell$, which keeps the label on $\llbracket \tau \rrbracket$ separate from ℓ . The corresponding translation of expressions uses `inl`, thus never actually returning the unit value during execution.

Second, in CG, side-effects are confined to the type CG $\ell_i \ell_o \tau$, so when translating CG’s remaining types, which represent pure terms, we can always use $pc = \top$ in FG (since there are no side-effects in the pure terms, \top is trivially the strictest lower-bound on the output effects). As a result, the control labels on \rightarrow , \Rightarrow and \forall in the translations of $\tau_1 \rightarrow \tau_2$, $c \Rightarrow \tau$ and $\forall \alpha. \tau$ are all \top .

The type CG $\ell_i \ell_o \tau$ represents a suspended computation whose effects are visible only after it is forced. This is emulated in FG using a thunk, a function that takes an argument of unit type. Specifically, CG $\ell_i \ell_o \tau$ translates to $(\text{unit} \xrightarrow{\ell_i} (\llbracket \tau \rrbracket + \text{unit})^{\ell_o})^\perp$, which is a decorated variant of the thunk type $\text{unit} \rightarrow \llbracket \tau \rrbracket$. The thunk can be forced when needed by applying it to $()$. The ℓ_i on the arrow means (in FG) that the write-effects of the computation (the thunk) are lower-bounded by ℓ_i , which is exactly the meaning of ℓ_i in CG $\ell_i \ell_o \tau$. The label ℓ_o on $(\llbracket \tau \rrbracket + \text{unit})$ implies that the result of the computation cannot be analyzed without raising the `pc` to ℓ_o in FG, which is exactly the consequence of having ℓ_o in the type CG $\ell_i \ell_o \tau$ in CG. (We note that the translation simulates CG $\ell_i \ell_o \tau$ using a combination of the type forms $\text{unit} \rightarrow \llbracket \tau \rrbracket$ and $\llbracket \tau \rrbracket + \text{unit}$, both of which are monads.)

Finally, in CG, a reference of type `ref $\ell \tau$` stores values of type Labeled $\ell \tau$. Hence, the translation of `ref $\ell \tau$` is $(\text{ref } (\llbracket \tau \rrbracket + \text{unit})^\ell)^\perp$.

The translation $\llbracket \cdot \rrbracket$ is lifted pointwise to contexts: $\llbracket \Gamma \rrbracket \triangleq \{x : \llbracket \tau \rrbracket \mid x : \tau \in \Gamma\}$. The translation of expressions is defined by induction on CG typing derivations. We write $\Sigma; \Psi; \Gamma \vdash e : \tau \rightsquigarrow \Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} e' : \llbracket \tau \rrbracket$ to mean that the well-typed CG expression e translates to the well-typed FG expression e' . Selected rules of the translation are shown in Figure 5. They should be unsurprising given the type translation.

The following theorem shows that this translation preserves types, in the sense that \rightsquigarrow always maps a valid CG typing derivation to a valid FG typing derivation.

Theorem 3.1 (Soundness, CG \rightsquigarrow FG). *If $\Sigma; \Psi; \Gamma \vdash e : \tau$ has a valid CG typing derivation, then there exists an e' such that $\Sigma; \Psi; \Gamma \vdash e : \tau \rightsquigarrow \Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} e' : \llbracket \tau \rrbracket$ and $\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} e' : \llbracket \tau \rrbracket$ has a valid FG typing derivation.*

3.2. Translating FG to CG

Next, we consider translating FG to CG. We start with an incorrect strawman translation, which we refine, eventually getting to a point where no further progress seems possible. At that point, we identify a fragment of FG for which the refined translation

$$\begin{array}{c}
\frac{\Sigma; \Psi; \Gamma \vdash e : \tau \quad \Sigma; \Psi \vdash \ell_i \sqsubseteq \ell}{\Sigma; \Psi; \Gamma \vdash \text{label}_\ell(e) : \text{CG } \ell_i \ell_i \text{ (Labeled } \ell \tau)} \rightsquigarrow \\
\frac{\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} e' : \llbracket \tau \rrbracket}{\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} \lambda_{\cdot}.\text{inl}(\text{inl}(e')) : (\text{unit} \xrightarrow{\ell_i} (\llbracket \tau \rrbracket + \text{unit})^\ell + \text{unit})^{\ell_i} \perp} \\
\hline
\frac{\Sigma; \Psi; \Gamma \vdash e : \text{Labeled } \ell \tau}{\Sigma; \Psi; \Gamma \vdash \text{unlabel}(e) : \text{CG } \ell_i (\ell_i \sqcup \ell) \tau} \rightsquigarrow \frac{\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} e' : (\llbracket \tau \rrbracket + \text{unit})^\ell}{\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} \lambda_{\cdot}.e' : (\text{unit} \xrightarrow{\ell_i} (\llbracket \tau \rrbracket + \text{unit})^{\ell_i \sqcup \ell}) \perp} \\
\hline
\frac{\Sigma; \Psi; \Gamma \vdash e : \text{CG } \ell_i \ell_o \tau}{\Sigma; \Psi; \Gamma \vdash \text{toLabeled}(e) : \text{CG } \ell_i \ell_i \text{ (Labeled } \ell_o \tau)} \rightsquigarrow \\
\frac{\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} e' : (\text{unit} \xrightarrow{\ell_i} (\llbracket \tau \rrbracket + \text{unit})^{\ell_o}) \perp}{\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} \lambda_{\cdot}.\text{inl}(e'()) : (\text{unit} \xrightarrow{\ell_i} (\llbracket \tau \rrbracket + \text{unit})^{\ell_o} + \text{unit})^{\ell_i} \perp} \\
\hline
\frac{\Sigma; \Psi; \Gamma \vdash e : \tau}{\Sigma; \Psi; \Gamma \vdash \text{ret}(e) : \text{CG } \ell_i \ell_i \tau} \rightsquigarrow \frac{\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} e' : \llbracket \tau \rrbracket}{\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} \lambda_{\cdot}.\text{inl}(e') : (\text{unit} \xrightarrow{\ell_i} (\llbracket \tau \rrbracket + \text{unit})^{\ell_i}) \perp} \\
\hline
\frac{\Sigma; \Psi; \Gamma \vdash e_1 : \text{CG } \ell_i \ell \tau \quad \Sigma; \Psi; \Gamma, x : \tau \vdash e_2 : \text{CG } \ell \ell_o \tau'}{\Sigma; \Psi; \Gamma \vdash \text{bind}(e_1, x.e_2) : \text{CG } \ell_i \ell_o \tau'} \rightsquigarrow \\
\frac{\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} e'_1 : (\text{unit} \xrightarrow{\ell_i} (\llbracket \tau \rrbracket + \text{unit})^\ell) \perp \quad \Sigma; \Psi; \llbracket \Gamma \rrbracket, x : \llbracket \tau \rrbracket \vdash_{\top} e'_2 : (\text{unit} \xrightarrow{\ell} (\llbracket \tau' \rrbracket + \text{unit})^{\ell_o}) \perp}{\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash_{\top} \lambda_{\cdot}.\text{case}(e'_1(), x.e'_2(), y.\text{inr}()) : (\text{unit} \xrightarrow{\ell_i} (\llbracket \tau' \rrbracket + \text{unit})^{\ell_o}) \perp}
\end{array}$$

Fig. 5. Type derivation-directed expression translation from CG into FG, selected rules.

works. The goal of going through this exercise is to impress upon the reader the difficulty of translating a fine-grained dependence analysis to a coarse-grained one, and to argue that there does not seem to be a straightforward translation from all of FG to CG, despite CG having the construct `toLabeled` to prevent label creep.

Strawman translation. We construct a strawman translation, $\llbracket \cdot \rrbracket$, from FG to CG that we soon show to be incorrect. We translate the type A^ℓ to `Labeled` ℓ $\llbracket A \rrbracket$ since this is the only type construct that adds a label in CG.

Next, consider the function type $\tau_1 \xrightarrow{\ell_e} \tau_2$ in FG. Since the body of a function of this type can have a write effect at level ℓ_e or higher, an intuitive translation of this type could have the form $\llbracket \tau_1 \rrbracket \rightarrow \text{CG } \ell_e \ell_o \llbracket \tau_2 \rrbracket$. For the translation of the function's body to be well-typed in CG, the label ℓ_o must be an upper-bound on the labels of everything the function's body analyzes. Nothing in the FG type specifies this upper-bound, so we must find some other alternative. Fortunately, it is possible to *confine* the effects of

value analysis using the construct `toLabeled` in CG. As a result, we may hope that we can choose $\ell_o = \ell_e$ and translate $\tau_1 \xrightarrow{\ell_e} \tau_2$ to $\llbracket \tau_1 \rrbracket \rightarrow \text{CG } \ell_e \ell_e \llbracket \tau_2 \rrbracket$.

Independent of what ℓ_o we choose, this translation has a label creep problem. Consider a FG function f of type $\text{unit} \xrightarrow{H} A^L$ in the lattice LH. This function may write high values to references but it eventually returns a low value. In FG, the *result* of f 's application can be written to a reference of type $\text{ref } A^L$. However, after translation, this write would be impossible because f 's type would translate to $\llbracket \text{unit} \rrbracket \rightarrow \text{CG } H H$ (Labeled $L A$). Applying this type would result in a computation, say c , of type $\text{CG } H H$ (Labeled $L A$). There is no way to extract a low labeled value from this computation. At best, we may use subtyping, `bind` and `toLabeled` as in `toLabeled(bind(c, x.unlabel(x)))` to coerce the type to $\text{CG } L L$ (Labeled $H A$), but the resulting value still has the label H .

Based on this, we may be tempted to translate $\tau_1 \xrightarrow{\ell_e} \tau_2$ to $\llbracket \tau_1 \rrbracket \rightarrow \text{CG } \perp \perp \llbracket \tau_2 \rrbracket$ instead (this is sound because \perp is trivially a lower bound on any write effect in the function's body). Although this translation would solve the label creep problem mentioned in the previous paragraph, it suffers from a different problem: Now, the translation cannot simulate an application of the previous paragraph's function f in a high context, i.e., in a case branch where the analyzed sum is labeled H . To see this, consider the FG expression `case(h, x.f(), ...)`, where $h : (\tau + \tau')^H$. In FG, the type of this expression is A^H . In CG, we would correspondingly like to construct a result of type Labeled $H \llbracket A \rrbracket$. However, this is impossible. Since h 's translation has type Labeled $H (\llbracket \tau \rrbracket + \llbracket \tau' \rrbracket)$, to perform a case analysis on it, we must first unlabel it. This will result in a computation of type $\text{CG } L H (\llbracket \tau \rrbracket + \llbracket \tau' \rrbracket)$. Next, we can bind this computation and case analyze the value of type $\llbracket \tau \rrbracket + \llbracket \tau' \rrbracket$. However, due to the restrictions in typing `bind`, any further binds we perform must be on values of type $\text{CG } H H _$. The body of f 's translation has the type $\text{CG } L L$ (Labeled $L \llbracket A \rrbracket$) ($L = \perp$ here) and there is no way to coerce this to the form $\text{CG } H H _$ because subtyping for $\text{CG } \ell_i \ell_e \tau$ is contravariant in ℓ_i . So, we cannot bind the body of f , and, hence, cannot obtain a value of type Labeled $_ \llbracket A \rrbracket$.

Using label polymorphism. The problems with the strawman translation above can be addressed using label polymorphism. For instance, we could translate $\tau_1 \xrightarrow{\ell_e} \tau_2$ to $\llbracket \tau_1 \rrbracket \rightarrow \forall \alpha. \text{CG } \alpha \alpha \llbracket \tau_2 \rrbracket$. This would allow us to use the earlier function f in both contexts, instantiating α with L in the first context and with H in the second context. However, this translation is unsound. Specifically, instantiating α with some $\ell'_e \not\sqsubseteq \ell_e$ allows us to establish that every write in the function's body is at the level ℓ'_e or higher, which is clearly false, since the function's body may write at level ℓ_e (according to the FG type $\tau_1 \xrightarrow{\ell_e} \tau_2$).

Consequently, we consider a revised translation that maps $\tau_1 \xrightarrow{\ell_e} \tau_2$ to $\llbracket \tau_1 \rrbracket \rightarrow \forall \alpha. (\alpha \sqsubseteq \ell_e) \Rightarrow \text{CG } \alpha \alpha \llbracket \tau_2 \rrbracket$. The entire type translation is shown below. (The translation of $c \xrightarrow{\ell_e} \tau$ and $\forall \alpha. (\ell_e, \tau)$ follows the same intuition as the translation of $\tau_1 \xrightarrow{\ell_e} \tau_2$.)

$$\begin{array}{ll} \llbracket b \rrbracket = b & \llbracket \text{unit} \rrbracket = \text{unit} \\ \llbracket \tau_1 \xrightarrow{\ell_e} \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \rightarrow \forall \alpha. (\alpha \sqsubseteq \ell_e) \Rightarrow \text{CG } \alpha \alpha \llbracket \tau_2 \rrbracket & \llbracket \tau_1 \times \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \times \llbracket \tau_2 \rrbracket \\ \llbracket c \xrightarrow{\ell_e} \tau \rrbracket = \forall \alpha. (\alpha \sqsubseteq \ell_e, c) \Rightarrow \text{CG } \alpha \alpha \llbracket \tau \rrbracket & \llbracket \tau_1 + \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket + \llbracket \tau_2 \rrbracket \\ \llbracket \forall \alpha. (\ell_e, \tau) \rrbracket = \forall \alpha. \forall \alpha'. (\alpha' \sqsubseteq \ell_e) \Rightarrow \text{CG } \alpha' \alpha' \llbracket \tau \rrbracket & \llbracket \text{ref } A^\ell \rrbracket = \text{ref } \ell \llbracket A \rrbracket \\ \llbracket A^\ell \rrbracket = \text{Labeled } \ell \llbracket A \rrbracket & \end{array}$$

The translation of contexts Γ is defined pointwise and a FG typing judgment $\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau$ translates to a CG judgment of the form $\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash e' : \forall \alpha. (\alpha \sqsubseteq pc) \Rightarrow \text{CG } \alpha \alpha \llbracket \tau \rrbracket$, mirroring the label polymorphism in the bodies of function types (e' is the translation of e).

Unfortunately, this translation has a different problem! Consider how we would (inductively) translate the rule FG-case from Figure 1. Inductively, from the premises we obtain e' , e'_1 and e'_2 (the translations of e , e_1 and e_2 , respectively) such that:

- (1) $\Sigma; \Psi; [\Gamma] \vdash e' : \forall \alpha. (\alpha \sqsubseteq pc) \Rightarrow \text{CG } \alpha \ \alpha \text{ (Labeled } \ell \text{ } ([\tau_1] + [\tau_2]))$
- (2) $\Sigma; \Psi; [\Gamma], x : [\tau_1] \vdash e'_1 : \forall \alpha_1. (\alpha_1 \sqsubseteq (pc \sqcup \ell)) \Rightarrow \text{CG } \alpha_1 \ \alpha_1 \text{ } [\tau]$
- (3) $\Sigma; \Psi; [\Gamma], y : [\tau_2] \vdash e'_2 : \forall \alpha_2. (\alpha_2 \sqsubseteq (pc \sqcup \ell)) \Rightarrow \text{CG } \alpha_2 \ \alpha_2 \text{ } [\tau]$

The goal is to construct a term e'' (the translation of $\text{case}(e, x.e_1, y.e_2)$) such that

$$\Sigma; \Psi; [\Gamma] \vdash e'' : \forall \alpha'. (\alpha' \sqsubseteq pc) \Rightarrow \text{CG } \alpha' \ \alpha' \text{ } [\tau]$$

We try to search for the appropriate term e'' (much as we would look for a proof in a formal proof system). We pick some α' such that $\alpha' \sqsubseteq pc$. We must construct a term of the type $\text{CG } \alpha' \ \alpha' \text{ } [\tau]$. Our only option is to case analyze the value of type $([\tau_1] + [\tau_2])$ in (1), so we must instantiate the quantified α in (1) and bind the resulting computation type. Since the eventual goal is to obtain something of type $\text{CG } \alpha' \ _ _$, we must pick $\alpha = \alpha'$. We instantiate $\alpha = \alpha'$, and bind the computation of type $\text{CG } \alpha' \ \alpha' \text{ (Labeled } \ell \text{ } ([\tau_1] + [\tau_2]))$ in (1), obtaining a local variable of type $\text{Labeled } \ell \text{ } ([\tau_1] + [\tau_2])$. We unlabel this to obtain a computation of type $\text{CG } \alpha' \ (\alpha' \sqcup \ell) \text{ } ([\tau_1] + [\tau_2])$, which we bind again to obtain a variable of type $[\tau_1] + [\tau_2]$. This variable can be case-analyzed. To construct the case branches we must instantiate and bind the computations in (2) and (3). We show only the operations on (2), those on (3) being similar. First, we must pick a suitable α_1 . Since the next computation we construct must have a type of the form $\text{CG } (\alpha' \sqcup \ell) \ _ _$, we must pick $\alpha_1 = \alpha' \sqcup \ell$ (which is indeed below $(pc \sqcup \ell)$, as required by the constraint in (2)). Second, we instantiate (2) with this substitution to obtain a computation of type $\text{CG } (\alpha' \sqcup \ell) \ (\alpha' \sqcup \ell) \text{ } [\tau]$. Repeating this process on (3), we obtain an end-to-end computation of type $\text{CG } \alpha' \ (\alpha' \sqcup \ell) \text{ } [\tau]$.

This is *almost* what we wanted. To complete the proof, we have to coerce the type $\text{CG } \alpha' \ (\alpha' \sqcup \ell) \text{ } [\tau]$ to the type $\text{CG } \alpha' \ \alpha' \text{ } [\tau]$. For this, we consider the cases $\alpha' \sqsubseteq \ell$ and $\alpha' \not\sqsubseteq \ell$ separately. Strictly speaking, CG does not allow a case analysis on constraints. However, we show below that the proof cannot even be completed in the second case, so the case analysis has expository value.

When $\alpha' \sqsubseteq \ell$, then $\text{CG } \alpha' \ (\alpha' \sqcup \ell) \text{ } [\tau] = \text{CG } \alpha' \ \ell \text{ } [\tau]$ and it is not difficult to write a coercion function from $\text{CG } \alpha' \ \ell \text{ } [\tau]$ to $\text{CG } \alpha' \ \alpha' \text{ } [\tau]$. The fourth premise of the FG-case rule is $\tau \searrow \ell$, so $\tau = A^{\ell'}$ for some $\ell' \sqsupseteq \ell$ and $[\tau] = \text{Labeled } \ell' \text{ } [A]$. The required coercion function is $\lambda x : (\text{CG } \alpha' \ \ell \text{ } [\tau]). \text{toLabeled}(\text{bind}(x, y.\text{unlabel}(y)))$.

However, in the case $\alpha' \not\sqsubseteq \ell$, such a coercion function may not exist. Concretely, consider the lattice $L \sqsubseteq \{M_1, M_2\} \sqsubseteq H$ with M_1, M_2 incomparable, $\alpha' = M_1$, $\ell = M_2$ and $\tau = A^{M_2}$. In this case, our goal is to coerce $\text{CG } M_1 \ H \text{ (Labeled } M_2 \text{ } [A])$ to $\text{CG } M_1 \ M_1 \text{ (Labeled } M_2 \text{ } [A])$. This is impossible in CG: Our only hope of getting rid of the H in the given type is to use toLabeled , but that would push the H into the label of the resulting value.

It follows, therefore, that even our revised translation does not work. However, on any fragment of FG where the second case $\alpha' \not\sqsubseteq \ell$ can never arise, this translation *would* work. In the following, we identify such a fragment, FG^- .

The fragment FG^- . Because α' is arbitrary and the only constraint on it is $\alpha' \sqsubseteq pc$, disallowing $\alpha' \not\sqsubseteq \ell$ is the same as always forcing $pc \sqsubseteq \ell$. One simple way of ensuring $pc \sqsubseteq \ell$ is to restrict FG to a fragment in which $\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau$ implies $\tau \searrow pc$. Then, (1) would force $pc \sqsubseteq \ell$. Defining such a fragment is straightforward. We only need to restrict the types in the conclusions of the typing rules for all introduction forms like pairing, functions, inl , inr , etc. to be labeled pc (currently, these rules allow the label \perp). Elimination rules do not require any changes (although some premises in the

elimination rules become redundant, e.g., the premise $\tau \searrow \ell$ in the rule FG-case). We can then show inductively that $\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau$ implies $\tau \searrow pc$.

For instance, the rules FG-var and FG-lam of Figure 1 are replaced with the following more restrictive rules.

$$\frac{\Sigma; \Psi \vdash \tau \sqsubseteq \tau' \quad \tau' \searrow pc}{\Sigma; \Psi; \Gamma, x : \tau \vdash_{pc} x : \tau'} \text{R-var} \qquad \frac{\Sigma; \Psi; \Gamma, x : \tau_1 \vdash_{\ell_e} e : \tau_2}{\Sigma; \Psi; \Gamma \vdash_{pc} \lambda x. e : (\tau_1 \xrightarrow{\ell_e} \tau_2)^{pc}} \text{R-lam}$$

Lemma 3.2. $\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau$ in FG^- implies $\Sigma; \Psi \vdash \tau \searrow pc$.

We can prove that on the fragment FG^- , the translation $\llbracket \cdot \rrbracket$ defined above is total and type-preserving. We have to first define a type derivation-directed translation of expressions, whose straightforward details we elide here (the details can be found in the accompanying technical report). This translation is written $\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau \rightsquigarrow \Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash e' : \forall \alpha. (\alpha \sqsubseteq pc) \Rightarrow CG \alpha \alpha \llbracket \tau \rrbracket$.

Theorem 3.3 (Soundness, $FG^- \rightsquigarrow CG$). *If $\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau$ has a valid FG^- typing derivation, then there exists an e' such that $\Sigma; \Psi; \Gamma \vdash_{pc} e : \tau \rightsquigarrow \Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash e' : \forall \alpha. (\alpha \sqsubseteq pc) \Rightarrow CG \alpha \alpha \llbracket \tau \rrbracket$ and $\Sigma; \Psi; \llbracket \Gamma \rrbracket \vdash e' : \forall \alpha. (\alpha \sqsubseteq pc) \Rightarrow CG \alpha \alpha \llbracket \tau \rrbracket$ has a valid CG typing derivation.*

4. OTHER TYPE SYSTEMS

Several other type systems for information flow control can be classified as either fine-grained [Pottier and Simonet 2003; Volpano et al. 1996; Heintze and Riecke 1998] or coarse-grained [Matos 2006; Russo 2015; Buiras et al. 2015]. Of particular note is the dependency core calculus (DCC) [Abadi et al. 1999]. DCC uses a monad to track dependencies, in a manner similar to CG, but is otherwise pure. [Abadi et al. 1999] show how several calculi for dependence analysis can be translated to DCC. One of these calculi is a first-order calculus with references [Smith and Volpano 1998]. This calculus has a rule very similar to the case analysis rule of FG, whose translation failed in Section 3.2. A priori, it seems that we ought to be able to examine the translation from [Smith and Volpano 1998] to DCC to understand how to translate FG's case analysis rule to CG. However, [Abadi et al. 1999]'s translation is not parametric in the security lattice: It is defined only for the lattice LH, and treats the (analogues of the) FG judgments $\Sigma; \Psi; \Gamma \vdash_L e : \tau$ and $\Sigma; \Psi; \Gamma \vdash_H e : \tau$ completely differently. Indeed, we expect that such a non-parametric translation would also exist from FG to CG, at least for the lattice LH.

5. CONCLUSION

At their core, type systems for information flow control perform dependence analysis. Moving from a fine-grained to a coarse-grained dependence analysis trades off precision for fewer type-label annotations. In this article, we have initiated a study of the relative expressiveness of these two approaches by considering type-preserving translations from a coarse-grained type system to a fine-grained type system and vice-versa. Our analysis indicates that the former is straightforward (as expected) whereas the latter is not.

In ongoing work, we are examining two problems that we have not yet addressed satisfactorily. First, we would like to prove that the translations are operationally sound (not just type-preserving). Ideally, we would like to derive the noninterference theorem for one system from the noninterference theorem of the other system and properties of the translation. Prior work has established similar results for other translations.

For example, [Abadi et al. 1999] establish similar results for the translation of several dependency-tracking calculi into DCC. In our setting, the problem is harder due to the presence of state, whose combination with higher-order functions would complicate any model of types. Second, we would like to find a translation from all of FG to CG or show that such a translation does not exist. Since Section 3.2 already shows a translation from FG^- to CG, the problem of translating FG to CG simplifies to that of finding a translation from FG to FG^- .

Acknowledgments. We would like to thank Alejandro Russo for discussions on coarse-grained dependence analysis and for feedback on a draft of this article. This work was supported in part by the Deutsche Forschungsgemeinschaft (DFG) grant “Information Flow Control for Browser Clients” under the priority program “Reliably Secure Software Systems” (RS³) and the DFG collaborative research center grant SFB 1223 “Methods and Tools for Understanding and Controlling Privacy”.

REFERENCES

- Martín Abadi, Anindya Banerjee, Nevin Heintze, and Jon G. Riecke. 1999. A core calculus of dependency. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, (POPL)*. 147–160.
- Thomas H. Austin and Cormac Flanagan. 2009. Efficient purely-dynamic information flow analysis. In *Proceedings of the 2009 Workshop on Programming Languages and Analysis for Security (PLAS)*. 113–124.
- Gilles Barthe, Pedro R. D’Argenio, and Tamara Rezk. 2011. Secure information flow by self-composition. *Mathematical Structures in Computer Science* 21, 6 (2011), 1207–1252.
- Pablo Buiras, Dimitrios Vytiniotis, and Alejandro Russo. 2015. HLIO: Mixing static and dynamic typing for information-flow control in Haskell. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming (ICFP)*. 289–301.
- Andrey Chudnov and David A. Naumann. 2015. Inlined Information Flow Monitoring for JavaScript. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 629–643.
- Eslam Elnikety, Aastha Mehta, Anjo Vahldiek-Oberwagner, Deepak Garg, and Peter Druschel. 2016. Thoth: Comprehensive Policy Compliance in Data Retrieval Systems. In *Proceedings of the 25th USENIX Security Symposium (USENIX Security)*. 637–654.
- Cédric Fournet, Gurvan Le Guernic, and Tamara Rezk. 2009. A security-preserving compiler for distributed programs: From information-flow policies to cryptographic mechanisms. In *Proceedings of the 16th ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 432–441.
- Joseph A. Goguen and José Meseguer. 1982. Security policies and security models. In *Proceedings of the 1982 IEEE Symposium on Security and Privacy (Oakland)*. 11–20.
- Daniel Hedin and Andrei Sabelfeld. 2012. Information-flow security for a core of JavaScript. In *Proceedings of the 25th IEEE Computer Security Foundations Symposium (CSF)*. 3–18.
- Nevin Heintze and Jon G. Riecke. 1998. The SLam calculus: Programming with secrecy and integrity. In *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 365–377.
- Sebastian Hunt and David Sands. 2006. On flow-sensitive security types. In *Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 79–90.
- Maxwell N. Krohn, Alexander Yip, Micah Z. Brodsky, Natan Cliffer, M. Frans Kaashoek, Eddie Kohler, and Robert Morris. 2007. Information flow control for standard OS abstractions. In *Proceedings of the 21st ACM Symposium on Operating Systems Principles (SOSP)*. 321–334.
- Peng Li and Steve Zdancewic. 2006. Encoding information flow in Haskell. In *Proceedings of the 19th IEEE Computer Security Foundations Workshop (CSFW)*.
- Ana Almeida Matos. 2006. *Typing secure information flow: Declassification and mobility*. Ph.D. Dissertation. École Nationale Supérieure des Mines de Paris.
- Andrew C. Myers. 1999. JFlow: Practical mostly-static information flow control. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 228–241.
- François Pottier and Vincent Simonet. 2003. Information flow inference for ML. *ACM Transactions on Programming Languages and Systems* 25, 1 (2003), 117–158.

- Vineet Rajani, Iulia Bastys, Willard Rafnsson, and Deepak Garg. 2016. *Fine-grained vs coarse-grained type systems for information flow control*. Technical Report MPI-SWS-2016-012. Max Planck Institute for Software Systems.
- Vineet Rajani, Abhishek Bichhawat, Deepak Garg, and Christian Hammer. 2015. Information flow control for event handling and the DOM in web browsers. In *Proceedings of the 28th IEEE Computer Security Foundations Symposium (CSF)*. 366–379.
- Alejandro Russo. 2015. Functional pearl: Two can keep a secret, if one of them uses Haskell. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming (ICFP)*. 280–288.
- Alejandro Russo, Koen Claessen, and John Hughes. 2008. A library for light-weight information-flow security in Haskell. In *Proceedings of the 1st ACM SIGPLAN Symposium on Haskell (Haskell)*. 13–24.
- Andrei Sabelfeld and David Sands. 2009. Declassification: Dimensions and principles. *Journal of Computer Security* 17, 5 (2009), 517–548.
- Geoffrey Smith and Dennis M. Volpano. 1998. Secure information flow in a multi-threaded imperative language. In *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 355–364.
- Deian Stefan, Alejandro Russo, John C. Mitchell, and David Mazières. 2011. Flexible dynamic information flow control in Haskell. In *Proceedings of the 4th ACM SIGPLAN Symposium on Haskell (Haskell)*. 95–106.
- Dennis M. Volpano, Cynthia E. Irvine, and Geoffrey Smith. 1996. A sound type system for secure flow analysis. *Journal of Computer Security* 4, 2/3 (1996), 167–188.
- Nickolai Zeldovich, Silas Boyd-Wickizer, Eddie Kohler, and David Mazières. 2006. Making information flow explicit in HiStar. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI)*. 263–278.