

Pure Information-Flow Control with Effects Made Simple

No Author Given

No Institute Given

Abstract. Information-flow control (IFC) is a promising technology to protect data confidentiality. The foundational work on the Dependency Core Calculus (DCC) positions monads as a suitable abstraction for enforcing IFC. Pure functional languages with effects, like Haskell, can provide IFC as a library (MAC, LIO etc.), a minor task compared to implementing compilers for IFC from scratch.

Previous works on IFC as a library introduce ad hoc primitives to type programs whose effects do not depend on the sensitive data in context. In this work, we start afresh and ask ourselves: what would we need to extend an effect-free language for IFC (e.g., DCC) with secure effects? The answer turns out to be elegant and simple. In a pure language with effects there is a natural place where information flows from sensitive data to effects need to be restricted, and when effects are tracked in a more fine-grained fashion than Haskell’s IO monad (e.g., with a graded monad) then a *single* primitive is enough to allow the secure flows!

To support our insight, we present and prove secure several IFC enforcement mechanisms based on extensions of the Sealing Calculus (SC) with effects using a graded monad. Effects that depend on sensitive data are secured through a novel primitive `distr`. All of our security guarantees are mechanized in the Agda proof assistant. Moreover, we provide an implementation of these mechanisms as a new Haskell library for IFC.

1 Introduction

Information-flow control (IFC) [35, 13] is a promising technology to protect data confidentiality. Many IFC approaches are designed to prevent sensitive data from influencing what attackers can observe from a program’s public behavior—a security property known as noninterference [11]. IFC mechanisms specify the sensitivity of data via *labels*, and then enforce security by controlling that the flows of information abide by the security policy. In the simplest scenario, there are two labels (alternatively, security levels or sensitivities) `public` and `secret` and the security policy specifies that every flow is allowed except from `secret` to `public`—i.e., flows from more to less sensitivity are forbidden. In static approaches to IFC, the sensitivity of data is known a priori, e.g., by specifying it using types, and the enforcement statically decides, e.g., during type checking, whether the program will leak information upon execution. To protect confidentiality of data, IFC mechanisms need to protect against two kinds of potentially malicious flows: 1) *explicit flows*, when `public` data directly depends on `secret` data; and

2) *implicit flows*, when the control flow and **public** outputs of the program are indirectly influenced by **secret** data, e.g., due to branching on a **secret**.

In recent years, the use of pure functional languages has been proliferating for tackling different IFC challenges (e.g., Vassena et al. [44], Parker, Vazou, and Hicks [29], Polikarpova et al. [31], and Rajani and Garg [32]). From a pragmatic perspective, pure functional languages can provide IFC security via libraries [22, 34, 39], which is less demanding than building compilers from scratch (e.g., Simonet [38] and Myers et al. [26]).

When it comes to controlling information flows, pure languages stand in a privileged position thanks to their abstraction facilities and strong encapsulation of effects. For instance, the popular Dependency Core Calculus (DCC) [1] utilizes the abstract type $T_{\text{secret}} A$ to label pieces of data of type A with sensitivity **secret** and then the type system ensures that data can only be eliminated into—or influence—data of equal or higher sensitivity. DCC’s security guarantees ensure that programs without effects are secure.

A different strand of work aims to provide security in a pure language by restricting the interplay between sensitive data and **public** effects (e.g., LIO [39], MAC [33] and HLIO [9]). In a pure language like Haskell, the *only* programs that can produce effects and thereby interact with the external world have to be of type $\text{IO } A$, for some type A . In this light, in order to protect against implicit flows through effects it is enough to control which programs of IO type are safe to execute. To realize this idea, the MAC library, for example, replaces the IO monad with a custom monad MAC_l of computations that is indexed by a type-level label l . The label l has two purposes: (i) akin to DCC, it is an *upper bound* on the sensitivity of the information “going into” the monad, as well as (ii) it is a *lower bound* on the observers’ effects—restricting information “leaving” the monad. Concretely, a computation of type $\text{MAC}_{\text{public}} \text{Bool}$ *cannot* branch on **secret** values but can perform **public** and **secret** effects; in contrast, a computation of type $\text{MAC}_{\text{secret}} \text{Bool}$ can branch on **secret** data but *cannot* perform **public** effects. In this manner, MAC prevents explicit and implicit flows by construction.

However, not everything in the garden is rosy. The label l in the MAC monad MAC_l does too many things at once. This leads to situations where the programmer needs to go through some contortions or use ad hoc primitives like join^1 : $\text{MAC}_l \text{Unit} \Rightarrow \text{MAC}_{l'} \text{Unit}$ (restricted to $l' \sqsubseteq l$) [41]. To illustrate this point, we extend the two-point security policy with two incomparable labels Alice and Bob such that $\text{public} \sqsubseteq \text{Alice} \sqsubseteq \text{secret}$ and $\text{public} \sqsubseteq \text{Bob} \sqsubseteq \text{secret}$ but neither $\text{Alice} \sqsubseteq \text{Bob}$ nor $\text{Bob} \sqsubseteq \text{Alice}$ (the relation \sqsubseteq specifies the permitted flows). With that in mind, let us consider the following program in MAC which receives an Alice-sensitive Bool , i.e., $\text{alice_sec} : \text{MAC}_{\text{Alice}} \text{Bool}$,² and prints a

¹ The type of join is more general but this simplified form suffices for our purposes.

² To the reader familiar with MAC: our point applies equally if one uses the Labeled_l type to protect sensitive Booleans.

string on the Alice-observable channel Ch_{Alice} :

```

prog1 : MACAlice Bool ⇒ MACAlice Unit
prog1 alice_sec = alice_sec >>= λb. if b then printChAlice("Alice is here!")
                                     else return(unit)
    
```

In the above program, the information flows according to the security policy (from the Alice value to Alice’s channel)—i.e., the program is secure. Consider now a different program that combines $prog_1$ with printing on the channel Ch_{Bob} :

```

prog2 : MACAlice Bool ⇒ MAC? Unit
prog2 alice_sec = prog1 alice_sec >> printChBob("Hi Bob")
    
```

Clearly, $prog_2$ is still secure since the decision to print to Bob’s channel and what is printed does not depend on the contents of the Alice-sensitive value $alice_sec$. What label then should replace $?$ in its type? First, the types of the computations on both sides of the bind (\gg) have to match. By (i) and (ii) the type of $prog_1$ sec has to be $\text{MAC}_{\text{Alice}} \text{Unit}$, and by (ii) $\text{print}_{\text{Ch}_{\text{Bob}}}(\text{"Hi Bob"})$ has to be of type $\text{MAC}_{\text{Bob}} \text{Unit}$. There is a type mismatch! Since the label **public** is a lower bound of Alice and Bob, we can use MAC’s join to fix the program:

```

prog2 : MACAlice Bool ⇒ MACpublic Unit
prog2 alice_sec = join(prog1 alice_sec) >> join(printChBob("Hi Bob"))
    
```

Now, something unexpected happened. To assign a type to $prog_2$, which only mentions labels Alice and Bob, we *have to* resort to another label **public**. What is worse, in the type of $prog_2$ only the label Alice appears and does so in an argument position which means we know nothing about the program’s possible effects. The design decision of indexing the monad MAC by a *single* label, while enabling a simple implementation of IFC as a library (cf. [33]), requires the application of ad hoc primitives (like join) to mitigate over-approximations of how information flows in and out of computations.

In this work, we take a step back and ask ourselves: what would it take to allow arbitrary effects in a pure language with an already existing mechanism for effect-free IFC (e.g., DCC)? The answer turns out to be elegant and simple. We observe that enforcing IFC in a pure language with effects can be reduced to the *single* point—which we will explain below—where effects and sensitive data interact. Based on this observation, we present a novel IFC mechanism which is arguably simpler than existing IFC libraries, allows to assign more natural types to programs; and overcomes the programming contortions discussed above.

To briefly present our idea, let us assume that effect-free IFC is achieved using a DCC-style abstract type $T_l A$, and we have at our disposal a more refined type constructor **Eff** of effectful programs akin to **IO** but annotated with concrete information about *observable* effects. For example, **Eff** could be annotated with the set of channels where the program might print, the set of exceptions the program might raise, or the set of memory references the program might modify. Moreover, we assume the security policy to specify the sensitivity of each effect.

Consider, for instance, a scenario with two output channels, namely $\text{Ch}_{\text{public}}$ and $\text{Ch}_{\text{secret}}$, that are assigned sensitivities **public** and **secret**, respectively.

In this scenario, a program $prog_3 : T_{\text{Alice}}(\text{Eff}_{\{\text{Ch}_{\text{Bob}}\}} \text{Unit})$ is a computation that might print to Bob’s channel Ch_{Bob} depending on Alice-sensitive data. This is where sensitive data and effects interact! Assuming $alice_sec : T_{\text{Alice}} \text{Bool}$ is in scope, $prog_3$, for instance, could have the following implementation:³

```

prog3 : TAlice (Eff{ChBob} Unit)
prog3 = bind b = alice_sec
      in returnAlice(if b then printChBob("So it was true, huh")
                    else return(unit))

```

In order to run the effects of the inner computation of type $\text{Eff}_{\{\text{Bob}\}} \text{Unit}$ we have to “extract” it first from the T_{Alice} value—recall that the language is pure. In general, extracting anything from T_{Alice} is prohibited as it would render the IFC enforcement unsound: in $prog_3$, the computation of type $\text{Eff}_{\{\text{Ch}_{\text{Bob}}\}} \text{Unit}$ would become executable, and if the Alice-sensitive Boolean $alice_sec$ is true, it will print “So it was true, huh” on Bob’s channel Ch_{Bob} —a flow that violates the security policy. On the contrary, let us consider a program of a different type, $prog_4 : T_{\text{Alice}}(\text{Eff}_{\{\text{Ch}_{\text{secret}}\}} \text{Unit})$. In this program, the computation is *definitely* safe to run since the only possible effects it might produce are printing to the **secret** channel $\text{Ch}_{\text{secret}}$ and the decision on what to print depends on data of at most sensitivity Alice—a flow permitted by the security policy. Securing effectful programs then amounts to allowing $prog_4$ to extract the computation of type $\text{Eff}_{\{\text{Ch}_{\text{secret}}\}} \text{Unit}$ from T_{Alice} and run its observable effects while forbidding $prog_3$ from doing so. To achieve this, we introduce a novel primitive `distr` which systematically permits computations to be extracted, and hence, executed only when they are known to depend on data less sensitive than their observers. With `distr` we can turn $prog_4$ into an executable program: $\text{distr}(prog_4) : \text{Eff}_{\{\text{Ch}_{\text{secret}}\}} \text{Unit}$. We have reduced the enforcement of IFC in pure languages with effects to a single primitive; this gives us *modularity*, *clarity* and *simplicity* in the language design and its possible implementations. Thanks to `distr`, we can express $prog_2$ with a more natural type:

```

prog'_2 : TAlice Bool ⇒ Eff{ChAlice, ChBob} Unit
prog'_2 alice_sec = distr(bind b = alice_sec
                        in return(if b then printChAlice("Alice is here!")
                                else return(unit)))
                >> printChBob("Hi Bob")

```

Our contributions In this work, we show that IFC in the context of pure languages with effects can be achieved through the combination of the following features: 1. an enforcement for effect-free IFC, 2. a type for tracking observable effects in a fine-grained fashion, and 3. a primitive `distr` which selectively permits to execute effectful computations which depend on sensitive data. We present our idea through an IFC enforcement mechanism for the simply-typed lambda-calculus (STLC). The mechanism is based on the Sealing Calculus (SC) [37], which is more expressive than DCC (cf. [40, 37]), although we could have chosen DCC instead. We then extend the language and IFC enforcement mechanism

³ `bind` is DCC’s eliminator for the type T_I .

in two different directions with printing and global store effects in the form of graded monads [18]. In these extensions, we include the `distr` primitive which allows to type check more effectful programs that are secure. Along with our informal argumentation for why our idea is secure for different instantiations of effects, we have mechanized proofs in the Agda proof assistant [2] about the strong security guarantees that the programs in the languages satisfy, namely termination-insensitive noninterference (TINI). Our proofs are based on the technique of logical relations [23]. Finally, we realize our idea in the form of a new Haskell library which we call SCLib. The conciseness of our implementation illustrates the elegance and simplicity of our insight: less than 10 lines of code for the effect-free fragment and less than 30 for the part with effects. In order to implement the effect-free IFC mechanism we also present a novel implementation of SC using an encoding of contextual information as type-level capabilities via higher-rank polymorphism [20]. Our library is at least as expressive as previous work on libraries for IFC in Haskell, which we evidence by showing implementations of SecLib [34], DCC (in its presentation by Algehed [4]) and MAC in terms of SCLib’s interface.

In summary, the technical contributions of this paper are:

- A reinterpretation of the Sealing Calculus as an IFC enforcement mechanism for STLC (Section 2)
- Two extensions of STLC and SC for enforcing IFC in pure languages with effects via graded monads. As examples, we consider printing (Section 3.1) and global store (Section 3.2) effects
- We present `distr`, a single primitive that can control the interaction of sensitive data and effects
- Security guarantees and proofs of TINI based on logical relations for all the enforcements (Section 4)
- A Haskell implementation using a novel encoding of contextual information as capabilities together with evidence that SCLib can encode existing monadic security libraries (Section 5)
- Mechanized proofs of all our security guarantees (approx. 1500 lines of Agda code submitted as accompanying material)

2 Effect-free Information-Flow Control

In this section, we briefly recall the Sealing Calculus (SC) and explain its role as an IFC enforcement mechanism for programs written in the STLC.

SC utilizes an abstract type $S_l A$ for protecting sensitive data.⁴ A value of type $S_l A$ is “sealed” at sensitivity l in the sense that it is only available to observers with sensitivity at least as high as l . Values of type $S_l A$ are introduced and eliminated using the primitives `seall` and `unseall`. SC enforces IFC by restricting in which *contexts* a sealed value can be “unsealed”. For example, the Alice-sensitive Boolean `sec :: SAlice Bool` can only be unsealed in contexts of at least sensitivity `Alice`.

⁴ In the original presentation, the authors use the notation $[A]^l$ instead.

Let us illustrate SC with a program that receives two Booleans with sensitivities `Alice` and `Bob`, respectively, and computes their conjunction with sensitivity `secret` (`and : Bool × Bool ⇒ Bool` implements conjunction):

$$\begin{aligned} \text{and}' &: S_{\text{Alice}} \text{Bool} \Rightarrow S_{\text{Bob}} \text{Bool} \Rightarrow S_{\text{secret}} \text{Bool} \\ \text{and}' &= \lambda(sb_1 sb_2. \text{seal}_{\text{secret}}(\text{and}(\text{unseal}_{\text{Alice}}(sb_1), \text{unseal}_{\text{Bob}}(sb_2)))) \end{aligned}$$

In the above program, the term `sealsecret` provides the context in which `sb1` and `sb2` can be unsealed: the term `unsealAlice(sb1)`, for instance, is only well-typed because its label `Alice` flows to `secret`, which is the highest label.

We take a variant of the STLC as the programming language on which we want to enforce IFC: it has call-by-name semantics and `Unit` and `Bool` as the only primitive types. We work directly with intrinsically well-typed terms, and thus, for us a typing derivation $\Gamma \vdash t : a$ is a term. The small-step semantics is specified by the relation $t \longrightarrow t'$, and we call values those terms t for which $t \not\longrightarrow$. For reference we include its complete definition in Appendix A. Since the STLC is well-understood we focus on SC.

Figure 1 presents the intrinsically-typed syntax of SC. The whole development is parameterized by a security policy specified in the form of a lattice structure on the set of labels $(\mathcal{L}, \sqsubseteq)$. The types reflect those in the STLC and include the new type constructor S_l . Typing judgements are of the form $\pi ; \Gamma \vdash^{\text{sc}} t : A$ where: π is a non-empty *finite* set of labels drawn from \mathcal{L} , i.e., $\pi \subseteq \mathcal{L}$; Γ is an SC typing context; and A is an SC type. The component π is analogous to protection contexts from related work by Tse and Zdancewic [40].

The set of labels π in the typing judgement represent the sensitivities of all the data on which the program may depend. In order to clarify the role of π , let us consider a program with a typing derivation indexed by the set of labels $\pi_1 := \{\text{secret}\}$, $\pi_1 ; \cdot \vdash^{\text{sc}} p : A$. This program may depend on data of types $S_{\text{public}} A$ and $S_{\text{secret}} A$ by unsealing. If, instead, p is indexed by the set $\pi_2 := \{\text{public}\}$, i.e., $\pi_2 ; \cdot \vdash^{\text{sc}} p : A$, then the only terms that the program can unseal are of type $S_{\text{public}} A$. This mechanism ensures that the flows of information are secure. It is useful to think that the labels that belong to π act as a kind of type-level *key* whose “possession” permits access to information at most as sensitive as the label itself.

The typing rules of the STLC fragment of SC, i.e., Rules `LAM`, `APP` and `IF`, are rather standard: they simply propagate the set of labels π to their premises. Observe that one has to explicitly unseal sensitive Booleans, i.e., of type $S_l \text{Bool}$, in order to branch on them using the Rule `IF`. Rules `SEAL` and `UNSEAL` are the most interesting since they enforce that information flows to the appropriate places. Rule `SEAL` serves a double purpose: from premise to conclusion, it introduces terms of type $S_l A$; and, from conclusion to premise, it extends the set of labels with the label l , i.e., $\pi \cup \{l\}$. The typing derivation above the premise can then unseal any term of type $S_{l'} A$ such that its label l' can flow to l . Rule `UNSEAL` allows unsealing a term with type $S_l A$ if the set π in its conclusion contains at least a label l' such that $l \sqsubseteq l'$. Continuing with the intuition of labels in π as keys, a key $l' \in \pi$ can be used to unseal terms of type $S_l A$ exactly when $l \sqsubseteq l'$.

Types	$A, B ::= \mathbf{Unit} \mid \mathbf{Bool} \mid A \Rightarrow B \mid \mathbf{S}_l A$	Sets of labels	$\pi \subseteq \mathcal{L}$
Typing contexts	$\Gamma ::= \cdot \mid \Gamma, x : A$		
$\pi ; \Gamma \vdash^{\text{sc}} t : A$			
VAR	$\frac{(x : A) \in \Gamma}{\pi ; \Gamma \vdash^{\text{sc}} x : A}$	LAM	$\frac{\pi ; \Gamma, x : A \vdash^{\text{sc}} t : B}{\pi ; \Gamma \vdash^{\text{sc}} \lambda(x.t) : A \Rightarrow B}$
		APP	$\frac{\pi ; \Gamma \vdash^{\text{sc}} t : A \Rightarrow B \quad \pi ; \Gamma \vdash^{\text{sc}} u : A}{\pi ; \Gamma \vdash^{\text{sc}} \mathbf{app}(t, u) : B}$
UNIT		TRUE	
$\frac{}{\pi ; \Gamma \vdash^{\text{sc}} \mathbf{unit} : \mathbf{Unit}}$		$\frac{}{\pi ; \Gamma \vdash^{\text{sc}} \mathbf{true} : \mathbf{Bool}}$	
		FALSE	
		$\frac{}{\pi ; \Gamma \vdash^{\text{sc}} \mathbf{false} : \mathbf{Bool}}$	
IF			
$\frac{\pi ; \Gamma \vdash^{\text{sc}} t : \mathbf{Bool} \quad \pi ; \Gamma \vdash^{\text{sc}} u_1 : A \quad \pi ; \Gamma \vdash^{\text{sc}} u_2 : A}{\pi ; \Gamma \vdash^{\text{sc}} \mathbf{ifte}(t, u_1, u_2) : A}$			
SEAL		UNSEAL	
$\frac{\pi \cup \{l\} ; \Gamma \vdash^{\text{sc}} t : A}{\pi ; \Gamma \vdash^{\text{sc}} \mathbf{seal}_l(t) : \mathbf{S}_l A}$		$\frac{\pi ; \Gamma \vdash^{\text{sc}} t : \mathbf{S}_l A \quad \exists l' \in \pi. l \sqsubseteq l'}{\pi ; \Gamma \vdash^{\text{sc}} \mathbf{unseal}_l(t) : A}$	

Fig. 1. Types and intrinsically-typed terms of SC

In order to use SC to enforce IFC on STLC programs, we define an erasure function from SC terms $\pi ; \Gamma \vdash^{\text{sc}} t : A$ to STLC programs $\varepsilon(\Gamma) \vdash \varepsilon(t) : \varepsilon(A)$:

$$\begin{array}{lll}
 \varepsilon(\mathbf{Unit}) = \mathbf{Unit} & \varepsilon(\cdot) = \cdot & \varepsilon(\lambda(x.t)) = \lambda(x.\varepsilon(t)) \\
 \varepsilon(\mathbf{Bool}) = \mathbf{Bool} & \varepsilon(\Gamma, x : A) = \varepsilon(\Gamma), x : \varepsilon(A) & \varepsilon(\mathbf{app}(t, u)) = \mathbf{app}(\varepsilon(t), \varepsilon(u)) \\
 \varepsilon(A \Rightarrow B) = \varepsilon(A) \Rightarrow \varepsilon(B) & & \varepsilon(\mathbf{seal}_l(t)) = \varepsilon(t) \\
 \varepsilon(\mathbf{S}_l A) = \varepsilon(A) & & \varepsilon(\mathbf{unseal}_l(t)) = \varepsilon(t)
 \end{array}$$

To clarify this point further, the noninterference property enforced by an SC term $\{\mathbf{public}\} ; \mathit{sec} : \mathbf{S}_{\mathbf{secret}} \mathbf{Bool} \vdash^{\text{sc}} t : \mathbf{Bool}$ on its underlying STLC program $\mathit{sec} : \mathbf{Bool} \vdash \varepsilon(f) : \mathbf{Bool}$ is that for all $\cdot \vdash s_1, s_2 : \mathbf{Bool}$ whenever both $\varepsilon(t)[s_1/\mathit{sec}]$ and $\varepsilon(t)[s_2/\mathit{sec}]$ terminate then they do so with the same Boolean. In this way, SC types and typing contexts play the role of security specifications, and SC terms of evidence that the underlying programs are secure, i.e., they satisfy the security specification. Indeed, for us SC terms do not have operational semantics, only their underlying STLC programs do. When convenient, we identify SC terms with the erased STLC programs. We observe that this interpretation of SC as an IFC enforcement mechanism is closed under the operational semantics of STLC:

Lemma 1. *Given a term $\pi ; \cdot \vdash^{\text{sc}} t : A$ such that $\varepsilon(t) \longrightarrow f$ then there exists a term $\pi ; \cdot \vdash^{\text{sc}} t' : A$ such that $\varepsilon(t') = f$.*

3 Effectful Information-Flow Control

In this section, we present the main contribution of this paper: the observation that a single primitive \mathbf{distr} is enough to extend IFC in pure languages with effects. We study two extensions of the programming language and the IFC enforcement mechanism from Section 2 with printing and global store effects.

In these extensions, we treat effects *explicitly* in the style of Haskell’s IO monad [16] and Moggi’s monadic metalanguage [25]: the only programs that can perform effects are of type $\text{Eff}_C a$ for some effect annotation C and type a , and sequencing of effects is made explicit through the primitive `bind`. Specifically, we consider printing and global store effects as suitable representatives of the two kinds of effects that need to be secured:

Printing Effects. Printing on a channel can be observed externally to the program by the channel’s observers. Observers can infer information about the program’s input from what is being printed to the channel. To secure printing effects one must ensure that the decision to print and what is being printed only depends on data less sensitive than the channel’s observers.

Global Store Effects. Reading from the store cannot be observed directly. However, reading effects need to be secured because what is read may influence the program’s subsequent behavior. To secure reading effects one must ensure that what has been read is tracked as sensitive data.

In contrast to Haskell and the metalanguage, we employ a graded monad [18] whose effect annotation C tracks precisely to which channels a computation might print and which store locations a computation might access.

3.1 Printing Effects

In Figure 2 we present the extension of STLC which allows programs to perform printing effects. We dub this language $\text{STLC}^{\text{PRINT}}$. We assume that the set of printing channels \mathcal{Ch} is fixed a priori, i.e., the channels are statically known. The set of types is extended with a new type for computations $\text{Eff}_C a$ that is indexed by a set of channels $C \subseteq \mathcal{Ch}$. A program of type $\text{Eff}_C a$ when executed might *only* print to the channels that appear in C and return a result of type a .

The typing rules of the standard monadic operations, `return` and `bind`, are as expected: in Rule `RETURN`, the computation does not perform any effects thus the type is indexed by the empty set of channels; and, in Rule `BIND`, the type is indexed by the union of the channels on which the computations $\Gamma \vdash t : \text{Eff}_{C_1} a$ and $\Gamma \vdash u : a \Rightarrow \text{Eff}_{C_2} b$ might print, that is, $C_1 \cup C_2$. We include subeffecting—casting from a smaller to a larger set of channels—as the term `subeff` in the language, see Rule `SUBEFF`. Printing is performed via a family of primitive operations, `printch`, one for each available channel $ch \in \mathcal{Ch}$ (Rule `PRINT`). We assume, for simplicity, that only Boolean values can be printed. Further, observe that the resulting monadic type is indexed by the singleton set that only contains the channel on which the printing is performed, i.e., $\text{Eff}_{\{ch\}} \text{Unit}$.

The operational semantics of $\text{STLC}^{\text{PRINT}}$ is defined as the combination of the small-step operational semantics of STLC—see Appendix A—and the small-step operational semantics of computations defined in Figure 2. The semantics of effect-free terms, inherited from STLC, $t \longrightarrow u$, treats computations, such as `return(t)`, as values even when their subterms are not values, e.g., `return(app(λ(x.x), true))` $\not\longrightarrow$. The semantics of computations (alternatively, monadic semantics) is of the form $t \rightsquigarrow u, o$ and is interpreted as follows: program $\cdot \vdash t : \text{Eff}_C a$ evaluates in one step to program $\cdot \vdash u : \text{Eff}_C a$ and produces

output o . The output is a function from channels to lists of Boolean values $o \in \mathcal{Ch} \rightarrow \text{List Bool}$, and it represents the outputs of the program during execution.

Types	$a, b ::= \dots \mid \text{Eff}_C a$	Sets of channels	$C, C_1, C_2 \subseteq \mathcal{Ch}$
Typing contexts	$\Gamma ::= \dots$	Outputs	$o, o_1, o_2 \in \mathcal{Ch} \rightarrow \text{List Bool}$
$\boxed{\Gamma \vdash t : a}$			
RETURN	$\frac{\Gamma \vdash t : a}{\Gamma \vdash \text{return}(t) : \text{Eff}_\emptyset a}$	BIND	$\frac{\Gamma \vdash t : \text{Eff}_{C_1} a \quad \Gamma \vdash u : a \Rightarrow \text{Eff}_{C_2} b}{\Gamma \vdash \text{bind}(t, u) : \text{Eff}_{C_1 \cup C_2} b}$
SUBEFF	$\frac{\Gamma \vdash t : \text{Eff}_{C_1} a \quad C_1 \subseteq C_2}{\Gamma \vdash \text{subeff}(t) : \text{Eff}_{C_2} a}$	PRINT	$\frac{\Gamma \vdash t : \text{Bool}}{\Gamma \vdash \text{print}_{ch}(t) : \text{Eff}_{\{ch\}} \text{Unit}}$
$\boxed{t \rightsquigarrow u, o \text{ with } \cdot \vdash t : \text{Eff}_C a \text{ and } \cdot \vdash u : \text{Eff}_C a}$			
BIND	$\frac{t \rightsquigarrow t', o}{\text{bind}(t, u) \rightsquigarrow \text{bind}(t', u), o}$	BIND-RET	$\frac{}{\text{bind}(\text{return}(t), u) \rightsquigarrow \text{app}(t, u), \epsilon}$
PRINT	$\frac{t \longrightarrow u}{\text{print}_{ch}(t) \rightsquigarrow \text{print}_{ch}(u), \epsilon}$	PRINT-TRUE	$\frac{}{\text{print}_{ch}(\text{true}) \rightsquigarrow \text{return}(\text{unit}), ch \mapsto [\text{true}]}$
PRINT-FALSE	$\frac{}{\text{print}_{ch}(\text{false}) \rightsquigarrow \text{return}(\text{unit}), ch \mapsto [\text{false}]}$	EFFECTFREE	$\frac{t \longrightarrow u}{t \rightsquigarrow u, \epsilon}$

Fig. 2. Types, well-typed terms and small-step semantics of $\text{STLC}^{\text{PRINT}}$ (excerpts)

We now briefly explain the semantics. Rule **BIND** reduces the left subterm of bind and executes its effects, o . Once the left subterm is a value, i.e., $\text{return}(t)$, **BIND-RET** applies the rest of the computation u to the result t . Applying the continuation u does not produce effects—recall that we are in a pure language—thus the step contains the empty output on the right, i.e., ϵ . The empty output maps every channel to the empty output, i.e., $\epsilon := \lambda ch. []$. Rule **PRINT** reduces the argument t of $\text{print}_{ch}(t)$ until it is a value of type **Bool**, either **true** or **false**, and then Rules **PRINT-TRUE** and **PRINT-FALSE** print the corresponding Boolean on the output channel ch . The output $ch \mapsto [v]$ is the function that maps the channel ch to the singleton list $[v]$ and every other channel to the empty list. Observe that these rules make $\text{print}_{ch}(t)$ strict in its argument. Rule **EFFECTFREE** serves to lift effect-free reductions to the level of computations. Since by definition

effect-free reductions do not produce effects, the right hand side of the effect contains the empty output ϵ . To complete the picture, we denote by $t \rightsquigarrow^* u, o$ the reflexive-transitive closure of the monadic reduction relation. To combine effects, we lift the monoid structure on List Bool to outputs. The rules that handle the term `subeff` (omitted) are straightforward.

Two reduction relations. While it might seem unnecessary to define the semantics using the combination of a small-step relation of effect-free programs and a small-step relation of computations, it is a natural form of expressing the operational semantics of pure languages with effects [45]. The effect-free relation evaluates programs that *cannot* perform effects, whilst the relation for computations evaluates programs which can, and computes those effects.

To conclude the presentation of $\text{STLC}^{\text{PRINT}}$, we enunciate the following lemma which states that the index C in the type of computations $\text{Eff}_C a$ is a sound approximation of the set of channels where a program $\cdot \vdash t : \text{Eff}_C a$ may print, i.e., t does not produce output in any channel not in C . In the security literature it is usually called the *confinement lemma*:

Lemma 2 (Confinement for $\text{STLC}^{\text{PRINT}}$). *For any $\text{STLC}^{\text{PRINT}}$ program of type $\cdot \vdash f : \text{Eff}_C a$, $\text{STLC}^{\text{PRINT}}$ value $\cdot \vdash v : \text{Eff}_C a$, and output o , if $t \rightsquigarrow^* v, o$ then $\forall ch \in \text{Ch}. ch \notin C \Rightarrow o(\text{ch}) = []$.*

Types $A, B ::= \dots \mid \text{Eff}_C A$	Sets of channels $C, C_1, C_2 \subseteq \text{Ch}$
Typing contexts $\Gamma ::= \dots$	Sensitivity of channels $\text{label} \in \text{Ch} \rightarrow \mathcal{L}$
$\pi ; \Gamma \vdash^{\text{sc}} t : A$	
$\frac{\text{RETURN} \quad \pi ; \Gamma \vdash^{\text{sc}} t : A}{\pi ; \Gamma \vdash^{\text{sc}} \text{return}(t) : \text{Eff}_\emptyset A}$	$\frac{\text{BIND} \quad \pi ; \Gamma \vdash^{\text{sc}} t : \text{Eff}_{C_1} A \quad \pi ; \Gamma \vdash^{\text{sc}} u : A \Rightarrow \text{Eff}_{C_2} B}{\pi ; \Gamma \vdash^{\text{sc}} \text{bind}(t, u) : \text{Eff}_{C_1 \cup C_2} B}$
$\frac{\text{SUBEFF} \quad \pi ; \Gamma \vdash^{\text{sc}} t : \text{Eff}_{C_1} A \quad C_1 \subseteq C_2}{\pi ; \Gamma \vdash^{\text{sc}} \text{subeff}(t) : \text{Eff}_{C_2} A}$	$\frac{\text{PRINT} \quad \pi ; \Gamma \vdash^{\text{sc}} t : \text{Bool}}{\pi ; \Gamma \vdash^{\text{sc}} \text{print}_{\text{ch}}(t) : \text{Eff}_{\{\text{ch}\}} \text{Unit}}$
$\frac{\text{DISTR} \quad \pi ; \Gamma \vdash^{\text{sc}} t : \text{S}_l(\text{Eff}_C A) \quad \forall ch \in C. l \sqsubseteq \text{label}(ch)}{\pi ; \Gamma \vdash^{\text{sc}} \text{distr}(t) : \text{Eff}_C(\text{S}_l A)}$	

Fig. 3. Types and intrinsically-typed terms of SC^{PRINT}

After having defined the programming language, we are in position to turn our attention to enforcing IFC. We assume that the security policy specifies the sensitivity of each printing channel, i.e., the greatest lower bound of the sensitivities of all its observers, in the form of a function $\text{label} \in \text{Ch} \rightarrow \mathcal{L}$. In

Figure 3 we present the extension of SC (Figure 1) that accommodates printing effects. We name it SC^{PRINT} hereafter. The types reflect those in SC and include the new type constructor $\text{Eff}_{\mathcal{C}}$ of computations. The typing rules for the $\text{STLC}^{\text{PRINT}}$ fragment of SC^{PRINT} , i.e., Rules RETURN to PRINT, simply propagate the set of labels π to their premises. Observe, again, that one has to explicitly unseal sensitive Booleans, i.e., of type $S_l \text{Bool}$, to apply the Rule PRINT. Before detailing Rule DISTR, we extend the erasure function ε from SC^{PRINT} -terms to $\text{STLC}^{\text{PRINT}}$ programs in the obvious way: i.e., the type former $\text{Eff}_{\mathcal{C}}$ erases to “itself” and, analogously to seal_l and unseal_l , distr is a no-op.

$$\begin{aligned} \varepsilon(\dots) &= \dots & \varepsilon(\dots) &= \dots \\ \varepsilon(\text{Eff}_{\mathcal{C}} A) &= \text{Eff}_{\mathcal{C}} \varepsilon(A) & \varepsilon(\text{distr}(t)) &= \varepsilon(t) \end{aligned}$$

Rule DISTR introduces one of the novelties of our work; an enforcement mechanism that selectively permits to execute the effects of computations that depend on sensitive data. In SC^{PRINT} , a term of type $S_l (\text{Eff}_{\mathcal{C}} A)$ describes a computation that might *only* print on the channels in \mathcal{C} and what is printed and the decision to print potentially depends on data of sensitivity l . Then, it is natural to ask, when it is secure to execute the effects of the inner computation of type $\text{Eff}_{\mathcal{C}} A$? Clearly, whenever the sensitivities of the computation’s effects, i.e., the channels in \mathcal{C} , are as high as the sensitivity of the data used to decide to perform the effects, i.e., sensitivity l . The premise of the rule exactly captures this condition: $\forall ch \in \mathcal{C}. l \sqsubseteq \text{label}(ch)$.

To illustrate DISTR in action, let us consider the following term in SC^{PRINT} that prints Alice’s sensitive input to the $\text{Ch}_{\text{secret}}$ channel:

$$\begin{aligned} \text{prog}_6 &: S_{\text{Alice}} \text{Bool} \Rightarrow \text{Eff}_{\{\text{Ch}_{\text{secret}}\}} \text{Unit} \\ \text{prog}_6 &= \lambda(sb. \text{distr}(\text{seal}_{\text{Alice}}(\text{print}_{\text{Ch}_{\text{secret}}}(\text{unseal}_{\text{Alice}}(sb)))) \gg \text{return}(\text{unit})) \end{aligned}$$

Inside the term λ , the primitive distr permits to execute the effects of the term $\text{seal}_{\text{Alice}}(\text{print}_{\text{Ch}_{\text{secret}}}(\text{unseal}_{\text{Alice}}(sb)))$ of type $S_{\text{Alice}} (\text{Eff}_{\{\text{Ch}_{\text{secret}}\}} \text{Bool})$. The term distr protects the return type of the computation at the same sensitivity as the premise’s type $S_l A$. This requirement is necessary to enforce IFC because the trivial computation that performs no effects and just returns has access to sensitive data, as exemplified by the following program:

$$\begin{aligned} \text{prog}_7 &: S_{\text{Alice}} \text{Bool} \Rightarrow S_{\text{Bob}} \text{Bool} \Rightarrow \text{Eff}_{\{\text{Ch}_{\text{secret}}\}} (S_{\text{secret}} \text{Bool}) \\ \text{prog}_7 &= \lambda(sb_1 sb_2. \text{distr}(\text{seal}_{\text{secret}}(\text{print}_{\text{Ch}_{\text{secret}}}(\text{unseal}_{\text{Alice}}(sb_1)))) \\ &\quad \gg \text{return}(\text{unseal}_{\text{Bob}}(sb_2))) \end{aligned}$$

3.2 Global Store Effects

We turn our attention to global store effects which combines printing effects from the previous section with reading from locations in the store. Following the same steps, in Figure 4 we present the extension of STLC (Appendix A) in which programs have access to a global store and can read from and write to it. We name this language $\text{STLC}^{\text{STORE}}$. Our development rests on two assumptions: 1. the set of locations in the store Loc is fixed during execution, and 2. only terms of ground type, i.e., Bool and Unit , can be stored. This helps simplify our presentation and permit us to side-step issues which are orthogonal to the point

of this work. Recent work by Rajani and Garg [32] and Gregersen et al. [12] hints to possible ways of lifting these two simplifications.

$\text{STLC}^{\text{STORE}}$ extends STLC with a type of computations $\text{Eff}_{\mathcal{S}} a$, indexed by a set of store locations \mathcal{S} (alternatively, references), and a type of references $\text{Ref}_s r$. By assumption 1, the whole development is parameterized by a fixed set of store locations Loc . A program of type $\text{Eff}_{\mathcal{S}} a$ when executed might *only* write to the references in \mathcal{S} and finally return a result of type a . A term of type $\text{Ref}_s r$ is a reference in the store s that contains terms of ground type r . References permit both "printing" effects via writing—like channels in $\text{STLC}^{\text{PRINT}}$ —but also reading effects. Note that the type of computations *does not* mention the store locations from which it might read. This asymmetry stems from how the execution of programs performing these effects interact with their environment: writing alters it whilst reading does not—at least directly.

Typing judgements are of the form $\Sigma, \Gamma \vdash t : a$ where Σ is a store typing which determines the shape of the store, i.e., which types it contains and in what locations. The rest of components are like those in STLC . The typing rules of the monadic operations `return`, `bind` and `subeff` follow the same pattern as in $\text{STLC}^{\text{PRINT}}$ (Figure 2), thus we do not discuss them any further. The term `refs` (Rule REF) is the runtime representation of references. Reading and writing is achieved via primitives `read` and `write` (Rules READ and WRITE respectively). Observe that in READ, the type of the computation $\text{Eff}_{\emptyset} r$ in the conclusion of the rule is indexed by the empty set of locations, while in WRITE is indexed by the singleton set $\{s\}$.

The operational semantics of the language is defined as in $\text{STLC}^{\text{PRINT}}$; a combination of the small-step semantics for STLC that treats effectful primitives as values, and a small-step semantics for computations. Given a store typing Σ , a store θ is a function from locations to typed terms according to Σ . Since the language considers a fixed-size store, we use the notation θ instead of $\theta(\Sigma)$. The semantics of computations is of the form $\theta_1, t \rightsquigarrow \theta_2, u$, and is interpreted as: program $\Sigma, \cdot \vdash t : \text{Eff}_{\mathcal{S}} a$ paired with store θ_1 evaluates in one step to program $\Sigma, \cdot \vdash u : \text{Eff}_{\mathcal{S}} a$ and store θ_2 .

We now explain the semantics. Rule READ reduces the argument of `read` and it does not modify the store. When the argument is a store location, i.e., `refs`, READ-REF retrieves the term t from the store, i.e., $\theta(s) = t$. The rules for writing Rules WRITE and WRITE-REF first reduce the left subterm of `write(u, t)` to a store location and then write the right subterm t on the store. Different from $\text{STLC}^{\text{PRINT}}$, we permit to write any term on the store, not only values.

To briefly illustrate $\text{STLC}^{\text{STORE}}$, consider the following program:

$$\begin{aligned} \text{prog}_s &: \text{Bool} \Rightarrow \text{Eff}_{\{s'\}} \text{Unit} \\ \text{prog}_s b &= \text{if } b \text{ then } (\text{read}(s) \gg \lambda(x. \text{write}(s', x))) \\ &\quad \text{else return(unit)} \end{aligned}$$

Based on the Boolean input, prog_s copies the contents of the store location s to s' . Observe that the type only mentions the location s' in its index.

We conclude the explanation of $\text{STLC}^{\text{STORE}}$ with a confinement lemma—similar to that of $\text{STLC}^{\text{PRINT}}$ (Lemma 2):

Store locations	$s \in \mathcal{Loc}$
Sets of s. locations	$S, S_1, S_2 \subseteq \mathcal{Loc}$
Ground types	$r ::= \text{Bool} \mid \text{Unit}$
Types	$a, b ::= r \mid a \Rightarrow b \mid \text{Eff}_s a \mid \text{Ref}_s r$
Typing contexts	$\Gamma ::= \dots$
Store typings	$\Sigma \in \mathcal{Loc} \rightarrow \text{Ground types}$
<div style="border: 1px solid black; display: inline-block; padding: 2px 10px;">$\Sigma, \Gamma \vdash t : a$</div>	
REF	$\frac{\Sigma(s) = r}{\Sigma, \Gamma \vdash \text{ref}_s : \text{Ref}_s r}$
READ	$\frac{}{\Sigma, \Gamma \vdash \text{read}(t) : \text{Eff}_\emptyset r}$
WRITE	$\frac{\Sigma, \Gamma \vdash t : \text{Ref}_s r \quad \Sigma, \Gamma \vdash u : r}{\Sigma, \Gamma \vdash \text{write}(t, u) : \text{Eff}_{\{s\}} \text{Unit}}$
Stores $\theta, \theta_1, \theta_2 \in (\Sigma : \text{Store typings}) \rightarrow (s : \mathcal{Loc}) \rightarrow \Sigma, \cdot \vdash t : \Sigma(s)$	
<div style="border: 1px solid black; display: inline-block; padding: 2px 10px;">$\theta_1(\Sigma), t \rightsquigarrow \theta_2(\Sigma), u$ with $\Sigma, \cdot \vdash t : \text{Eff}_s a$ and $\Sigma, \cdot \vdash u : \text{Eff}_s a$</div>	
READ	$\frac{t \longrightarrow u}{\theta, \text{read}(t) \rightsquigarrow \theta, \text{read}(u)}$
READ-REF	$\frac{\theta(s) = t}{\theta, \text{read}(\text{ref}_s) \rightsquigarrow \theta, \text{return}(t)}$
WRITE	$\frac{t \longrightarrow t'}{\theta, \text{write}(t, u) \rightsquigarrow \theta, \text{write}(t', u)}$
WRITE-REF	$\frac{\theta_2 = \theta_1[s \mapsto t]}{\theta_1, \text{write}(\text{ref}_s, t) \rightsquigarrow \theta_2, \text{return}(\text{unit})}$

Fig. 4. Types, well-typed terms and small-step semantics of $\text{STLC}^{\text{STORE}}$ (excerpts)

Lemma 3 (Confinement for $\text{STLC}^{\text{STORE}}$). *For any $\text{STLC}^{\text{STORE}}$ program $\Sigma, \cdot \vdash f : \text{Eff}_s a$, $\text{STLC}^{\text{STORE}}$ value $\Sigma, \cdot \vdash v : \text{Eff}_s a$, and stores $\theta_1, \theta_2 : \Sigma$, if $\theta_1, f \rightsquigarrow^* \theta_2, v$ then $\forall s \in \mathcal{Loc}. s \notin S \Rightarrow \theta_1(s) = \theta_2(s)$.*

Now its turn to explain the IFC enforcement mechanism SC^{STORE} (Figure 5). As in SC^{PRINT} (Figure 3) we assume that the security policy specifies for each store location its sensitivity as a function $\text{label} \in \mathcal{Loc} \rightarrow \mathcal{L}$. The typing rules for the monadic primitives are analogous to SC^{PRINT} thus we have omitted them. The rule for references is straightforward (Rule REF). More interesting is the typing rule for reading from the store (Rule READ). In the conclusion of the rule, the return type of the computation is the SC type for sensitive data $S_l R$. The sensitivity of the location is l , i.e., $\text{label}(s) = l$ in the premise of the rule, thus to protect the flow of information is necessary to wrap also the return type.

The type of `read` diverges from usual presentations of IFC libraries (e.g., LIO [39], MAC [33] with the exception of SLIO [32]) in that reading from the store is wrapped in the type S_l . These libraries incorporate the data into their monad of computations, which keeps track of the sensitivities of the observed values.

We conclude the section with a concrete example of SC^{STORE} that shows that prog_8 is secure with respect to the following specification: $\text{label}(s) = \text{Alice}$,

<p>Ground types $R ::= \text{Bool} \mid \text{Unit}$ Types $A, B ::= R \mid A \Rightarrow B \mid \text{Eff}_s A \mid \text{Ref}_s R$ Store typings $\Sigma \in \text{Loc} \rightarrow \text{Ground types}$ Sensitivity of loc. $\text{label} \in \text{Loc} \rightarrow \mathcal{L}$</p> <div style="border: 1px solid black; padding: 2px; display: inline-block;"> $\pi ; \Sigma, \Gamma \vdash^{\text{sc}} t : A$ </div>
$\frac{\text{REF} \quad \Sigma(s) = R}{\pi ; \Sigma, \Gamma \vdash^{\text{sc}} \text{ref}_s : \text{Ref}_s R}$
$\frac{\text{READ} \quad \pi ; \Sigma, \Gamma \vdash^{\text{sc}} t : \text{Ref}_s R \quad \text{label}(s) = l}{\pi ; \Sigma, \Gamma \vdash^{\text{sc}} \text{read}(t) : \text{Eff}_\emptyset (S_l R)}$
$\frac{\text{WRITE} \quad \pi ; \Sigma, \Gamma \vdash^{\text{sc}} t : \text{Ref}_s R \quad \pi ; \Sigma, \Gamma \vdash^{\text{sc}} u : R}{\pi ; \Sigma, \Gamma \vdash^{\text{sc}} \text{write}(t, u) : \text{Eff}_{\{s\}} \text{Unit}}$
$\frac{\text{DISTR} \quad \pi ; \Sigma, \Gamma \vdash^{\text{sc}} t : S_l (\text{Eff}_s A) \quad \forall s \in \mathcal{S}. l \sqsubseteq \text{label}(s)}{\pi ; \Gamma \vdash^{\text{sc}} \text{distr}(t) : \text{Eff}_s (S_l A)}$

Fig. 5. Types and intrinsically-typed terms of SC^{STORE} (excerpts)

$\text{label}(s') = \text{secret}$ and the Boolean argument is **secret**, i.e., $sb : S_{\text{secret}} \text{Bool}$:

$$\begin{aligned} \text{prog}'_s &: S_{\text{secret}} \text{Bool} \Rightarrow \text{Eff}_{\{s'\}} (S_{\text{secret}} \text{Unit}) \\ \text{prog}'_s sb &= \text{distr}(\text{seal}_{\text{secret}}(\text{if unseal}_{\text{Alice}}(sb) \\ &\quad \text{then } (\text{read}(s) \gg \lambda(x. \text{write}(s', \text{unseal}_{\text{secret}}(x)))) \\ &\quad \text{else return}(\text{unit}))) \end{aligned}$$

The above program exemplifies how SC^{STORE} enforces that flows from the store to the program and back to the store are secure. Note that $\varepsilon(\text{prog}'_s) = \text{prog}_s$.

4 Security Guarantees

In this section, we prove that the IFC mechanisms SC , SC^{PRINT} and SC^{STORE} can be used to enforce noninterference for the programming languages STLC , $\text{STLC}^{\text{PRINT}}$ and $\text{STLC}^{\text{STORE}}$ presented in Sections 2, 3.1 and 3.2, respectively. Our noninterference proofs employ the technique of logical relations (LR) [23]. For each language, we construct a family of LRs parameterized by the sensitivity of the attacker l_{atk} and the SC types. In the effect-free setting, the LR interprets each SC type A as a binary relation over STLC programs of the erased type $\varepsilon(A)$. The relation captures the idea of indistinguishable programs: if the type is *public* enough, e.g., $S_{\text{public}} \text{Bool}$ and **public** flows to l_{atk} , then two programs are related when they evaluate to the same value. Noninterference then follows as a corollary of the so called fundamental theorem of the LR.

Although the languages we consider do not have primitives for writing general recursive programs, the noninterference properties that we prove, namely termination-insensitive noninterference (TINI), do not assume that programs terminate. A non-terminating loop in a language with printing effects can leak secrets, however, the leakage bandwidth is exponential in the size of the secret [6],

and hence, most IFC tools (e.g., Lifty [31]) ignore such leaks and enforce TINI. From an IFC perspective, it is straightforward to consider general recursion (e.g., as done in Russo, Claessen, and Hughes [34] and Schoepe, Hedin, and Sabelfeld [36]).

TINI states that when two runs of a program with different **secret** inputs terminate, then its **public** outputs agree. In the effect-free setting, the inputs to a program are its arguments, and the output is its return value. In the effectful setting, what we need to consider as inputs and outputs changes: in $\text{STLC}^{\text{STORE}}$, for instance, the store must be considered an additional input to the program.

Definition 1 (TINI for STLC). *An STLC program $sec : \text{Bool} \vdash f : \text{Bool}$ satisfies TINI if for any two STLC terms $\cdot \vdash s_1, s_2 : \text{Bool}$, and any two STLC values $\cdot \vdash v_1, v_2 : \text{Bool}$, if $f[s_1/sec] \longrightarrow^* v_1$ and $f[s_2/sec] \longrightarrow^* v_2$ then $v_1 \equiv v_2$.*

In the definition $v_1 \equiv v_2$ denotes that v_1 and v_2 are syntactically equal values. Note that programs that diverge for any input vacuously satisfy TINI: the assumption that the substituted programs terminate will never hold.

SC can be used to enforce TINI:

Theorem 1. *Given any SC term $\{l_{\text{atk}}\}; sec : S_{l_{\text{sec}}} \text{Bool} \vdash^{\text{sc}} t : S_{l_{\text{atk}}} \text{Bool}$ where $l_{\text{sec}} \not\sqsubseteq l_{\text{atk}}$, the erased program $sec : \text{Bool} \vdash \varepsilon(t) : \text{Bool}$ satisfies TINI.*

In practice, we are concerned that information does not flow from l_{sec} -protected data to the attacker with sensitivity l_{atk} . Thus, to show that an STLC program $sec : \text{Bool} \vdash f : \text{Bool}$ does not leak information from l_{sec} to l_{atk} , it is enough to find an SC term $\{l_{\text{atk}}\}; sec : S_{l_{\text{sec}}} \text{Bool} \vdash^{\text{sc}} t : S_{l_{\text{atk}}} \text{Bool}$ such that $f = \varepsilon(t)$.

Logical relation. In order to prove that SC enforces TINI (Theorem 1), we construct an LR parameterized by the attacker’s sensitivity l_{atk} and the SC types—see Figure 6. The proof, as we will show, then falls out as a consequence of the fundamental theorem of the LR.

At each SC type the LR defines what the attacker can observe about pairs of STLC programs of erased type—alternatively, the same program with different secrets. We split the definition depending on whether programs are evaluated, $\mathcal{R}_{\mathcal{V}}^{l_{\text{atk}}}[-]$ and $\mathcal{R}_{\mathcal{E}}^{l_{\text{atk}}}[-]$ respectively. We briefly explain these. At SC type Bool , for instance, see $\mathcal{R}_{\mathcal{E}}^{l_{\text{atk}}}[\text{Bool}]$ and $\mathcal{R}_{\mathcal{V}}^{l_{\text{atk}}}[\text{Bool}]$, the LR states that if the programs terminate then the attacker can observe if they return equal values. At higher types, i.e., $A \Rightarrow B$, two functions are related if whenever they reduce to a value, see $\mathcal{R}_{\mathcal{E}}^{l_{\text{atk}}}[A \Rightarrow B]$, they map related inputs $\mathcal{R}_{\mathcal{E}}^{l_{\text{atk}}}[A](u_1, u_2)$ to related outputs $\mathcal{R}_{\mathcal{E}}^{l_{\text{atk}}}[A](\text{app}(t_1, u_1), \text{app}(t_2, u_2))$, see $\mathcal{R}_{\mathcal{V}}^{l_{\text{atk}}}[A \Rightarrow B]$. Lastly, at SC type $S_l A$, see $\mathcal{R}_{\mathcal{V}}^{l_{\text{atk}}}[S_l A]$, the LR compares the sensitivity of the attacker with the label l , and in case it is less sensitive, i.e., $l \sqsubseteq l_{\text{atk}}$, the programs have to be related at SC type A . If the label l is more sensitive than the attacker’s label, i.e., $l \not\sqsubseteq l_{\text{atk}}$ then the programs do not need to be related.

Definitions $\mathcal{R}_{\mathcal{V}}^{l_{\text{atk}}}[-]$ and $\mathcal{R}_{\mathcal{E}}^{l_{\text{atk}}}[-]$ work on closed terms, however, in order to prove the fundamental theorem we have to lift them to substitutions. A substitution assigns to each type a in a typing context Γ a closed term of

that type, i.e., $\cdot \vdash f : a$. We denote substitutions by γ and use $\gamma : \cdot \vdash \Gamma$ to mean that γ is in the set of substitutions over Γ . We define the LR for substitutions, $\mathcal{R}_S^{l_{\text{atk}}}\llbracket - \rrbracket$, by induction on SC typing contexts. At the empty context \cdot the empty substitutions (ϵ, ϵ) are trivially related—denoted by \top . Two non-empty substitutions (γ_1, t_1) and (γ_2, t_2) are related whenever they are pointwise related, i.e., $\mathcal{R}_E^{l_{\text{atk}}}\llbracket A \rrbracket(t_1, t_2)$ and $\mathcal{R}_S^{l_{\text{atk}}}\llbracket \Gamma \rrbracket(\gamma_1, \gamma_2)$. The LR for open terms, written $\mathcal{R}_T^{l_{\text{atk}}}\llbracket (\Gamma, A) \rrbracket$, is indexed by a pair consisting of an SC typing context Γ and an SC type A . It states that two STLC terms are related if for any two related substitutions the substituted terms are related at type A .

$$\begin{aligned}
& \mathcal{R}_V^{l_{\text{atk}}}\llbracket - \rrbracket \in (A : \text{SC type}) \rightarrow (t_1 : \cdot \vdash \varepsilon(A)) \rightarrow (t_2 : \cdot \vdash \varepsilon(A)) \rightarrow \text{Set} \\
& \mathcal{R}_V^{l_{\text{atk}}}\llbracket \text{Unit} \rrbracket(t_1, t_2) :\Leftrightarrow t_1 \equiv t_2 \\
& \mathcal{R}_V^{l_{\text{atk}}}\llbracket \text{Bool} \rrbracket(t_1, t_2) :\Leftrightarrow t_1 \equiv t_2 \\
& \mathcal{R}_V^{l_{\text{atk}}}\llbracket A \Rightarrow B \rrbracket(t_1, t_2) :\Leftrightarrow \\
& \quad \forall (u_1, u_2 : \cdot \vdash \varepsilon(A)). \mathcal{R}_E^{l_{\text{atk}}}\llbracket A \rrbracket(u_1, u_2) \Rightarrow \mathcal{R}_E^{l_{\text{atk}}}\llbracket B \rrbracket(\text{app}(t_1, u_1), \text{app}(t_2, u_2)) \\
& \mathcal{R}_V^{l_{\text{atk}}}\llbracket S_i A \rrbracket(t_1, t_2) :\Leftrightarrow l \sqsubseteq l_{\text{atk}} \Rightarrow \mathcal{R}_V^{l_{\text{atk}}}\llbracket A \rrbracket(t_1, t_2) \\
& \mathcal{R}_E^{l_{\text{atk}}}\llbracket A \rrbracket(t_1, t_2) :\Leftrightarrow \forall (u_1, u_2 : \cdot \vdash \varepsilon(A)). t_1 \longrightarrow^* u_1 \wedge t_2 \longrightarrow^* u_2 \Rightarrow \mathcal{R}_V^{l_{\text{atk}}}\llbracket A \rrbracket(u_1, u_2) \\
& \mathcal{R}_S^{l_{\text{atk}}}\llbracket \cdot \rrbracket(\epsilon, \epsilon) :\Leftrightarrow \top \\
& \mathcal{R}_S^{l_{\text{atk}}}\llbracket \Gamma, x : A \rrbracket((\gamma_1, t_1), (\gamma_2, t_2)) :\Leftrightarrow \mathcal{R}_E^{l_{\text{atk}}}\llbracket A \rrbracket(t_1, t_2) \wedge \mathcal{R}_S^{l_{\text{atk}}}\llbracket \Gamma \rrbracket(\gamma_1, \gamma_2) \\
& \mathcal{R}_T^{l_{\text{atk}}}\llbracket (\Gamma, A) \rrbracket(t_1, t_2) :\Leftrightarrow \\
& \quad \forall (\gamma_1, \gamma_2 : \cdot \vdash \varepsilon(\Gamma)). \mathcal{R}_S^{l_{\text{atk}}}\llbracket \Gamma \rrbracket(\gamma_1, \gamma_2) \Rightarrow \mathcal{R}_E^{l_{\text{atk}}}\llbracket A \rrbracket(t_1[\gamma_1], t_2[\gamma_2])
\end{aligned}$$

Fig. 6. Logical relation for STLC-SC

The fundamental theorem of the LR states that the underlying program of an SC term is related to itself. Formally:

Theorem 2 (Fundamental Theorem of the LR for STLC). *For any attacker with sensitivity l_{atk} , and SC term $\{l_{\text{atk}}\}; \Gamma \vdash^{\text{sc}} t : A$, it is the case that $\mathcal{R}_T^{l_{\text{atk}}}\llbracket (\Gamma, A) \rrbracket(\varepsilon(t), \varepsilon(t))$.*

Proof. By induction on the typing derivation.

TINI (Theorem 1) follows as a corollary of the fundamental theorem:

Proof. Assume two secrets $\cdot \vdash s_1, s_2 : \text{Bool}$. Since $l_{\text{sec}} \not\sqsubseteq l_{\text{atk}}$, the secrets are related, i.e., $\mathcal{R}_E^{l_{\text{atk}}}\llbracket S_{l_{\text{sec}}}\text{Bool} \rrbracket(s_1, s_2)$, and thus the two substitutions $\gamma_1 = \{\text{sec} \mapsto$

$s_1\}$ and $\gamma_2 = \{sec \mapsto s_2\}$ are related. By the fundamental theorem, the term $\varepsilon(t)$ is related to itself, i.e., $\mathcal{R}_{\varepsilon}^{l_{\text{atk}}} \llbracket S_{l_{\text{atk}}} \text{Bool} \rrbracket (\varepsilon(t)[\gamma_1], \varepsilon(t)[\gamma_2])$. Unfolding the definitions of we obtain that $\forall (v_1, v_2 : \cdot \vdash \text{Bool}). \varepsilon(t)[\gamma_1] \longrightarrow^* v_1 \wedge \varepsilon(t)[\gamma_2] \longrightarrow^* v_2 \Rightarrow v_1 \equiv v_2$.

4.1 Noninterference for Printing Effects: $\text{STLC}^{\text{Print}}$ and SC^{Print}

In order to formalize noninterference for $\text{STLC}^{\text{Print}}$ we look at effectful programs that might print. For that, we first define indistinguishability of outputs with respect to a subset of channels:

Definition 2 (Output indistinguishability). *Let $C \subseteq \text{Ch}$. Two outputs o_1 and o_2 are C -indistinguishable, denoted by $o_1 =_C o_2$, if the two outputs agree in C , i.e., $\forall ch \in C. o_1(ch) = o_2(ch)$.*

We define TINI for $\text{STLC}^{\text{Print}}$ for programs from Bool to $\text{Eff}_C \text{Unit}$, i.e., programs that depending on a Boolean produce output in an arbitrary set of channels C :

Definition 3 (TINI for $\text{STLC}^{\text{Print}}$). *An $\text{STLC}^{\text{Print}}$ program $sec : \text{Bool} \vdash f : \text{Eff}_C \text{Unit}$ satisfies TINI with respect to $C' \subseteq C$, if for any two $\text{STLC}^{\text{Print}}$ terms $\cdot \vdash s_1, s_2 : \text{Bool}$, any two $\text{STLC}^{\text{Print}}$ values $\cdot \vdash v_1, v_2 : \text{Eff}_C \text{Unit}$, and any two outputs o_1 and o_2 , if $f[s_1/sec] \rightsquigarrow^* v_1, o_1$ and $f[s_2/sec] \rightsquigarrow^* v_2, o_2$ then $o_1 =_{C'} o_2$.*

TINI is parameterized by a subset of the channels where the the outputs have to agree. We will instantiate C' with the set of channels observable by the attacker.

SC^{Print} can be used to enforce TINI on $\text{STLC}^{\text{Print}}$ programs:

Theorem 3. *Given any SC^{Print} term $\{l_{\text{atk}}\}$; $sec : S_{l_{\text{sec}}} \text{Bool} \vdash^{\text{sc}} t : \text{Eff}_C \text{Unit}$ where $l_{\text{sec}} \not\sqsubseteq l_{\text{atk}}$ the erased program $sec : \text{Bool} \vdash \varepsilon(t) : \text{Eff}_C \text{Unit}$ satisfies TINI with respect to C_{atk} where $C_{\text{atk}} := \{ch \mid ch \in C, \text{label}(ch) \sqsubseteq l_{\text{atk}}\}$.*

$$\begin{aligned} \mathcal{R}_{\mathcal{V}}^{l_{\text{atk}}} \llbracket \dots \rrbracket (t_1, t_2) &:\Leftrightarrow \dots \\ \mathcal{R}_{\mathcal{V}}^{l_{\text{atk}}} \llbracket \text{Eff}_S A \rrbracket (t_1, t_2) &:\Leftrightarrow \forall (u_1, u_2 : \cdot \vdash \varepsilon(A)), o_1, o_2. t_1 \rightsquigarrow^* \text{return}(u_1), o_1 \\ &\quad \wedge t_2 \rightsquigarrow^* \text{return}(u_2), o_2 \Rightarrow \mathcal{R}_{\varepsilon}^{l_{\text{atk}}} \llbracket A \rrbracket (u_1, u_2) \wedge o_1 =_{C_{\text{atk}}} o_2 \end{aligned}$$

Fig. 7. Logical relation for $\text{STLC}^{\text{Print}}$ (excerpts)

Logical relation. The LR for $\text{STLC}^{\text{Print}}$ (Figure 7) is very similar to that of STLC (Figure 6) so we skip over the commonalities and directly discuss its definition at the type of computations. The LR relates two computations $\mathcal{R}_{\mathcal{V}}^{l_{\text{atk}}} \llbracket \text{Eff}_C A \rrbracket (t_1, t_2)$ if whenever they terminate, i.e., $t_1 \rightsquigarrow^* \text{return}(u_1), o_1$ and $t_2 \rightsquigarrow^* \text{return}(u_2), o_2$, the resulting terms are related, i.e., $\mathcal{R}_{\varepsilon}^{l_{\text{atk}}} \llbracket \tau \rrbracket (u_1, u_2)$, and the outputs are indistinguishable by the attacker, i.e., $o_1 =_{C_{\text{atk}}} o_2$. The fundamental theorem of the LR states that erased terms are related to themselves:

Theorem 4 (Fundamental Theorem of the LR for $\text{STLC}^{\text{Print}}$). *For any attacker with sensitivity l_{atk} , and SC^{Print} term $\{l_{\text{atk}}\}; \Gamma \vdash^{\text{sc}} t : A$, it is the case that $\mathcal{R}_{\mathcal{T}}^{l_{\text{atk}}} \llbracket (\Gamma, A) \rrbracket (\varepsilon(t), \varepsilon(t))$.*

Proof. By induction on the typing derivation with use of the Lemma 2.

The proof that SC^{Print} enforces TINI (Theorem 3) follows as a corollary:

Proof. Let $\cdot \vdash s_1, s_2 : \text{Bool}$ be two secrets. Since $l_{\text{sec}} \not\sqsubseteq l_{\text{atk}}$, the secrets are related, i.e., $\mathcal{R}_{\mathcal{E}}^{l_{\text{atk}}} \llbracket \mathcal{S}_{l_{\text{sec}}} \text{Bool} \rrbracket (s_1, s_2)$, and thus the substitutions $\gamma_1 = \{sec \mapsto s_1\}$ and $\gamma_2 = \{sec \mapsto s_2\}$ are related. By the fundamental theorem, the term $\varepsilon(t)$ is related to itself $\mathcal{R}_{\mathcal{E}}^{l_{\text{atk}}} \llbracket \text{Eff}_C \text{Unit} \rrbracket (\varepsilon(t)[\gamma_1], \varepsilon(t)[\gamma_2])$. By assumption, $\varepsilon(t)[s_1/sec] \rightsquigarrow^* v_1, o_1$ and $\varepsilon(t)[s_2/sec] \rightsquigarrow^* v_2', o_2$. By reasoning about the operational semantics, there exists two intermediate programs, t_1', t_2' such that: $\varepsilon(t)[s_1/sec] \longrightarrow^* t_1'$ and $t_1' \rightsquigarrow^* v_1, o_1$; and $\varepsilon(t)[s_2/sec] \longrightarrow^* t_2'$ and $t_2' \rightsquigarrow^* v_2, o_2$. We apply $\mathcal{R}_{\mathcal{E}}^{l_{\text{atk}}} \llbracket \text{Eff}_C \text{Unit} \rrbracket (\varepsilon(t)[s_1/sec], \varepsilon(t)[s_2/sec])$ to the two effect-free reductions which gives us that $\mathcal{R}_{\mathcal{V}}^{l_{\text{atk}}} \llbracket \text{Eff}_C \text{Unit} \rrbracket (t_1', t_2')$. We apply this to the monadic reductions and obtain that $o_1 =_{c_{\text{atk}}} o_2$.

4.2 Noninterference for Global Store Effects: $\text{STLC}^{\text{Store}}$ and SC^{Store}

We formalize noninterference for $\text{STLC}^{\text{Store}}$ by looking at effectful programs which receive a store as input and produce a store as output. The contents of the store are possibly unevaluated $\text{STLC}^{\text{Store}}$ programs of ground type (see Figure 4). In order to compare stores, we define an indistinguishability relation for programs. This is a sort of LR parameterized by $\text{STLC}^{\text{Store}}$ ground types:

$$\begin{aligned} \mathcal{I}[\!-\!] &: (r : \text{STLC}^{\text{Store}} \text{ ground type}) \rightarrow (t_1 : \cdot \vdash r) \rightarrow (t_2 : \cdot \vdash r) \rightarrow \text{Set} \\ \mathcal{I}[\text{Unit}] &: (t_1, t_2) :\Leftrightarrow \forall (v_1, v_2 : \cdot \vdash \text{Unit}). t_1 \longrightarrow^* v_1 \wedge t_2 \longrightarrow^* v_2 \Rightarrow v_1 \equiv v_2 \\ \mathcal{I}[\text{Bool}] &: (t_1, t_2) :\Leftrightarrow \forall (v_1, v_2 : \cdot \vdash \text{Bool}). t_1 \longrightarrow^* v_1 \wedge t_2 \longrightarrow^* v_2 \Rightarrow v_1 \equiv v_2 \end{aligned}$$

In some sense it resembles the LRs for Unit and Bool types in Figure 6.

Stores are parameterized by store typings that determine the type of the contents at each location. Since stores neither grow nor shrink, assumption 1 (Section 3.2), we define an indistinguishability relation for stores of the same store typing. Indistinguishability is parameterized by a subset of the locations.

Definition 4 (Store indistinguishability). *Let $S \subseteq \text{Loc}$. Two stores θ_1 and θ_2 of store typing a are S -indistinguishable, denoted by $\theta_1 =_S \theta_2$, if the two stores are indistinguishable at each location in S , i.e., $\forall s \in S. \mathcal{I}[\Sigma(s)](\theta_1(s), \theta_2(s))$.*

TINI for $\text{STLC}^{\text{Store}}$ programs is:

Definition 5 (TINI for $\text{STLC}^{\text{Store}}$). *An $\text{STLC}^{\text{Store}}$ program $\Sigma, \cdot \vdash f : \text{Eff}_S \text{Unit}$ satisfies TINI with respect to $S' \subseteq \text{Loc}$, if for any two stores $\theta_1, \theta_2 : \Sigma$, any two $\text{STLC}^{\text{Store}}$ values $\cdot \vdash v_1, v_2 : \text{Eff}_S \text{Unit}$, and any two stores $\theta_1', \theta_2' : \Sigma$, if $\theta_1 =_S \theta_2$ and $\theta_1, f \rightsquigarrow^* \theta_1', v_1$ and $\theta_2, f \rightsquigarrow^* \theta_2', v_2$ then $\theta_1' =_S \theta_2'$.*

Again, SC^{Store} enforces TINI on $\text{STLC}^{\text{Store}}$ programs:

$$\begin{aligned}
 \mathcal{R}_V^{l_{\text{atk}}} \llbracket \dots \rrbracket (t_1, t_2) &:\Leftrightarrow \dots \\
 \mathcal{R}_V^{l_{\text{atk}}} \llbracket \text{Eff}_c A \rrbracket (t_1, t_2) &:\Leftrightarrow \forall (u_1, u_2 : \cdot \vdash \varepsilon(A)) (\theta_1, \theta_2, \theta'_1, \theta'_2 : \Sigma). \theta_1 =_{\mathcal{S}_{\text{atk}}} \theta_2 \wedge \\
 &\theta_1, t_1 \rightsquigarrow^* \theta'_1, \text{return}(u_1) \wedge \theta_2, t_2 \rightsquigarrow^* \theta'_2, \text{return}(u_2) \Rightarrow \mathcal{R}_E^{l_{\text{atk}}} \llbracket A \rrbracket (u_1, u_2) \wedge \theta'_1 =_{\mathcal{S}_{\text{atk}}} \theta'_2
 \end{aligned}$$

Fig. 8. Logical relation for $\text{STLC}^{\text{STORE}}$ (excerpts)

Theorem 5. *Given any SC^{STORE} term $\{l_{\text{atk}}\}$; $\Sigma, \cdot \vdash^{\text{sc}} f : \text{Eff}_S \text{Unit}$ the erased program $\Sigma, \cdot \vdash \varepsilon(f) : \text{Eff}_S \text{Unit}$ satisfies TINI with respect to \mathcal{S}_{atk} where $\mathcal{S}_{\text{atk}} := \{s \mid s \in \text{Loc}, \text{label}(s) \sqsubseteq l_{\text{atk}}\}$.*

The LR that we construct to prove TINI (Figure 8) is largely similar to that of $\text{STLC}^{\text{PRINT}}$ (Figures 6 and 7), with the difference that effectful programs take as argument and produce as result indistinguishable pairs of stores. TINI follows as a consequence of the fundamental theorem of the LR.

5 Implementation

In this section, we present an implementation of SC and SC^{PRINT} (Sections 2 and 3.1) as a Haskell library, which we call SCLib. We omit SC^{STORE} (Section 3.2) for lack of space. However, its implementation is similar to that of SC^{PRINT} . Furthermore, we demonstrate that existing Haskell libraries for static IFC can be reimplemented in terms of the interface that SCLib exposes.

The main characteristic of SC is that typing judgements $\pi ; \Gamma \vdash^{\text{sc}} t : A$ are indexed by a set of labels π . Onwards, we refer to the left part the judgement, i.e., $\pi ; \Gamma$, as the context of the term t . The set of labels plays an important role in enforcing IFC. However, individual labels are not first-class citizens: there is no type of labels and, hence, labels can neither be introduced nor eliminated. Further, some typing rules in SC modify the set of labels in their context: e.g., the rule for unseal_l (cf. Figure 1) augments the set in its premise with l . When shallowly embedding in Haskell any calculus that manipulates the context in this fashion, there is a natural problem to overcome: Haskell does not allow library implementors to have access to a program’s context. For instance, to embed a linear type system in Haskell, Bernardy et al. [7] need to change the compiler.

To overcome these difficulties, our implementation resorts to a combination of Haskell’s module system to hide the implementation details’ from users, and a well-known trick of higher-rank polymorphism [20] to prevent first-class access to labels. We hope to convince the reader that our simple implementation matches the studied enforcement mechanisms, and that it can shed light on previous work on static IFC in Haskell.

5.1 Implementation of SC

When we introduced SC, we mentioned the intuition that labels in π are some sort of type-level keys whose possession permits access to sealed data. Our implementation takes this intuition literally: there is a type for keys whose elements are attached with type-level labels, and the primitive to unseal, i.e., unseal , is

parameterized by a key. The elements of this type are like capabilities [17] which need to be explicitly exercised.

```

1  module SCLib
2    (Key (), Label (..), FlowsTo (..), S (S), seal, unseal, ...)
3  where
4    -- Enumeration of security labels for the two-point lattice
5    data Label = H | L
6    -- "Flows to" relation as a type-class
7    class FlowsTo (l :: Label) (l' :: Label)
8    -- Instances
9    instance FlowsTo l l
10   instance FlowsTo L H
11   -- Type-level keys
12   data Key l s = Key
13   -- Security type
14   data S l a = S (forall s. Key l s -> a)
15   -- Sealing
16   seal :: (forall s. Key l s -> a) -> S l a
17   seal = Seal
18   -- Unsealing
19   unseal :: FlowsTo l' l => Key l s -> S l' a -> a
20   unseal k@Key (S f) = f Key

```

Fig. 9. Implementation of SC (e.g., for the two-point security lattice)

Figure 9 shows the *complete* implementation of SC in Haskell. Without loss of generality, we assume the two point-security lattice. As previous work (e.g., MAC [33], HLIO [9], and DCC [5]) we represent labels as types of kind `Label` (line 5, and the use of the GHC extension `DataKinds`) and encode the “flows to” relation via a type-class (lines 7–10). For simplicity we show the encoding of the two-point security lattice, however, this can be generalized (cf. [9]). In line 12 we introduce a new datatype `Key` which is parameterized by a type `l` of kind `Label` and a phantom type `s`. Then, line 14 introduces the type `S`, which is a wrapper over the function space between the types `Key l s` and `a`, i.e., `forall s. Key l s -> a`. The constructor of `S l a`, i.e., `S`, uses higher-rank polymorphism, as evidenced by the `forall` keyword, to quantify over the type variable `s` in `Key l s`. This ensures that the type of any newly bound key will be unique: the type `s` will not coincide with any other type `s'`. Launchbury and Jones [20] use the same trick, for example, to provide pure mutable references in Haskell where the type variable represents memory regions. Since the type variable `s` does not coincide with any other, elements of the type `Key l s` are not first-class values, and, for example, cannot be stored: the program `S (\k -> k)` does not type check.

We now implement the primitives `seal` and `unseal`—Rules `SEAL` and `UNSEAL` from Figure 1. The combinator `seal` (lines 16 and 17) is a wrapper over the

constructor of the type `S l a`. The combinator `unseal` (lines 19–20) allows unsealing terms of type `S l a` in case we *have* a value of type `Key l' s` and the label in its type, i.e., `l'`, flows to `l`—see the constraint `FlowsTo l' l`. In order to enforce IFC, it is important that the constructors of `Key` are kept abstract from the user—observe `Key ()` in the export list of the module (line 2). Otherwise, anyone could extract the underlying term of type `a` from an `l`-sensitive value `secret :: S l a` by applying `unseal Key`. Similar to Russo, Claessen, and Hughes [34], the combinator `unseal` is strict in its argument (`k@Key`) in order to forbid forged keys like `undefined :: Key l s`. We remark that the noninterference property—recall TINI from Definitions 1 and 3—rules out programs that force `undefined` and halt with error.

The implementation discussed so far consists of the *trusted computing base* (TCB) of SCLib. From now on, users of the library can derive functionalities from the library’s interface. For example, programmers can implement the `Functor`, `Applicative` and `Monad` instances for the type `S l a` as follows:

```
instance Functor (S l) where
  fmap f x = seal (\k -> f (unseal k x))

instance Applicative (S l) where
  pure x = seal (\k -> x)
  f <*> a = seal (\k -> (unseal k f) (unseal k a))

instance Monad (S l) where
  return = pure
  m >>= f = seal (\k -> unseal k (f (unseal k m)))
```

Note that the programmer does not need access to the TBC in order to implement these instances—as opposed to MAC (cf. [41]). This phenomenon, we believe, is a sign of the simplicity and generality of our implementation.

5.2 Implementation of SC^{Print}

Figure 10 shows the implementation of SC^{PRINT} which builds on the implementation of SC (Figure 9). The datatype `Eff` wraps `IO` computations and is indexed by a type-level set [28] `ls` of channels where the computation can write to. For simplicity, we will omit the map label (Figure 3) from channels to labels and identify channels with labels so that the index `ls` has kind `[Label]`. This makes the implementation simpler.

In lines 15–18 we implement the return and bind of the graded monad. `returnEff` does not produce effects, thus, its type is indexed by the empty set `[]` (cf. Figure 3). `bindEff` type is indexed by the union of the sets of labels of the computations (cf. Figure 3). The implementation of subeffecting is the identity function (lines 20–21). Printing effects can be performed by the combinator `printEff` (lines 26–30). The argument of type `SLabel l` is a term level representation of a the type-level label of printing channel.

Lines 36–38 show the implementation of the novel primitive `distr`. The type constraint `FlowsToSet l ls` (see the definitions in lines 32–35) ensures that

```

1  module SCLib
2      (... , Eff (), FlowsToSet (), pureEff, appEff, returnEff
3      , bindEff, distr, subeff, printEff)
4  where
5
6  newtype Eff (ls :: [Label]) a = Eff { runEff :: IO a }
7
8  -- Functor
9  instance Functor (Eff ls) where
10     fmap f (Eff io) = Eff (fmap f io)
11
12  -- Applicative
13  pureEff :: a -> Eff [] a
14  pureEff = returnEff
15
16  appEff :: Eff ls1 (a -> b) -> Eff ls2 a -> Eff (Union ls1 ls2) b
17  appEff (Eff ioff) (Eff ioa) = Eff (ioff <*> ioa)
18
19  -- Graded monad
20  returnEff :: a -> Eff [] a
21  returnEff a = Eff (return a)
22
23  bindEff :: Eff ls1 a -> (a -> Eff ls2 b) -> Eff (Union ls1 ls2) b
24  bindEff (Eff m) f = Eff (m >>= runEff . f)
25
26  -- Subeffecting
27  subeff :: Subset ls1 ls2 => Eff ls1 a -> Eff ls2 a
28  subeff (Eff m) = Eff m
29
30  -- Print
31  data SLabel :: Label -> * where
32      SH :: SLabel H
33      SL :: SLabel L
34
35  printEff :: (Show a) => SLabel l -> a -> Eff [l] ()
36  printEff l x = Eff (print (header l) >> print x)
37     where header :: SLabel l -> String
38           header SH = "Channel H:"
39           header SL = "Channel L:"
40
41  -- Distr
42  type family FlowsToSet (l :: Label) (ls :: [Label]) :: Constraint
43  where
44      FlowsToSet l1 [] = ()
45      FlowsToSet l1 (l2 : ls) = (FlowsTo l1 l2, FlowsToSet l1 ls)
46
47  distr :: FlowsToSet l ls => S l (Eff ls a) -> Eff ls (S l a)
48  distr m = Eff (do res <- (runEff (unseal Key m))
49                return (seal (\k -> res)))

```

Fig. 10. Implementation of SC^{PRINT} (excerpts)

$l \sqsubseteq l'$ for every label l' in ls . The implementation uses the value `Key`, which pertains to the TCB, to unseal the `Eff` action, i.e., `unseal Key m`; and then it runs its effects, i.e., `res <- runEff (unseal Key m)`; finally it seals the result at the same label, i.e., `return (seal (\k -> res))`.

5.3 Implementing existing libraries for IFC

We conclude this section by showing that we can reimplement some of the existing libraries in Haskell for IFC. We show implementations of `SecLib` [34], `SDCC` [4] (an alternative presentation of `DCC`) and a variation of `MAC` [33]⁵ using `SCLib` interface. In some sense the implementations help to explain the mentioned libraries. Further, this shows that the programmer can choose to write programs against `SCLib`'s “low-level” interface; or a more “high-level” interface, e.g., `MAC`; or a combination of both. We declare future work to compare the performance among implementations.

For each library we briefly explain its interface and show its implementation in terms of `SCLib`.

SecLib and SDCC. `SecLib` is one of the pioneers of static IFC libraries in the context of Haskell. Its main feature is a family of security monads `Sec` indexed by labels from the security lattice, each equipped with `>>=` (bind) and `return`. `SecLib`'s special ingredient is a combinator `up` that allows coercions from lower to higher labels in the security monad.

`SDCC` is an alternative presentation of `DCC` which favors a simple set of combinators instead of `DCC`'s non-standard bind and *protected at* relation. Similar to `DCC`, `SDCC` sports a family of monads indexed by security labels. These support `fmap`, `return` and `>>=` (bind). Further, `SDCC` implements two combinators `up` and `com`, that allow to relabel in the style of `SecLib` and commute terms of monadic type with different labels. `SDCC`'s interface is strictly a superset of that of `SecLib`, thus we directly show the implementation of the former. The left column displays the interface whilst the right its implementation in `SCLib`.

<pre> type T l a instance Functor (Sec l) where instance Monad (Sec l) where up :: FlowsTo l l' -> T l a -> T l' a com :: T l (T l' a) -> T l' (T l a) </pre>	<pre> type T l a = S l a up :: FlowsTo l l' => T l a -> T l' a up lv = seal (\k -> unseal k lv) com :: T l (T l' a) -> T l' (T l a) com lv = seal (\k' -> seal (\k -> unseal k' (unseal k lv))) </pre>
--	---

MAC which we discussed in the introduction, is one of the state-of-the-art libraries for effectful IFC in Haskell. At its core, `MAC` defines two types: `Labeled l a` for pure sensitive values, and `MAC l a` for secure computations. `MAC l a` is a monad for each label l where the label: l . protects the data in

⁵ In our variation the type `Labeled` is a monad. This is “unsafe” in `MAC` (cf. [42]).

context; and 2. restricts the permitted effects. (cf. [43]) MAC’s functionality stems from the interaction between `Labeled` and `MAC` through the primitives `label` and `unlabel`. In order to label a value one needs to do so within the `MAC l` monad: the `Labeled l a` type does not export a combinator `a -> Labeled l a`. Our implementation, however, permits to do so. The left column shows MAC’s interface, with the combinator `join` in its full generality, and the right column hints its implementation in terms of SCLib.

<pre> type Labeled l a type MAC l a label :: FlowsTo l l' => a -> MAC l (Labeled l' a) unlabel :: FlowsTo l l' => Labeled l a -> MAC l' a join :: FlowsTo l l' => MAC l' a -> MAC l (Labeled l' a) </pre>	<pre> type Labeled l a = S l a type MAC l a = forall ls. FlowsToSet l ls => Eff ls (S l a) label :: FlowsTo l l' => ... label a = returnEff (seal (\k -> a)) unlabel :: FlowsTo l l' => ... unlabel lv = returnEff lv join :: FlowsTo l l' => ... join m = bindEff m (\x -> seal (\k -> x)) </pre>
---	---

6 Related work

IFC for effect-free and effectful languages. Algehed and Russo [5] add effects to their embedding of DCC in Haskell but argue that their approach only works for those effects that can be implemented within Haskell. Hirsch and Cecchetti [15] develop a formal framework based on productors and type-and-effect systems to characterize secure programs in impure languages with IFC. They give semantics to traditional security-type systems based on controlling implicit flows using PC labels. In contrast, our approach considers from starters a pure language, where the type of effectful computation is separated from that of effect-free programs.

Modalities for IFC. The languages and IFC enforcement mechanisms that we present are based on the Sealing Calculus (SC) of Shikuma and Igarashi [37]. Differently from them, we think of SC terms as evidence that STLC programs satisfy noninterference. The work by Miyamoto and Igarashi [24] gives an informal connection between a classical type system for IFC and a certain modal logic. Their type system is very different from our enforcement mechanism in that a typing judgement has two separate variable contexts. Recently, the work by Abel and Bernardy [3] presents a unified treatment of modalities in typed lambda calculi. The authors present a effect-free lambda calculus parameterized by family of modalities with certain mathematical structure, and show that many PL analyses, including IFC, are instantiations of their framework. In contrast to our work, it is not very clear how one would implement theirs system in Haskell, since it would require a fine-grained control over the variables in the context.

Kavvos [19] studies modalities for IFC in the classified sets model, which they use to prove noninterference properties for a range of calculi that includes SC.

Coeffectful type systems for IFC. A recent line of work suggests using coeffect type systems to enforce IFC. Petricek, Orchard, and Mycroft [30] develop a calculus to capture different granularity demands on contexts, i.e., flat whole-context coeffects (like implicit parameters [21]) or structural per-variable ones (like usage or data access patterns). The work by Gaboardi et al. [10] expands on that and uses graded monads and comonads to combine effects and coeffects. The authors describe *distributivity laws* that are similar to our primitive `distr` addresses. The article suggests IFC as an application where the coeffect system captures the IFC constraints and the effect system gives semantics to effects. The distributive laws explains how both are combined. However, their work does not state neither proves a security property for their calculus. Different from it, our work does not use comonads as the underlying structure for IFC, and further considers printing and global store effects. Granule is a recent programming language [27] based on *graded modal types* that impose usage constraints on the variables.

Logical relations for noninterference. Both Heintze and Riecke [14] and Zdancewic [46] use logical relation arguments to prove noninterference for a simply-typed security lambda calculus. Tse and Zdancewic [40] use logical relations to prove soundness of a translation from DCC [1] to System F and obtain noninterference from parametricity. Unfortunately, their translation is unsound (cf. [37]). Bowman and Ahmed [8] fix this by using “open” logical relations show their translation from DCC to System F_ω is sound. Different from the cited work so far, Rajani and Garg [32] use logical relations to prove noninterference for a language with references. Gregersen et al. [12] extend the use of logical relation to prove noninterference for languages with inpredicative polymorphism. Different from Rajani and Garg [32] and Gregersen et al. [12], we consider first-order references for simplicity. Otherwise, we should have had to utilize a step-indexed Kripke-style logical-relations model, which would have introduced technical complications that are orthogonal to the main contribution of our work.

7 Conclusions

In this paper, we have demonstrated that to enforce IFC in pure languages with a single primitive `distr` suffices to securely control what information flows from sensitive data to effects. To support our claim, we have presented IFC enforcement mechanisms for several kinds of effects and proved that they satisfy noninterference. Our development rests on the insight that effect-free IFC for pure languages can already express that a computation will not leak sensitive data through its effects when executed. Then, a single primitive, `distr`, to execute these is enough to extend IFC to effects and retain the security guarantees. We hope that this work brings a new perspective to IFC research for pure languages with effects.

A STLC

Types $a, b ::= \mathbf{Unit} \mid \mathbf{Bool} \mid a \Rightarrow b$
 Typing contexts $\Gamma ::= \cdot \mid \Gamma, x : a$

$$\boxed{\Gamma \vdash t : a}$$

$$\begin{array}{c}
 \text{VAR} \\
 \frac{(x : a) \in \Gamma}{\Gamma \vdash x : a} \\
 \\
 \text{LAM} \\
 \frac{\Gamma, x : a \vdash t : b}{\Gamma \vdash \lambda(x.t) : a \Rightarrow b} \\
 \\
 \text{APP} \\
 \frac{\Gamma \vdash t : a \Rightarrow b \quad \Gamma \vdash u : a}{\Gamma \vdash \mathbf{app}(t, u) : b} \\
 \\
 \text{UNIT} \\
 \frac{}{\Gamma \vdash \mathbf{unit} : \mathbf{Unit}} \\
 \\
 \text{TRUE} \\
 \frac{}{\Gamma \vdash \mathbf{true} : \mathbf{Bool}} \\
 \\
 \text{FALSE} \\
 \frac{}{\Gamma \vdash \mathbf{false} : \mathbf{Bool}} \\
 \\
 \text{IF} \\
 \frac{\Gamma \vdash t : \mathbf{Bool} \quad \Gamma \vdash u_1 : a \quad \Gamma \vdash u_2 : a}{\Gamma \vdash \mathbf{ifte}(t, u_1, u_2) : a}
 \end{array}$$

$$\boxed{t \rightarrow u \text{ with } \cdot \vdash t : a \text{ and } \cdot \vdash u : a}$$

$$\begin{array}{c}
 \text{APP} \\
 \frac{t \rightarrow t'}{\mathbf{app}(t, u) \rightarrow \mathbf{app}(t', u)} \\
 \\
 \text{BETA} \\
 \frac{}{\mathbf{app}(\lambda(x.t), u) \rightarrow t[u/x]} \\
 \\
 \text{IF} \\
 \frac{t \rightarrow t'}{\mathbf{ifte}(t, u_1, u_2) \rightarrow \mathbf{ifte}(t', u_1, u_2)} \\
 \\
 \text{IF-TRUE} \\
 \frac{}{\mathbf{ifte}(\mathbf{true}, u_1, u_2) \rightarrow u_1} \\
 \\
 \text{IF-FALSE} \\
 \frac{}{\mathbf{ifte}(\mathbf{false}, u_1, u_2) \rightarrow u_2}
 \end{array}$$

Fig. 11. Types, intrinsically-typed terms and small-step semantics of (call-by-name) STLC.

References

1. Abadi, M., Banerjee, A., Heintze, N., and Riecke, J.G.: A Core Calculus of Dependency. In: Appel, A.W., and Aiken, A. (eds.) POPL '99, Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Antonio, TX, USA, January 20-22, 1999, pp. 147–160. ACM (1999). DOI: 10.1145/292540.292555
2. [SW Rel.] Abel, A., Allais, G., Cockx, J., Danielsson, N.A., Hausmann, P., Nordvall Forsberg, F., Norell, U., López Juan, V., Sicard-Ramírez, A., and Vezzosi, A., *Agda 2* version 2.6.1.3, 2005–2021. Chalmers University of Technology and Gothenburg

- University. LIC: BSD3. URL: <https://wiki.portal.chalmers.se/agda/pmwiki.php>, VCS: <https://github.com/agda/agda>.
3. Abel, A., and Bernardy, J.: A unified view of modalities in type systems. *Proc. ACM Program. Lang.* 4(ICFP), 90:1–90:28 (2020). DOI: 10.1145/3408972
 4. Algehed, M.: A Perspective on the Dependency Core Calculus. In: Alvim, M.S., and Delaune, S. (eds.) *Proceedings of the 13th Workshop on Programming Languages and Analysis for Security, PLAS@CCS 2018*, Toronto, ON, Canada, October 15-19, 2018, pp. 24–28. ACM (2018). DOI: 10.1145/3264820.3264823
 5. Algehed, M., and Russo, A.: Encoding DCC in Haskell. In: *Proceedings of the 2017 Workshop on Programming Languages and Analysis for Security, PLAS@CCS 2017*, Dallas, TX, USA, October 30, 2017, pp. 77–89. ACM (2017). DOI: 10.1145/3139337.3139338
 6. Askarov, A., Hunt, S., Sabelfeld, A., and Sands, D.: Termination-Insensitive Non-interference Leaks More Than Just a Bit. In: Jajodia, S., and López, J. (eds.) *Computer Security - ESORICS 2008*, 13th European Symposium on Research in Computer Security, Málaga, Spain, October 6-8, 2008. *Proceedings. LNCS*, vol. 5283, pp. 333–348. Springer, Heidelberg (2008). DOI: 10.1007/978-3-540-88313-5_22
 7. Bernardy, J., Boespflug, M., Newton, R.R., Jones, S.P., and Spiwack, A.: Linear Haskell: practical linearity in a higher-order polymorphic language. *CoRR* abs/1710.09756 (2017)
 8. Bowman, W.J., and Ahmed, A.: Noninterference for free. In: Fisher, K., and Reppy, J.H. (eds.) *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015*, Vancouver, BC, Canada, September 1-3, 2015, pp. 101–113. ACM (2015). DOI: 10.1145/2784731.2784733
 9. Buiras, P., Vytiniotis, D., and Russo, A.: HLIO: mixing static and dynamic typing for information-flow control in Haskell. In: Fisher, K., and Reppy, J.H. (eds.) *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015*, Vancouver, BC, Canada, September 1-3, 2015, pp. 289–301. ACM (2015). DOI: 10.1145/2784731.2784758
 10. Gaboardi, M., Katsumata, S., Orchard, D.A., Breuvar, F., and Uustalu, T.: Combining effects and coeffects via grading. In: Garrigue, J., Keller, G., and Sumii, E. (eds.) *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016*, Nara, Japan, September 18-22, 2016, pp. 476–489. ACM (2016). DOI: 10.1145/2951913.2951939
 11. Goguen, J.A., and Meseguer, J.: Security Policies and Security Models. In: *1982 IEEE Symposium on Security and Privacy*, Oakland, CA, USA, April 26-28, 1982, pp. 11–20. IEEE Computer Society (1982). DOI: 10.1109/SP.1982.10014
 12. Gregersen, S.O., Bay, J., Timany, A., and Birkedal, L.: Mechanized logical relations for termination-insensitive noninterference. *Proc. ACM Program. Lang.* 5(POPL), 1–29 (2021). DOI: 10.1145/3434291
 13. Hedin, D., and Sabelfeld, A.: A Perspective on Information-Flow Control. In: *Software Safety and Security - Tools for Analysis and Verification*. Ed. by T. Nipkow, O. Grumberg, and B. Hauptmann, pp. 319–347. IOS Press (2012). DOI: 10.3233/978-1-61499-028-4-319
 14. Heintze, N., and Riecke, J.G.: The SLam Calculus: Programming with Secrecy and Integrity. In: MacQueen, D.B., and Cardelli, L. (eds.) *POPL '98*, *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, San Diego, CA, USA, January 19-21, 1998, pp. 365–377. ACM (1998). DOI: 10.1145/268946.268976

15. Hirsch, A.K., and Cecchetti, E.: Giving semantics to program-counter labels via secure effects. *Proc. ACM Program. Lang.* 5(POPL), 1–29 (2021). doi: 10.1145/3434316
16. Jones, S.L.P., and Wadler, P.: Imperative Functional Programming. In: Deusen, M.S.V., and Lang, B. (eds.) *Conference Record of the Twentieth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Charleston, South Carolina, USA, January 1993, pp. 71–84. ACM Press (1993). doi: 10.1145/158511.158524
17. Kain, R.Y., and Landwehr, C.E.: On Access Checking in Capability-Based Systems. *IEEE Trans. Software Eng.* 13(2), 202–207 (1987). doi: 10.1109/TSE.1987.232892
18. Katsumata, S.: Parametric effect monads and semantics of effect systems. In: Jagannathan, S., and Sewell, P. (eds.) *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14*, San Diego, CA, USA, January 20–21, 2014, pp. 633–646. ACM (2014). doi: 10.1145/2535838.2535846
19. Kavvos, G.A.: Modalities, cohesion, and information flow. *Proc. ACM Program. Lang.* 3(POPL), 20:1–20:29 (2019). doi: 10.1145/3290333
20. Launchbury, J., and Jones, S.L.P.: Lazy Functional State Threads. In: Sarkar, V., Ryder, B.G., and Soffa, M.L. (eds.) *Proceedings of the ACM SIGPLAN'94 Conference on Programming Language Design and Implementation (PLDI)*, Orlando, Florida, USA, June 20–24, 1994, pp. 24–35. ACM (1994). doi: 10.1145/178243.178246
21. Lewis, J.R., Launchbury, J., Meijer, E., and Shields, M.: Implicit Parameters: Dynamic Scoping with Static Types. In: Wegman, M.N., and Reps, T.W. (eds.) *POPL 2000, Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Boston, Massachusetts, USA, January 19–21, 2000, pp. 108–118. ACM (2000). doi: 10.1145/325694.325708
22. Li, P., and Zdancewic, S.: Encoding Information Flow in Haskell. In: *19th IEEE Computer Security Foundations Workshop, (CSFW-19 2006)*, 5–7 July 2006, Venice, Italy, p. 16. IEEE Computer Society (2006). doi: 10.1109/CSFW.2006.13
23. Mitchell, J.C.: *Foundations for programming languages*. MIT Press (1996)
24. Miyamoto, K., and Igarashi, A.: A modal foundation for secure information flow. In: *Proceedings of IEEE Foundations of Computer Security (FCS)*, pp. 187–203 (2004)
25. Moggi, E.: Notions of Computation and Monads. *Inf. Comput.* 93(1), 55–92 (1991). doi: 10.1016/0890-5401(91)90052-4
26. [SW] Myers, A.C., Zheng, L., Zdancewic, S., Chong, S., and Nystrom, N., *Jif: Java information flow* version 3.0, 2006. URL: <http://www.cs.cornell.edu/jif>.
27. Orchard, D., Liepelt, V., and III, H.E.: Quantitative program reasoning with graded modal types. *Proc. ACM Program. Lang.* 3(ICFP), 110:1–110:30 (2019). doi: 10.1145/3341714
28. Orchard, D.A., and Petricek, T.: Embedding effect systems in Haskell. In: Swierstra, W. (ed.) *Proceedings of the 2014 ACM SIGPLAN symposium on Haskell*, Gothenburg, Sweden, September 4–5, 2014, pp. 13–24. ACM (2014). doi: 10.1145/2633357.2633368
29. Parker, J., Vazou, N., and Hicks, M.: LWeb: information flow security for multi-tier web applications. *Proc. ACM Program. Lang.* 3(POPL), 75:1–75:30 (2019). doi: 10.1145/3290388
30. Petricek, T., Orchard, D.A., and Mycroft, A.: Coeffects: a calculus of context-dependent computation. In: Jeuring, J., and Chakravarty, M.M.T. (eds.) *Proceed-*

- ings of the 19th ACM SIGPLAN international conference on Functional programming, Gothenburg, Sweden, September 1-3, 2014, pp. 123–135. ACM (2014). DOI: 10.1145/2628136.2628160
31. Polikarpova, N., Stefan, D., Yang, J., Itzhaky, S., Hance, T., and Solar-Lezama, A.: Liquid information flow control. *Proc. ACM Program. Lang.* 4(ICFP), 105:1–105:30 (2020). DOI: 10.1145/3408987
 32. Rajani, V., and Garg, D.: On the expressiveness and semantics of information flow types. *J. Comput. Secur.* 28(1), 129–156 (2020). DOI: 10.3233/JCS-191382
 33. Russo, A.: Functional pearl: two can keep a secret, if one of them uses Haskell. In: Fisher, K., and Reppy, J.H. (eds.) *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015, Vancouver, BC, Canada, September 1-3, 2015*, pp. 280–288. ACM (2015). DOI: 10.1145/2784731.2784756
 34. Russo, A., Claessen, K., and Hughes, J.: A library for light-weight information-flow security in haskell. In: Gill, A. (ed.) *Proceedings of the 1st ACM SIGPLAN Symposium on Haskell, Haskell 2008, Victoria, BC, Canada, 25 September 2008*, pp. 13–24. ACM (2008). DOI: 10.1145/1411286.1411289
 35. Sabelfeld, A., and Myers, A.C.: Language-based information-flow security. *IEEE J. Sel. Areas Commun.* 21(1), 5–19 (2003). DOI: 10.1109/JSAC.2002.806121
 36. Schoepe, D., Hedin, D., and Sabelfeld, A.: SeLINQ: tracking information across application-database boundaries. In: Jeuring, J., and Chakravarty, M.M.T. (eds.) *Proceedings of the 19th ACM SIGPLAN international conference on Functional programming, Gothenburg, Sweden, September 1-3, 2014*, pp. 25–38. ACM (2014). DOI: 10.1145/2628136.2628151
 37. Shikuma, N., and Igarashi, A.: Proving Noninterference by a Fully Complete Translation to the Simply Typed Lambda-Calculus. *Log. Methods Comput. Sci.* 4(3) (2008). DOI: 10.2168/LMCS-4(3:10)2008
 38. [SW] Simonet, V., *Flow Caml* 2003. URL: <http://cristal.inria.fr/~simonet/soft/flowcaml/>.
 39. Stefan, D., Russo, A., Mitchell, J.C., and Mazières, D.: Flexible dynamic information flow control in Haskell. In: Claessen, K. (ed.) *Proceedings of the 4th ACM SIGPLAN Symposium on Haskell, Haskell 2011, Tokyo, Japan, 22 September 2011*, pp. 95–106. ACM (2011). DOI: 10.1145/2034675.2034688
 40. Tse, S., and Zdancewic, S.: Translating dependency into parametricity. In: Okasaki, C., and Fisher, K. (eds.) *Proceedings of the Ninth ACM SIGPLAN International Conference on Functional Programming, ICFP 2004, Snow Bird, UT, USA, September 19-21, 2004*, pp. 115–125. ACM (2004). DOI: 10.1145/1016850.1016868
 41. Vassena, M., Buiras, P., Waye, L., and Russo, A.: Flexible Manipulation of Labeled Values for Information-Flow Control Libraries. In: Askoxylakis, I.G., Ioannidis, S., Katsikas, S.K., and Meadows, C.A. (eds.) *Computer Security - ESORICS 2016 - 21st European Symposium on Research in Computer Security, Heraklion, Greece, September 26-30, 2016, Proceedings, Part I. LNCS, vol. 9878*, pp. 538–557. Springer, Heidelberg (2016). DOI: 10.1007/978-3-319-45744-4_27
 42. Vassena, M., and Russo, A.: On Formalizing Information-Flow Control Libraries. In: Murray, T.C., and Stefan, D. (eds.) *Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security, PLAS@CCS 2016, Vienna, Austria, October 24, 2016*, pp. 15–28. ACM (2016). DOI: 10.1145/2993600.2993608
 43. Vassena, M., Russo, A., Buiras, P., and Waye, L.: MAC A verified static information-flow control library. *Journal of Logical and Algebraic Methods in Programming* 95, 148–180 (2018). DOI: <https://doi.org/10.1016/j.jlamp.2017.12.003>

44. Vassena, M., Russo, A., Garg, D., Rajani, V., and Stefan, D.: From fine- to coarse-grained dynamic information flow control and back. *Proc. ACM Program. Lang.* 3(POPL), 76:1–76:31 (2019). DOI: [10.1145/3290389](https://doi.org/10.1145/3290389)
45. Wadler, P., and Thiemann, P.: The marriage of effects and monads. *ACM Trans. Comput. Log.* 4(1), 1–32 (2003). DOI: [10.1145/601775.601776](https://doi.org/10.1145/601775.601776)
46. Zdancewic, S.A.: *Programming languages for information security*. Cornell University (2002)