

MAC

A Verified Static Information-Flow Control Library

Marco Vassena, Alejandro Russo

Chalmers University of Technology, Sweden

Pablo Buiras, Lucas Wayne

Harvard University, USA

Abstract

The programming language Haskell plays a unique, privileged role in information-flow control (IFC) research: *it is able to enforce information security via libraries*. Many state-of-the-art IFC libraries (e.g., **LIO** and **HLIO**) support a variety of advanced features like mutable data structures, exceptions, and concurrency, whose subtle interaction makes verification of security guarantees challenging. In this work, we focus on **MAC**, a statically-enforced IFC library for Haskell. In **MAC**, like other IFC libraries, computations have a well-established algebraic structure for computations (i.e., monads) responsible to manipulate *labeled values*—values coming from an abstract data type which associates a sensitivity label to a piece of information. In this work, we enrich labeled values with a *functor* structure and provide an *applicative functor* operator which encourages a more functional programming style and simplifies code. Furthermore, we present a full-fledged, mechanically-verified model of **MAC**. Specifically, we show progress-insensitive noninterference for our sequential calculus and pinpoint sufficient requirements on the scheduler to prove progress-sensitive noninterference for our concurrent calculus. For that, we study the security guarantees of **MAC** using *term erasure*, a proof technique that ensures that the same public output should be produced if secrets are erased before or after program execution. As another contribution, we extend term erasure with *two-steps erasure*, a flexible novel technique that greatly simplifies the noninterference proof and helps to prove many advanced features of **MAC**.

Keywords: Information Flow Control, Non Interference, Functional Programming, Haskell, Agda.

Email addresses: `vassena@chalmers.se` (Marco Vassena), `russo@chalmers.se` (Alejandro Russo), `pbuiras@seas.harvard.edu` (Pablo Buiras), `lwayne@seas.harvard.edu` (Lucas Wayne)

1. Introduction

Nowadays, many applications (apps) manipulate users’ private data. Such apps *could have been written by anyone* and users who wish to benefit from their functionality are forced to grant them access to their data—something that most users will do without a second thought (Meurer and Wismler, 2012). Once apps collect users’ information, there are no guarantees about how they handle it, thus leaving room for data theft and data breach by malicious apps. The key to guaranteeing security without sacrificing functionality is not granting or denying access to sensitive data, but rather ensuring that *information flows only into appropriate places*.

Language-based Information-Flow Control (IFC) (Sabelfeld and Myers, 2003a) is a promising approach to enforcing information security in software. A traditional IFC enforcement scrutinizes how data of different sensitivity levels (e.g., public or private) flows within a program, detects when an unsafe flow of information occurs and take action to suppress the leakage. To do that, most IFC tools require the design of new languages, compilers, interpreters, or modifications to the runtime, e.g., (Myers, 1999; Pottier and Simonet, 2002; Roy et al., 2009; Broberg et al., 2013). Nonetheless, in the functional programming language Haskell, the strict separation between side-effect free and side-effectful code enables lightweight security enforcements. Specifically, it is possible to build a secure programming language atop Haskell, as an embedded domain-specific language that gets distributed and used as a Haskell library (Li and Zdancewic, 2006). Many of the state-of-the-art Haskell security libraries, namely **LIO** (Stefan et al., 2011a), **HLIO** (Buiras et al., 2015), and **MAC** (Russo, 2015), bring ideas from Mandatory Access Control (Bell and La Padula, 1976) into a language-based setting.¹ These libraries promote a *secure-by-construction* programming model: any program written against their API does not leak secrets. This model is attractive, because it protects not only against benign code that leaks accidentally, e.g., due to a software bug, but also against a malicious program designed to do so. Every computation in such libraries has a *current label* which is used to (i) approximate the sensitivity level of all the data in scope and (ii) restrict subsequent side-effects which might compromise security. IFC uses labels to model the sensitivity of data, which are then organized in a security lattice (Denning and Denning, 1977) specifying the allowed flows of information, i.e., $\ell_1 \sqsubseteq \ell_2$ means that data with label ℓ_1 can flow into entities labeled with ℓ_2 .

Although these libraries are parameterized on the security lattice, for simplicity we focus on the classic two-point lattice with labels H and L to respectively denote secret (high) and public (low) data, and where $H \not\sqsubseteq L$ is the only disallowed flow. Figure 1 shows a graphical representation of a public computation in these libraries, i.e., a computation with current label L . The computation can read or write data in scope, which is considered




Figure 1: Public computation

¹From now on, we simply use the term libraries when referring to **LIO**, **HLIO**, and **MAC**.

public (e.g., average temperature of 17°C in the Swedish summer), and it can write to public (L -) or secret (H -) sinks. By contrast, a secret computation, i.e., a computation with current label H , can also read and write data in its scope, which is considered sensitive, but in order to prevent information leaks it can *only* write to sensitive/secret sinks. Structuring computations in this manner ensures that sensitive data does not flow into public entities, a policy known as noninterference (Goguen and Meseguer, 1982). While secure, programming in this model can be overly restrictive for users who want to manipulate differently-labeled values.

To address this shortcoming, libraries introduce the notion of a *labeled value* as an abstract data type which protects values with explicit labels, in addition to the current label. Figure 2 shows a public computation with access to both public and sensitive pieces of information, such as a password (pwd). Public computations can freely manipulate sensitive




Figure 2: Labeled values

labeled values provided that they are treated as black boxes, i.e., they can be stored, retrieved, and passed around as long as its content is not inspected. Libraries **LIO** and **HLIO** even allow public computations to inspect the contents of sensitive labeled values, raising the current label to H to keep track of the fact that a secret is in scope—this variant is known as a *floating-label* system.

Reading sensitive data usually amounts to “tainting” the entire context or ensuring the context is as sensitive as the data being observed. As a result, the system is susceptible to an issue known as *label creep*: reading too many secrets may cause the current label to be so high in the lattice that the computation can no longer perform any useful side effects. To address this problem, libraries provide a primitive which enables public computations to spawn sub-computations that access sensitive labeled values without tainting the parent. In a sequential setting, such sub-computations are implemented by special function calls. In the presence of concurrency, however, they must be executed in a different thread to avoid compromising security through *internal timing* and *termination covert channels* (Stefan et al., 2012a).

Practical programs need to manipulate sensitive labeled values by transforming them. It is quite common for these operations to be naturally free of I/O or other side effects, e.g., arithmetical or algebraic operations, especially in applications like image processing, cryptography, or data aggregation for statistical purposes. Writing such functions, known as *pure* functions, is the bread and butter of functional programming style, and is known to improve programmer productivity, encourage code reuse, and reduce the likelihood of bugs (Hughes, 1984). Nevertheless, the programming model involving sub-computations that manipulate secrets forces an imperative style, whereby computations must be structured into separate compartments that must communicate explicitly. While side-effecting instructions have an underlying algebraic structure, called monad (Moggi, 1991), research literature has neglected studying the algebraic structure of labeled values and their consequences for the programming model. To empower programmers with the simpler, functional style, we propose additional operations that allow pure functions to securely manipulate labeled values, specifically by

means of a structure similar to *applicative functors* (Mcbride and Paterson, 2008). In particular, this structure is useful in concurrent settings where it is no longer necessary to spawn threads to manipulate sensitive data, thus making the code less imperative (i.e., side-effect free).

Additionally, practical programs often aggregate information from heterogeneous sources. For that, programs need to upgrade labeled values to an upper bound of the labels being involved before data can be combined. In previous incarnations of the libraries, such relabelings require to spawn threads just for that purpose. As before, the reason for that is libraries decoupling every computation which manipulate sensitive data—even those for simply relabeling—so that the internal timing and termination covert channels imposed no threats. In this light, we introduce a primitive to securely relabel labeled values, which can be applied irrespective of the computation’s current label and does not require spawning threads.

We provide a mechanized security proof for the security library **MAC**² and claim our results also apply to **LIO** and **HLIO**. **MAC** has fewer lines of code and leverages types to enforce confidentiality, thus making it ideal to model its semantics in a dependently-typed language like Agda. The contributions of this paper are:

1. We develop the first exhaustive full-fledged formalization of **MAC**, a state-of-the-art library for Information-Flow Control, in a call-by-need λ -calculus and prove progress-insensitive noninterference (PINI) for the sequential calculus.
2. We enrich the calculus with scheduler-parametric concurrency and prove progress-sensitive noninterference (PSNI) (Askarov et al., 2008) for a wide-range of deterministic schedulers, by formally identifying sufficient requirements on the scheduler to ensure PSNI—a novel aspect if compared with previous work (Stefan et al., 2012a; Heule et al., 2015). We leverage on the generality of our result and prove that **MAC** is secure by instantiating our PSNI theorem with a round-robin scheduler, i.e., the scheduler used by GHC’s runtime system.
3. We corroborate our results with an extensive mechanized proof developed in the Agda proof assistant that counts more than 4000 lines of code. The mechanization has provided us with stimulating insights and pinpointed problems in proofs of similar works.
4. We improve and simplify the term-erasure proof technique by proposing a novel flexible technique called *two-steps* erasure, which we utilize systematically to prove that many advanced features are secure, especially those that change the security level of other terms and detect exceptions.
5. We introduce a *functor* structure, a relabeling primitive and an *applicative* operator that give flexibility to programmers, by upgrading labeled values and conveniently aggregating heterogeneously labeled data.
6. We have released a prototype of our ideas in the **MAC** library³.

Highlights. This work builds on our previous papers “Flexible Manipulation of Labeled Values for Information-Flow Control Libraries” (Vassena et al., 2016) and “On

²Available at <https://bitbucket.org/MarcoVassena/mac-model/>.

³<https://hackage.haskell.org/package/mac>

```

data Labeled  $\ell$   $a = \text{Labeled}^{\text{TCB}} a$ 
data MAC  $\ell$   $a = \text{MAC}^{\text{TCB}} \{ \text{run}^{\text{TCB}} :: \text{IO } a \}$ 
instance Monad (MAC  $\ell$ )
label ::  $\ell_{\text{L}} \sqsubseteq \ell_{\text{H}} \Rightarrow a \rightarrow \text{MAC } \ell_{\text{L}} (\text{Labeled } \ell_{\text{H}} a)$ 
unlabel ::  $\ell_{\text{L}} \sqsubseteq \ell_{\text{H}} \Rightarrow \text{Labeled } \ell_{\text{L}} a \rightarrow \text{MAC } \ell_{\text{H}} a$ 

```

Figure 3: Core API for **MAC**.

Formalizing Information-Flow Control Libraries” (Vassena and Russo, 2016), which we have blended and significantly rewritten and corrected in a few technical inaccuracies. We have integrated these works with several examples and shaped them into a uniform, coherent and comprehensive story of this line of work. We summarize the novel contributions of this article as follows:

- Uniform, coherent and comprehensive account of a formal model of **MAC**;
- Integration of examples in the description of the features of the library;
- Fixed several technical inaccuracies in the semantics of the calculus;
- Simplification and full account of the scheduler-parametric PSNI proof.

In the following, we point out the technical differences between this article and the conference version in footnotes.

This paper is organized as follows. Section 2 gives an overview of **MAC**. Section 3 formalizes the core of **MAC** in a simply-typed call-by-need lambda-calculus. Section 4 presents a secure primitive that regulates the interaction between computations at different security levels. Sections 5 and 6 extend the calculus with other advanced practical features, namely exceptions and mutable references. Section 7 proves that the sequential calculus satisfies progress-insensitive noninterference (PINI). Section 8 extends the calculus with concurrency and Section 9 presents functor, applicative, and relabeling operations. Section 10 gives the security guarantee of the concurrent calculus, which satisfies progress-sensitive noninterference (PSNI). Section 11 gives related work and Section 12 concludes.

2. Overview

In **MAC**, each label is represented as an abstract data type. Figure 3 shows the core part of **MAC**’s API. Abstract data type *Labeled* ℓ a classifies data of type a with a security label ℓ . For instance, $\text{pwd} :: \text{Labeled } H \text{ String}$ is a sensitive string, while $\text{rating} :: \text{Labeled } L \text{ Int}$ is a public integer. (Symbol $::$ is used to describe the type of terms in Haskell.) Abstract data type *MAC* ℓ a denotes a (possibly) side-effectful secure computation which handles information at sensitivity level ℓ and yields a value of type a as a result. A *MAC* ℓ a computation enjoys a monadic structure, i.e., it is built using the fundamental operations $\text{return} :: a \rightarrow \text{MAC } \ell a$ and $(\gg=) :: \text{MAC } \ell a \rightarrow$

$(a \rightarrow MAC \ell b) \rightarrow MAC \ell b$ (read as “bind”). The operation *return* x produces a computation that returns the value denoted by x and produces no side-effects. The function $(\gg=)$ is used to *sequence* computations and their corresponding side-effects. Specifically, $m \gg= f$ takes a computation m and function f which will be applied to the *result* produced by running m and yields the resulting computation. We sometimes use Haskell’s **do**-notation to write such monadic computations. For example, the program $m \gg= \lambda x \rightarrow return (x + 1)$, which adds 1 to the value produced by m , can be written as shown in Figure 4.

<pre>do x ← m return (x + 1)</pre> <p>Figure 4: do-notation</p>

2.1. Secure flows of information

Generally speaking, side-effects in a $MAC \ell a$ computation can be seen as actions which either read or write data. Such actions, however, need to be conceived in a manner that respects the sensitivity of the computations’ results as well as the sensitivity of sources and sinks of information modeled as labeled values. The functions *label* and *unlabel* allow $MAC \ell a$ computations to securely interact with labeled values. To help readers, we indicate the relationship between type variables in their subindexes, i.e., we use ℓ_L and ℓ_H to attest that $\ell_L \sqsubseteq \ell_H$. If a $MAC \ell_L$ computation writes data into a sink, the computation should have at most the sensitivity of the sink itself. This restriction, known as *no write-down* (Bell and La Padula, 1976), respects the sensitivity of the data sink, e.g., the sink never receives data more sensitive than its label. In the case of function *label*, it creates a fresh labeled value, which from the security point of view can be seen as allocating a fresh location in memory and immediately writing a value into it—thus, it applies the no write-down principle. In the type signature of *label*, what appears on the left-hand side of the symbol \Rightarrow are *type constraints*. They represent properties that must be statically fulfilled about the types appearing on the right-hand side of \Rightarrow . Type constraint $\ell_L \sqsubseteq \ell_H$ ensures that when calling *label* x (for some x in scope), the computation creates a labeled value only if ℓ_L , i.e. the current label of the computation, is no more confidential than ℓ_H , i.e. the sensitivity of the created labeled value. In contrast, a computation $MAC \ell_H a$ is only allowed to read labeled values at most as sensitive as ℓ_H —observe the type constraint $\ell_L \sqsubseteq \ell_H$ in the type signature of *unlabel*. This restriction, known as *no read-up* (Bell and La Padula, 1976), protects the confidentiality degree of the result produced by $MAC \ell_H a$, i.e. the result might only involve data ℓ_L which is, at most, as sensitive as ℓ_H .

We remark that **MAC** is an embedded domain specific language (EDSL), implemented as a Haskell *library* of around 200 lines of code and programs written in **MAC** are secure-by-construction. What makes it possible to provide strong security guarantees via a library is the fact that Haskell type-system enforces a strict separation between side-effect free code, which is guaranteed *not* to perform side effects, and side-effectful code, where side-effects may occur⁴. Specifically side-effects, i.e., input-output operations, can only occur in monadic computations of type $IO a$. Crucially *pure* computations are inherently secure, while *IO* computations are potentially leaky.

⁴In the functional programming community, they are also known as *pure* and *impure* code respectively.

In **MAC**, a secure computation of type $MAC\ \ell\ a$ is internally represented as a wrapper around an $IO\ a$ computation, that is used to implement side-effectful features, such as references and concurrency. **MAC** provides security-by-construction because *impure* operations, i.e., those of type IO , can only be constructed using **MAC** label-annotated API, which accepts only those that are statically deemed secure. Function run^{TCB} extracts the underlying $IO\ a$ computation from a secure computation of type $MAC\ \ell\ a$. Thanks to the secure-by-construction design, the IO computation so obtained is secure and can be executed directly, without the need of additional protection mechanism, such as monitors. Note that the function run^{TCB} is part of the Trusted Computing Base (TCB), i.e., it is available only to trusted code. In what follows, we describe an example which illustrates **MAC**'s programming model, particularly the use of *label*, *unlabel*.

Example. The most common use of *label* is to classify data to be protected. As an example, consider the Haskell program listed in Figure 5, which prompts the user for a password through the terminal and then passes it to a routine to check if the password is listed on dictionaries of commonly used passwords. Observe that the program performs input-output operations: $putStrLn :: String \rightarrow IO\ ()$ prints to standard

```
p :: IO Bool
p = do
  putStrLn "Choose a password:"
  pwd ← getLine
  return (isWeak pwd)
```

Figure 5: The password is exposed in *isWeak*.

output and $getLine :: IO\ String$ reads from standard input. Clearly the content of variable *pwd* should be handled with care by $isWeak :: String \rightarrow IO\ Bool$. In particular a computation of type $IO\ Bool$ can also perform arbitrary output operations and potentially leak the password. One way to protect *pwd* is by writing all password-related operations, like *isWeak*, within **MAC**, where *pwd* is marked as sensitive data. Adjusting the type of *isWeak* appropriately, **MAC** prevents intentional or accidental leakage of the password. Several secure designs are possible, depending on how *isWeak* provides its functionality. For example a secure interface could be $isWeak :: Labeled\ H\ String \rightarrow MAC\ L\ (Labeled\ H\ Bool)$, where the outermost computation ($MAC\ L$) accounts for reading public data, e.g., fetching online dictionaries of common passwords, while the labeled result ($Labeled\ H\ Bool$), protects the sensitivity of this piece of information about the password, namely if it is weak or not. The type $isWeak :: Labeled\ H\ String \rightarrow MAC\ H\ Bool$ is also secure and additionally allows to read from secret channels, e.g., file `/etc/shadow`, to check that the password is not reused. Figure 6 shows the modifications to the code needed to use a secure password strength checker. Observe how *label* is used to mark *pwd* as sensitive by wrapping it inside a labeled expression of type $Labeled\ H\ String$. After that, the labeled password is passed to function *isWeak* by bind (\gg), function run^{TCB} executes the whole computation, whose labeled result is then pattern matched with $Labeled^{TCB}$,

```

isWeak :: Labeled H String → MAC L (Labeled H Bool)

p :: IO Bool
p = do putStrLn "Choose a password:"
      pwd ← getLine
      let lpwd = label pwd :: MAC L (Labeled H String)
          LabeledTCB b ← runTCB (lpwd >>= isWeak)
      return b

```

Figure 6: Label H protects the password in *isWeak*.

exposing the boolean value, that is finally returned⁵.

2.2. Implicit flows

The interaction between the current label of a computation and the no read-up and no write-down security policies makes implicit flow ill-typed. Consider for instance, the ill-typed program in Figure 7, that attempts to leak the value of the secret boolean in a public boolean. Unlike other IFC system, the code cannot branch on *secret* directly, because it is explicitly labeled, i.e., it has type *Labeled H Bool*, instead of *Bool*. In order to branch on sensitive data, the program needs first to unlabel it, thus incurring in the no read-up restriction that requires the computation to be sensitive as well, that is the program must have type *MAC H a* (for some type *a*). The only primitive that produces labeled data is *label*, which according to the no write-down restriction, prevents a sensitive computation from creating a public labeled value. Regardless of the branch taken, but for that reason, i.e., trying to label a piece of data with L in a computation labeled with H , the program in Figure 7 is ill-typed.

```

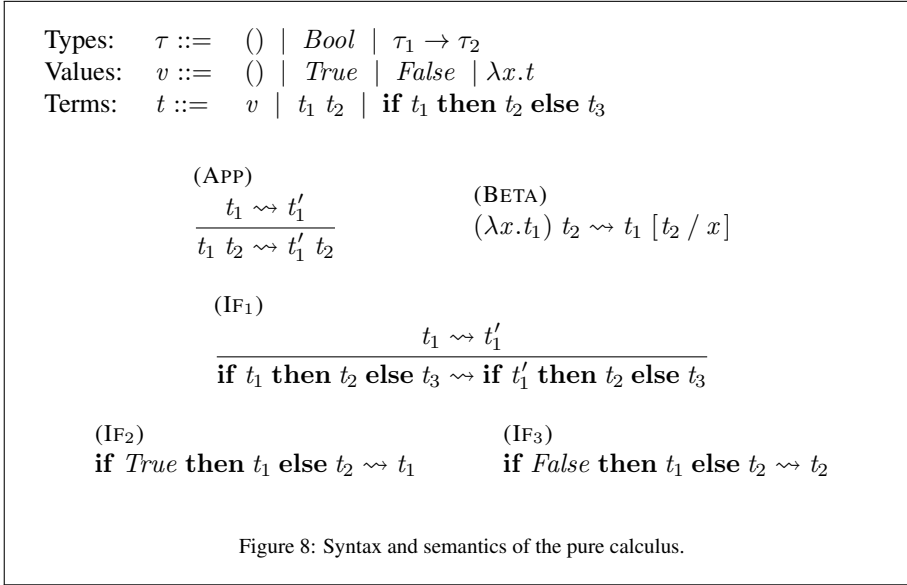
impl :: Labeled H Bool →
      MAC H (Labeled L Bool)
impl secret = do
  bool ← unlabel secret
  -- H ⊈ L
  if bool then label True
         else label False

```

Figure 7: Implicit flows are ill-typed.

Features Overview. Modern programming languages provide many abstractions that simplify the development of complex software. In the following, we extend **MAC** with additional primitives that make software development within **MAC** practical, without sacrificing security. The list of programming features securely supported in **MAC** include exception handling (Section 5), references (Section 6), concurrency (Section 8), functors 9 and synchronization primitives (Appendix B).

⁵In Figure 6, the code in the IO monad is trusted, hence the use of *run^{TCB}* and *Labeled^{TCB}*. The function *isWeak* is not trusted and the password is protected by **MAC** secure API.



3. The Core Calculus

This section formalizes **MAC** as a simply typed call-by-name λ -calculus extended with unit and boolean values and security primitives.

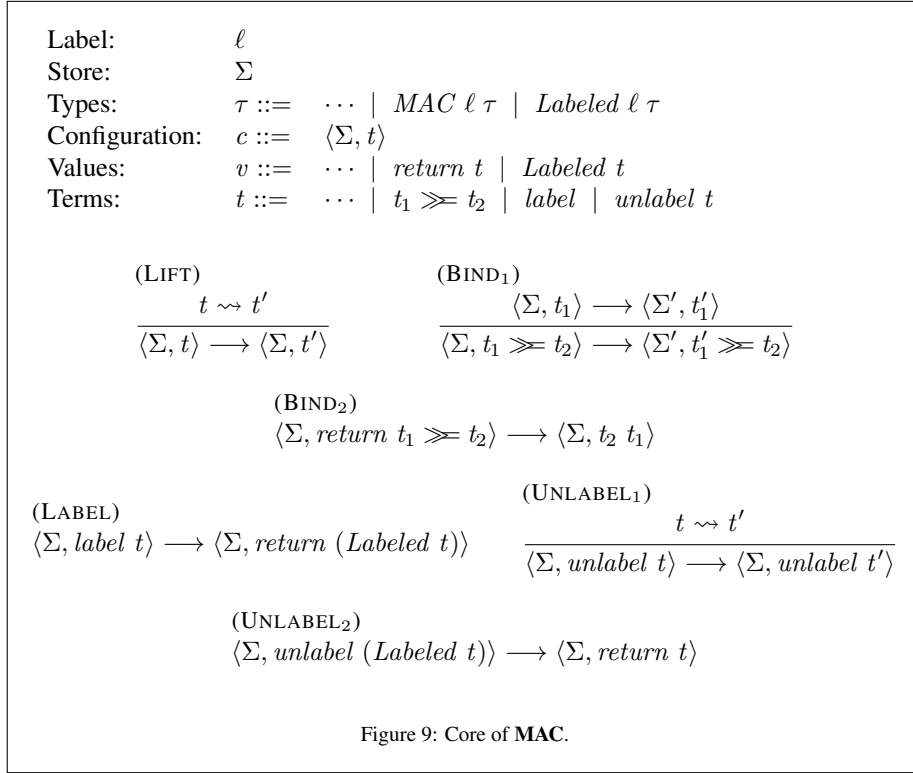
3.1. Pure Calculus

Figure 8 shows the formal syntax of the pure calculus underlying **MAC**, where meta variables τ , v and t denote respectively types, values, and terms. The typing judgment $\Gamma \vdash t : \tau$ denotes that term t has type τ assuming typing environment Γ . The typing rules of the pure calculus are standard and therefore omitted. The small-step semantics of the the calculus is represented by the relation $t_1 \rightsquigarrow t_2$, which denotes that term t_1 reduces to t_2 . Rule [BETA] indicates that the calculus has *call-by-name* semantics, because the argument of a function, evaluated to weak-head normal form by rule [APP], is not evaluated upon function application, but rather substituted in the body—we write $t_1 [x / t_2]$ for *capture-avoiding substitution*⁶. Rule [IF₁] evaluates the conditional of an if-then-else expression and rules [IF₂,IF₃] take the appropriate branch.

3.2. Core of **MAC**

We now extend this standard calculus with the security primitives of **MAC** as shown in Figure 9. Meta variable ℓ ranges over labels, which are assumed to form a lattice $(\mathcal{L}, \sqsubseteq)$. Labels are types in **MAC** despite we place them in a different

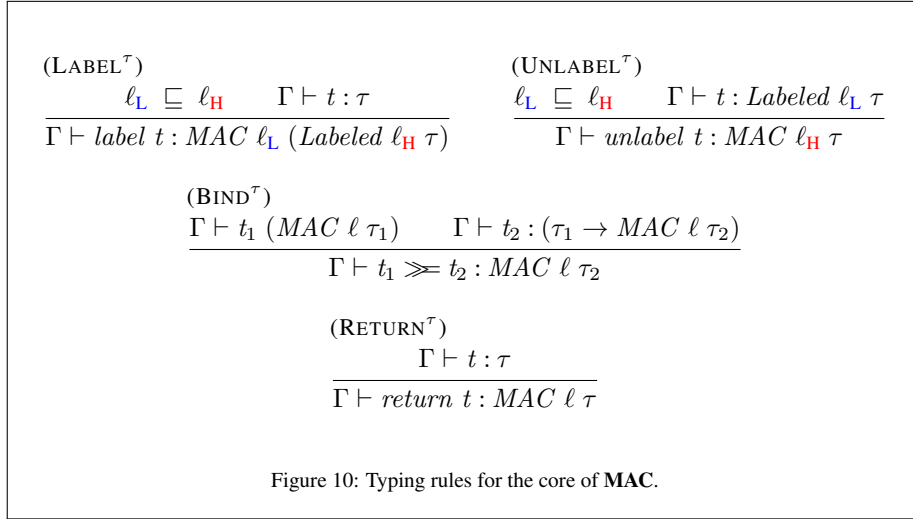
⁶In the machine-checked proofs all variables are De Bruijn indexes.



syntactic category named ℓ —this decision is made merely for clarity of exposition. The new type $\text{Labeled } \ell \tau$ represents a (possibly side-effect free) resource, which annotates with the security level ℓ a value $t :: \tau$ wrapped in Labeled . For example, $\text{Labeled } 42 :: \text{Labeled } L \text{Int}$ is a public integer. In the following, we introduce further forms of labeled resources, in particular mutable references in Section 6 and synchronization variables in Appendix B. The actual **MAC** implementation handles more labeled resources and provides a uniform implementation for them (Russo, 2015)⁷. The constructor Labeled is not available to the user, who can only use label and unlabel to create and inspect labeled resources, respectively.

A configuration $\langle \Sigma, t \rangle$ consists of a store Σ and a term t describing a computation of type $\text{MAC } \ell \tau$ and represents a secure computation at sensitivity level ℓ , which yields a value of type τ as result. For the moment, we ignore the store in the configuration (explained in Section 6). In order to enforce the security invariants, functions label

⁷ In our conference version (Vassena and Russo, 2016; Vassena et al., 2016), we follow the original **MAC** paper (Russo, 2015) and represent all labeled resources using the same labeled data type $\text{Res } t :: \text{Res } \ell \tau$, where $t :: \tau$ determines the kind of resource. For example $\text{Res } (\text{Id } 42) :: \text{Res } \ell (\text{Id } \text{Int})$ is a term representing a public integer. Here, for clarity of exposition, we use separate data types for each labeled resource. This design choice does not affect our results.



and *unlabel* live in the *MAC* monad and the typing rules in Figure 10 ensure that the label of the resource is compatible with the security level of the current computation, as explained in the previous section. We explain the relation between those typing rules and their corresponding type signatures given as Haskell API in Figure 3 as follows. The typing rules in Figure 10 are *type scheme rules*, i.e., there is such a judgment for every label ℓ_L and $\ell_H \in \mathcal{L}$, such that $\ell_L \sqsubseteq \ell_H$, where labels come from either type signatures or explicit type annotations in programs, as we showed in the previous section. The *type constraints* in the API, i.e., what appears before the symbol \Rightarrow , is placed as a premise of the corresponding typing rule. We remark that type constraints are built using *type classes*, a well-established feature of Haskell type system, therefore we do not discuss them any further (Wadler and Blott, 1989). Besides those primitives, computations are created using the standard monad operations *return* and \gg . The primitive *return* lifts a term in the monad at security level ℓ by means of typing rule [RETURN $^\tau$]. Unlike the Dependency Core Calculus (DCC) (Abadi et al., 1999), secure computations at different security levels do not mix in **MAC**: the typing rule [BIND $^\tau$] prevents that from happening—note the same label ℓ is expected both in the types of t_1 and t_2 . Just like rules [LABEL $^\tau$, UNLABEL $^\tau$], the typing rules [RETURN $^\tau$, BIND $^\tau$] are type scheme rules, i.e., there is such a rule for each label $\ell \in \mathcal{L}$. For easy exposition, in the following we give the type of **MAC**'s constructs as Haskell APIs.

We explicitly distinguish pure-term evaluation from top-level monadic-term evaluation, similarly to (Stefan et al., 2012b). The extended semantics is represented as the relation $c_1 \longrightarrow c_2$, which extends \rightsquigarrow via [LIFT]. The semantics rules in Figure 9 are fairly straight-forward and follow the pattern seen in the pure semantics, where some *context-rules*, e.g. [BIND $_1$, UNLABEL $_1$] reduce a redex subterm, and then the interesting rule fires, e.g. [BIND $_2$, UNLABEL $_2$]. In particular rule [BIND $_1$] executes the computation on the left of the bind and rule [BIND $_2$] extracts the result of the computation and feeds it to the right-side argument of (\gg). Rule [UNLABEL $_1$] evaluates the ar-

```

savePwd :: Labeled H String → MAC L (MAC H ())
savePwd lpwd = do putStrLnMAC "Saving new password"
                 return (passwdMAC pwd)

```

Figure 11: A nested computation that writes at security level L and H .

```

join :: ℓL ⊆ ℓH ⇒ MAC ℓH τ → MAC ℓL (Labeled ℓH τ)

```

Figure 12: Primitive *join*.

gument to labeled expression and rule [UNLABEL₂] returns its content. Rule [LABEL] creates a labeled expression by wrapping the argument in *Labeled* and returns it in the security monad. It is worth noting that thanks to the static nature of **MAC**, no run-time checks are needed to prevent insecure flows of information in these rules.

4. Label Creep

Let us continue the password example from the introduction. After checking that the password is strong enough, the program replaces the old password with the new one by updating file `/etc/shadow` with the new hashed password, using primitive $passwd^{MAC} :: Labeled\ H\ String \rightarrow MAC\ H\ ()$ —note that the label of the computation is H , in order to unlabel the password and hash it. (We treat password hashes as confidential data as well, because they could enable *offline* dictionary attacks otherwise.) The program should also inform the user that the password is being saved by printing on the screen a message. We consider printing on the screen as a public write operation, i.e., $putStrLn^{MAC} :: String \rightarrow MAC\ L\ ()$. Figure 11 shows the code of the discussed routine. Observe that $putStrLn^{MAC}\ "Saving\ new\ password" :: MAC\ L\ ()$ and $passwd^{MAC}\ pwd :: MAC\ H\ ()$ belong to different *MAC* computations. Therefore, both operations cannot coexist together, otherwise secret data, e.g., the password, could be unlabelled and then leaked on a public channel, e.g., standard output. Specifically the program in Figure 13 is rejected as ill-typed. Programs that handle data and channels with heterogeneous labels necessarily involve *nested MAC ℓ a* computations in its return type. In this case, the type of $savePwd\ lpwd :: MAC\ L\ (MAC\ H\ ())$ indicates that it is a public computations, which prints on the screen, and that *produces* as value a sensitive computation $MAC\ H\ Int$, which lastly writes to the sensitive file. Obviously having to nest computations complicates the programming model

```

do
  putStrLnMAC "Saving ..."
  passwdMAC lpwd

```

Figure 13: Ill-typed ($L \not\subseteq H$).

Terms: $t ::= \dots \mid \text{join } t$

$$\text{(JOIN)} \quad \frac{\langle \Sigma, t \rangle \Downarrow \langle \Sigma', \text{return } t' \rangle}{\langle \Sigma, \text{join } t \rangle \longrightarrow \langle \Sigma', \text{return } (\text{Labeled } t') \rangle}$$

Figure 14: Calculus with *join*.

of **MAC** and hinders its applicability⁸. For example, *savePwd lpwd* requires to run the public computation to completion first, and then execute the resulting sensitive computation. We recognize this pattern of returning nested computations as a static version of a problem known in dynamic systems as *label creep* (Sabelfeld and Myers, 2003b; Buiras et al., 2014)—which occurs when the context gets tainted to the point that no useful operations are allowed anymore. In a *sequential setting*, **MAC** provides the primitive *join*⁹, which alleviates this problem by safely *integrating* more sensitive computations into less sensitive ones.

4.1. Primitive *join*

Figure 12 shows the type signature of *join*. Intuitively, function *join* runs the computation of type $\text{MAC } \ell_{\mathbf{H}} \tau$ and wraps the result into a labeled expression to protect its sensitivity. As we will show in Section 7.5, programs written using the monadic API, *label*, *unlabel*, and *join* satisfy *progress-insensitive noninterference* (PINI), where leaks due to non-termination of programs are ignored. This design decision is similar to that taken by mainstream IFC compilers (e.g., (Myers et al., 2001; Simonet, 2003; Hedin et al., 2014)), where the most effective manner to exploit termination takes exponential time in the size (of bits) of the secret (Askarov et al., 2008).

In the semantics, Figure 14 extends terms with the new primitive *join t*. Rule [JOIN] formalizes the semantics of *join* using big-step semantics—similar to other work (Stefan et al., 2011b; Russo, 2015), we restrict ourselves to terminating computations. We write $\langle \Sigma, t \rangle \Downarrow \langle \Sigma', v \rangle$ if and only if v is a value and $\langle \Sigma, t \rangle \longrightarrow^* \langle \Sigma', v \rangle$, where relation \longrightarrow^* denotes the reflexive transitive closure of \longrightarrow . Rule [JOIN] executes the secure computation $t \Downarrow \text{return } t'$ and wraps the result t in *Labeled* to protect its sensitivity.¹⁰

⁸Remember that Haskell employs lazy evaluation, therefore the inner computations is not automatically evaluated, but needs to be explicitly executed. Only trusted code, using run^{TCB} can force evaluation of *MAC* computations.

⁹Not to be confused with the monadic $\text{join} :: \text{Monad } m \Rightarrow m (m a) \rightarrow m a$.

¹⁰We refrain from using *label t'* because we will soon add exceptions to secure computations.

```

savePwd :: Labeled H String → MAC L ()
savePwd lpwd = do putStrLnMAC "Saving new password"
                  join (passwdMAC pwd)
                  putStrLnMAC "Password saved"

```

Figure 15: Example revisited with *join*.

```

throw :: χ → MAC ℓ τ
catch :: MAC ℓ τ → (χ → MAC ℓ τ) → MAC ℓ τ

```

Figure 16: API for exceptions.

Revisited Example. By replacing *return* with *join*, we can simplify the program *savePwd* from the previous section: compare the two versions of the program in Figure 11 (using *return*) and in Figure 15 (using *join*). In Figure 15 the return type of *savePwd* does not involve nested computations, therefore the execution of the sensitive computation is not suspended, but rather follows directly after the public print statement.

5. Exception handling

Exception handling is a common programming language mechanism used to signal some anomalous condition and stop the execution of a program. It is sometimes possible to recover from such exceptional circumstances and resume execution afterwards. For instance, consider again the program *savePwd* in Figure 15. If primitive *passwd^{MAC}* fails due to some IO exception, e.g., file *etc/shadow* has already been opened or has not been found, the whole program crashes. Not supporting exceptions in the context of input-output operations, is not only impeding our programming model, but it is also insecure. In fact, exceptions change the control flow of a program, and an uncaught exception can propagate throughout a program and eventually crash it, potentially suppressing public events. For example, if *passwd^{MAC}* throws an exception, the program aborts before printing "Password saved" on the screen. Observe that, such behavior constitutes a leak, because the failure comes from a sensitive context, i.e., *passwd^{MAC}*, and therefore can depend on the value of the secret, i.e., the password. In this section, we incorporate exception handling primitives in **MAC** to remedy this situation, see Figure 16. Intuitively, *catch t₁ t₂* runs the computation *t₁* and recovers from a failure by passing the exception to the exception handler *t₂*. Section 5.2 discusses some subtleties between exception handling primitives and *join*, which may propagate exceptions from sensitive contexts to less sensitive ones, if neglected.

5.1. Calculus

For simplicity, we consider only one exception $\xi :: \chi$, where χ denotes an exception type. In the calculus, we extend terms with ξ , *throw t*, and *catch t₁ t₂*—see Figure 17.

Types: $\tau ::= \dots \mid \chi$
 Values: $v ::= \dots \mid \xi \mid \text{throw } t$
 Terms: $t ::= \dots \mid \text{catch } t_1 t_2$

(BIND _{χ})
 $\langle \Sigma, \text{throw } t_1 \gg t_2 \rangle \longrightarrow \langle \Sigma, \text{throw } t_1 \rangle$

(CATCH₁)

$$\frac{\langle \Sigma, t_1 \rangle \longrightarrow \langle \Sigma', t'_1 \rangle}{\langle \Sigma, \text{catch } t_1 t_2 \rangle \longrightarrow \langle \Sigma', \text{catch } t'_1 t_2 \rangle}$$

(CATCH₂)
 $\langle \Sigma, \text{catch } (\text{return } t_1) t_2 \rangle \longrightarrow \langle \Sigma, \text{return } t_1 \rangle$

(CATCH₃)
 $\langle \Sigma, \text{catch } (\text{throw } t_1) t_2 \rangle \longrightarrow \langle \Sigma, t_2 t_1 \rangle$

Figure 17: Exception handling primitives.

Term *throw* t aborts the current *MAC* computation with exception t , see rule [BIND _{χ}]. Term *catch* $t_1 t_2$ evaluates computation t_1 via rule [CATCH₁], and either it returns the result, if the computation succeeds, i.e., rule [CATCH₂], or it attempts to recover a failure by running exception handler t_2 , if the computation throws an exception, i.e., rule [CATCH₃].

5.2. Join and exceptions

The interplay between exceptions and *join* is delicate and security might be at stake if these two features were naively combined (Stefan et al., 2012b; Hritcu et al., 2013). Observe that type signatures in Figure 16 hint that exceptions can be thrown and caught among computations with the same label—a design decision which does not break security guarantees. Nevertheless, information can be leaked if exceptions thrown in sensitive computations are propagated to less sensitive ones. From now on, we refer to exceptions raised in a sensitive *MAC* computation as *sensitive exceptions*. In fact, sensitive exceptions can affect the control-flow of less sensitive computations and thus suppressing observable events, giving place to an implicit flow¹¹. In our calculus, *join* is the only primitive that combines computations with different labels and thus is potentially vulnerable to this attack. In order to close leaks via exceptions, **MAC** modifies the semantics of *join* to *mask* exceptions, preventing them to propagate to less sensitive computations—this solution is similar to previous work (Stefan et al., 2012b; Hritcu et al., 2013).

¹¹We refer interested readers to (Russo, 2015) for further details about this attack.

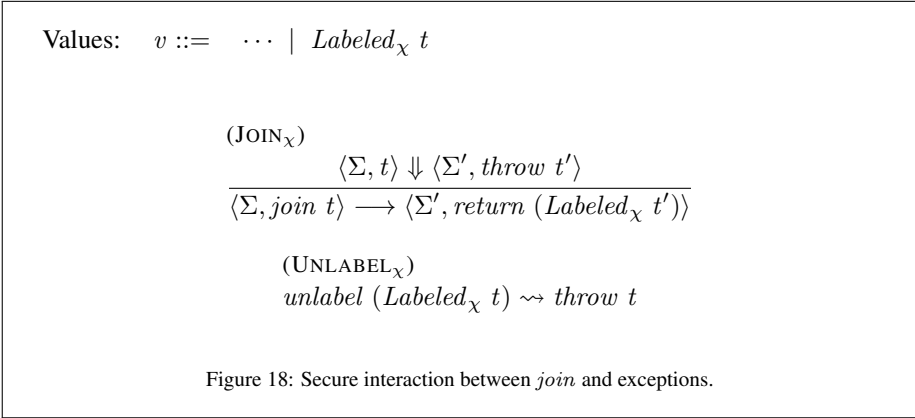


Figure 18 implements this countermeasure. Firstly it adds a new internal constructor $Labeled_{\chi} t$ denoting a labeled value (of type $Labeled \ell \tau$) which contains inside the exception ($t :: \chi$). Rule [JOIN $_{\chi}$] shows the semantics for *join* t when exceptions are triggered: *exceptions are not propagated further but rather returned inside a labeled expression*. Under this programming model, it is necessary to inspect the return value of *join* to determine if the computation terminated abnormally. The attacker must then *unlabel* the result to observe the exception, see rule [UNLABEL $_{\chi}$]. Observe that, since this operation is subject to *no read-up*, sensitive exceptions are not observable from less sensitive computations. As a consequence of this programming model, *only* sensitive computations can handle sensitive exceptions. Consider again the program *savePwd* in the example from Figure 15. The program prints "Password Saved" even though $passwd^{MAC}$ might have actually failed: it would be insecure to do otherwise! The only way to observe and recover from a failure of $passwd^{MAC}$, without compromising security, is to explicitly surround it with a *catch* block, i.e., *catch* ($passwd^{MAC} pwd$) *handler*, and lift that computation with *join*.

6. References

Mutable references are an imperative feature often needed to boost the performance of algorithms. Following the password example from the previous sections, we might want to reject weak passwords that are vulnerable to dictionary attacks. To do that, in Figure 19, function *fetchDict* fetches a list of words from a dictionary available in the system—we consider the content of a dictionary to be public information therefore the computation has security level L . Depending on the local system language, we can tweak the function to pick an appropriate dictionary, for example *fetchDict* "en" fetches English words from dictionary "usr/share/dict-en". A password-strength checker application could test a password against multiple dictionaries, which would require to call *fetchDict* multiple times. Since dictionaries are seldom changed, it is wasteful to fetch the same dictionary multiple times, therefore, using references, we implement a simple caching mechanism that avoids the overhead.


```

fetchDict :: String → MAC L [String]
fetchDict lang = readFile "usr/share/dict" ++ "-" ++ lang

fetchCacheDict : Ref L (Map String [String]) → String → MAC L [String]
fetchCacheDict r lang = do
  dicts ← read r
  case lookup lang dicts of
    Just dict → return dict
    Nothing → do
      dict ← fetchDict lang
      write (insert dict dicts) r
      return dict

```

Figure 19: *fetchCacheDict* is a cached version of *fetchDict*.

```

data Ref l τ
new ::  $l_L \sqsubseteq l_H \Rightarrow \tau \rightarrow \text{MAC } l_L (\text{Ref } l_H \tau)$ 
read ::  $l_L \sqsubseteq l_H \Rightarrow \text{Ref } l_L \tau \rightarrow \text{MAC } l_H \tau$ 
write ::  $l_L \sqsubseteq l_H \Rightarrow \tau \rightarrow \text{Ref } l_H \tau \rightarrow \text{MAC } l_L ()$ 

```

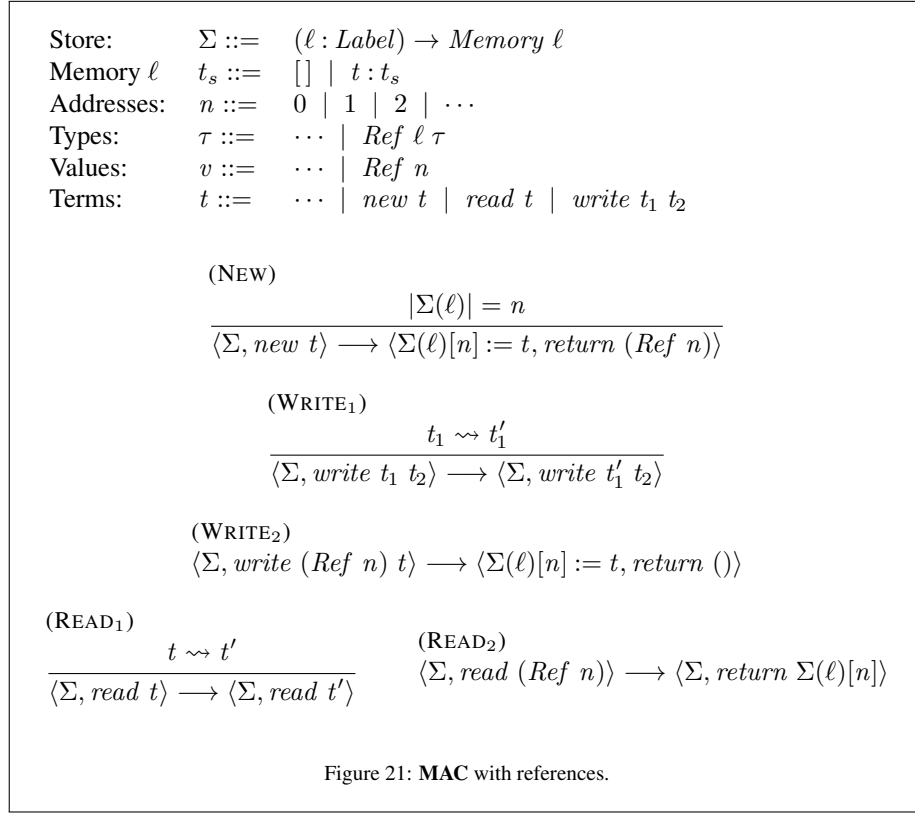
Figure 20: API for references.

Function *fetchCacheDict* takes as an extra argument a reference to a table of cached dictionaries, i.e., *Ref L (Map String [String])*. When the language *lang* dictionary is needed, the function reads the cached table (*dicts*) from the reference (*read r*) and checks if it has already been fetched (*lookup lang dicts*). If it is a hit (case *Just dict*), the dictionary is returned directly without the need of any IO operation. Otherwise (case *Nothing*), the dictionary is fetched with *fetchDict*, the result cached (*write (insert dict dicts) r*) and returned.

6.1. Calculus

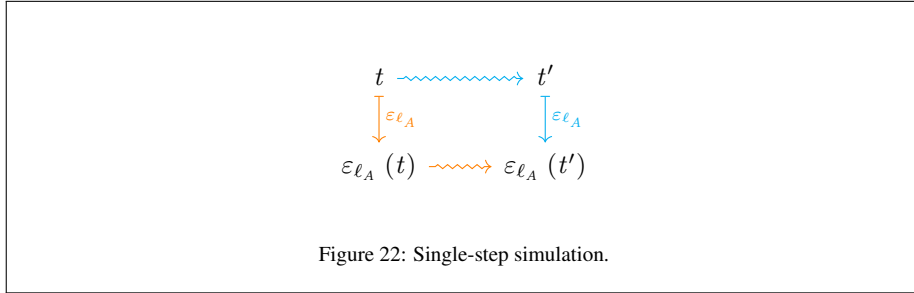
Figure 21 extends the calculus with mutable references, another feature available in **MAC**. Memory is compartmentalized into isolated labeled segments¹², one for each label of the lattice, and accessed exclusively through the store Σ . A memory in the category *Memory l* contains terms at security level l . We use the standard list interface $[], t : t_s$ and $t_s[n]$ for the empty list, the insertion of a term into an existing list and accessing the *n*th-element, respectively. We write $\Sigma(l)[n]$ to retrieve the *n*th-cell in the l -memory. The notation $\Sigma(l)[n] := t$ denotes the store obtained by performing

¹²A split memory model simplifies the proofs because allocation in one segment does not affect allocation in another. We argue why this model is reasonable and discuss alternatives in Section 7.



the update $\Sigma(\ell)[n \mapsto t]$. Secure computations create, read and write references using primitives *new*, *read* and *write* respectively. Observe that their types are restricted according to the *no read-up* and *no write-down* rules, like those of *label* and *unlabel*—see Figure 20. A reference is represented as a value $\text{Ref } n :: \text{Ref } \ell \tau$ where n is an address¹³, pointing to the n -th cell of the ℓ -memory, which contains a term of type τ . Rule [NEW] extends the ℓ -labeled memory with the new term and returns a reference to it. The notation $|t_s|$ denotes the length of a list and is used to compute the address of a new reference—memories are zero-indexed. Rule [WRITE₁] evaluates its first argument to a reference and rule [WRITE₂] overwrites the content of the memory cell pointed by the reference and returns unit. Similarly, rule [READ₂] retrieves the term stored in memory and pointed to by the reference, which is evaluated via rule [READ₁].

¹³MACs implementation of labeled reference is a simple wrapper around Haskell's type *IORef*. However, we denote references as a simple index into the labeled memory. This design choice does not affect our results.



7. Soundness

This section formally presents the security guarantees of the sequential calculus. Section 7.1 gives an overview of the proof technique (*term erasure*), Section 7.2 describes *two-steps erasure*, a novel technique that overcomes some shortcomings of vanilla *term erasure*, Section 7.3 defines the erasure function and Section 7.5 concludes with the *progress-insensitive noninterference* theorem (PINI).

7.1. Term Erasure

Term erasure is a proof technique to prove noninterference in functional programs. It was firstly introduced by Li and Zdancewic (Li and Zdancewic, 2010) and then used in a subsequent series of work on information-flow libraries (Russo et al., 2008; Stefan et al., 2011b, 2012b,a; Heule et al., 2015). The technique relies on an erasure function on terms, which we denote by ε_{ℓ_A} . This function essentially rewrites data above the attacker’s security level, denoted by label ℓ_A , to the special syntax node \bullet . Once ε_{ℓ_A} is defined, the core of the proof technique consists of proving an essential relationship about the erasure function and reduction steps. The diagram in Figure 22 highlights this intuition. It shows that erasing sensitive data from a term t and then taking a step (orange path) is the same as firstly taking a step and then erasing sensitive data (cyan path), i.e., the diagram *commutes*. If term t leaks data whose sensitivity label is above ℓ_A , then erasing all sensitive data first and then taking a step might not be the same as taking a step and then erasing secret values—the sensitive data that has been leaked into t' might remain in $\varepsilon_{\ell_A}(t')$ after all. From now on, we refer to this relationship as the *single-step simulation* between regular terms and erased ones.

7.2. Two Steps Erasure

Unfortunately, the simulation property is often a too strong requirement: there are primitives of **MAC** that do not leak intuitively, and yet there is no definition of ε_{ℓ_A} that respects the commutativity of their simulation diagram. Typically, the erasure function erases either too much and breaks simulation of some other steps, or too little and some secret data remains after the erased term steps (see Section 10.1 for several concrete examples). To formally prove that those primitives are in fact secure, we have devised a technique called *two-steps erasure*, which performs erasure in two steps—a novel approach if compared with previous papers (Stefan et al., 2011a). Rather than being

$$\varepsilon_{\ell_A}(\text{Labeled } t :: \text{Labeled } \ell_H \tau) = \begin{cases} \text{Labeled } \bullet & \text{if } \ell_H \not\sqsubseteq \ell_A \\ \text{Labeled } \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\text{label } t :: \text{MAC } \ell_L (\text{Labeled } \ell_H \tau)) = \begin{cases} \text{label } \bullet & \text{if } \ell_H \not\sqsubseteq \ell_A \\ \text{label } \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

Figure 23: Term erasure for labeled values.

a pure syntactic procedure, erasure is also performed by additional evaluation rules, triggered by special constructs introduced by the erasure function. Erasure occurs in two stages following the **orange** path in Figure 22, firstly by rewriting a problematic primitive to an ad hoc construct (along the vertical solid arrow), secondly through the reduction step of that construct (along the horizontal curly arrow). In the following, we apply two-steps erasure systematically to gain the extra flexibility needed to prove that several advanced primitives, such as *new*, *write*, *join*, satisfy *single-step simulation*.

7.3. Erasure Function

We proceed to define the erasure function for the pure calculus. Since security levels are at the type-level, the erasure function is type-driven. We write $\varepsilon_{\ell_A}(t :: \tau)$ for the erasure of term t with type τ of data not observable by the attacker. We omit the type annotation when it is either irrelevant or clear from the context. Ground values (e.g., *True*) are unaffected by the erasure function and, for most terms, the function is homomorphically applied, e.g., $\varepsilon_{\ell_A}(t_1 t_2 :: \tau) = \varepsilon_{\ell_A}(t_1 :: \tau' \rightarrow \tau) \varepsilon_{\ell_A}(t_2 :: \tau')$. Figure 23 shows the definition of the erasure functions for the interesting cases. The *content* of a resource of type *Labeled* $\ell_H \tau$ is rewritten to \bullet if the label is sensitive, i.e., it is not visible to the attacker's label ($\ell_H \not\sqsubseteq \ell_A$), otherwise it is erased homomorphically¹⁴. Similarly the erasure function rewrites the argument of *label* to \bullet , if it gets labeled with a sensitive label or otherwise erased homomorphically. Observe that this definition respects the commutativity of the diagram in Figure 22 for rule [LABEL].

Figure 24 shows the erasure function for configuration, store and memory primitives. A configuration $\langle \Sigma, t \rangle$ is erased by erasing the store Σ and by rewriting term t to \bullet , if it represents a *sensitive* computation, i.e., if term t has type *MAC* $\ell_H \tau$, where ($\ell_H \not\sqsubseteq \ell_A$), and homomorphically otherwise, see Figure 24a. It is worth pointing out that the erasure of a term $t :: \text{MAC } \ell_H \tau$, where $\ell_H \not\sqsubseteq \ell_A$ is homomorphic if the term is considered in *isolation*¹⁵, but aggressively erased to \bullet as shown in Figure

¹⁴The special term \bullet can have any type τ . We give the typing rules for the extended calculus in Figure C.52 in Appendix C.

¹⁵This is different from the conference version of this work (Vassena and Russo, 2016), where $\varepsilon_{\ell_A}(\text{MAC } \ell_H \tau :: t) = \bullet$ if $\ell_H \not\sqsubseteq \ell_A$. Erasing such terms homomorphically simplifies the formal-

$$\varepsilon_{\ell_A}(\langle \Sigma, t :: MAC \ell_H \tau \rangle) = \begin{cases} \langle \varepsilon_{\ell_A}(\Sigma), \bullet \rangle & \text{if } \ell_H \not\sqsubseteq \ell_A \\ \langle \varepsilon_{\ell_A}(\Sigma), \varepsilon_{\ell_A}(t) \rangle & \text{otherwise} \end{cases}$$

(a) Erasure for configuration.

$$\varepsilon_{\ell_A}(t_s :: Memory \ell_H) = \begin{cases} \bullet & \text{if } \ell_H \not\sqsubseteq \ell_A \\ map \ \varepsilon_{\ell_A} \ t_s & \text{otherwise} \end{cases}$$

(b) Erasure for memory.

$$\varepsilon_{\ell_A}(Ref \ n :: Ref \ell_H \tau) = \begin{cases} Ref \ \bullet & \text{if } \ell_H \not\sqsubseteq \ell_A \\ Ref \ n & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(new \ t :: MAC \ell_L (Ref \ell_H \tau)) = \begin{cases} new \ \bullet \ \varepsilon_{\ell_A}(t) & \text{if } \ell_H \not\sqsubseteq \ell_A \\ new \ \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(write \ t_1 \ t_2) = \begin{cases} write \ \bullet \ \varepsilon_{\ell_A}(t_1) \ \varepsilon_{\ell_A}(t_2 :: Ref \ell_H \tau) & \text{if } \ell_H \not\sqsubseteq \ell_A \\ write \ \varepsilon_{\ell_A}(t_1) \ \varepsilon_{\ell_A}(t_2) & \text{otherwise} \end{cases}$$

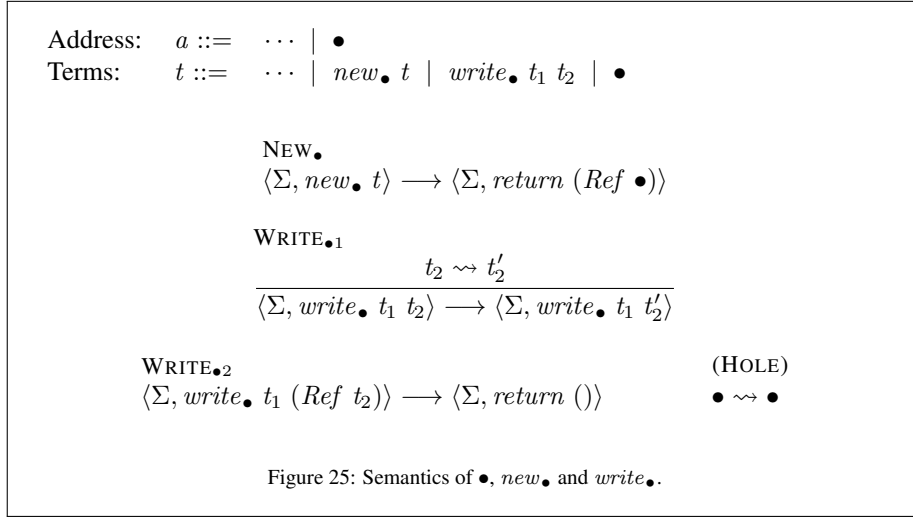
(c) Erasure for references and memory primitives.

Figure 24: Erasure for configuration, store and memory primitives.

24a if paired with a store in a configuration. Intuitively the term alone is just the *description* of a secure computation¹⁶, which can be executed only if paired with a store in a configuration. The store Σ is erased pointwise by erasing the memories at each security level, i.e., $\varepsilon_{\ell_A}(\Sigma) = \lambda \ell. \varepsilon_{\ell_A}(\Sigma(\ell))$, see Figure 24b. The erasure function collapses sensitive memories completely by rewriting them to \bullet and erase non-sensitive ones homomorphically. Figure 24c shows the erasure of references, whose address is rewritten to \bullet if sensitive, and primitive *new* and *write*, which is non-standard. Observe that these primitive perform a *write* effect and due to the *no write-down* rule they

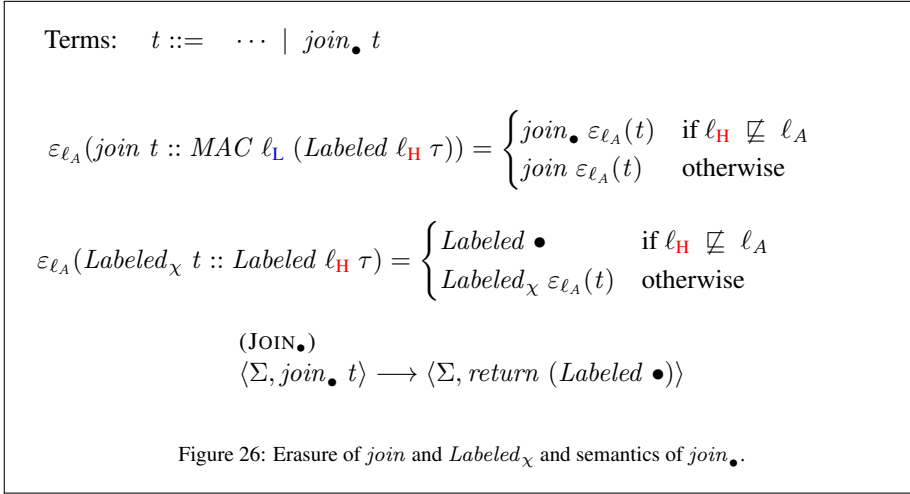
ization.

¹⁶Observe that in (Vassena and Russo, 2016) this was not the case, because rule [UNLABEL₂] and [BIND₂] were given as pure reductions (\rightsquigarrow). By separating the pure semantics from the top-level monadic semantics, we simplify the formalization of applicative functors, see Section 10.1.



can only affect memories at least as sensitive as the current secure computation. When these operations constitute a *sensitive write*, i.e., they involve memories not visible to the attacker ($\ell_H \not\sqsubseteq \ell_A$), we employ our *two-steps erasure* technique. Specifically the erasure function replaces constructs new and $write$ with special constructs new_{\bullet} and $write_{\bullet}$, whose semantics simulates that of the original terms with a *no-operation*—see Figure 25. In particular rule [NEW $_{\bullet}$] leaves the store Σ unchanged (the argument to new_{\bullet} is ignored), and returns a dummy reference with address \bullet . The same principle applies to $write_{\bullet}$. Rule [WRITE $_{\bullet,1}$] evaluates the second argument to a reference, simulating rule [WRITE $_1$] and [WRITE $_{\bullet,2}$] skips the write and just returns unit, simulating rule [WRITE $_2$]. The presence of these rules ensures that any sensitive write operation can be simulated in a lock-step fashion. Note that the semantics of new_{\bullet} and $write_{\bullet}$ correctly captures the unchanged observational power of an attacker performing *sensitive write* operations. We remark that \bullet , new_{\bullet} and $write_{\bullet}$ and their semantics rules are introduced in the calculus due to mere technical reasons (as explained above)—they are not part of the surface syntax nor **MAC**.

Figure 26 shows the erasure function for the remaining terms of the sequential calculus, that is $join$ and $Labeled_{\chi}$. Using the same technique that we have described previously, we replace $join$ with special term $join_{\bullet}$, when it is used to run a sensitive computation ($\ell_H \not\sqsubseteq \ell_A$). Erasure is then performed by means of rule [JOIN $_{\bullet}$], which immediately returns a dummy labeled value ($Labeled \bullet$) and the store unchanged. The rule captures the observational power of an attacker that runs a *terminating* sensitive computation. Observe in particular that the rule does not need to run the sensitive computation: the store can only be changed in sensitive memories (*no write-down*), which are not visible to the attacker, and the result of the computation is irrelevant—the attacker cannot unlabel it (*no read-up*), because it is marked as sensitive. What about computations that fail with an exception? In Figure 26, the erasure function not only rewrites the content of a sensitive exception to \bullet , as expected, but it also *masks* its



exceptional nature, by replacing the constructor Labeled_{χ} with Labeled , thus ensuring that rule $[\text{JOIN}_{\bullet}]$ simulates rule $[\text{JOIN}_{\chi}]$ as well. Crucially, we have the freedom of choosing this definition without breaking *simulation*, because no other construct can detect, either explicitly or implicitly, the difference. For instance, rule $[\text{UNLABEL}_{\chi}]$ operates on labeled expressions containing exceptions. In this case, if the labeled exception is not visible to the attacker, then *unlabel* must be performed in a non-visible computation as well, due to the typing rules. Operation *unlabel* then gets rewritten to \bullet and the step is then simulated by rule $[\text{HOLE}]$ instead. As a result of that, and unlike the approach taken by Stefan et al. in (Stefan et al., 2012b), there are no sensitive labeled exceptions in erased terms.

7.4. Discussion

Term Erasure. We prove the single-step simulation directly over the small-step reduction relation. Instead, other works (Li and Zdancewic, 2010; Russo et al., 2008; Stefan et al., 2011b, 2012b,a; Heule et al., 2015) prove the simulation by relating operational semantics step reductions (upper part in Figure 22) with reductions on a ℓ_A -indexed small-step relation of the form $c \longrightarrow_{\ell} \varepsilon_{\ell_A}(c')$, i.e., a relation which applies erasure at every reduction step. The reason for that is wired deeply in the dynamic nature of the enforcement. For instance, **LIO** considers labels as terms, which makes difficult to know what data is sensitive until run-time. In contrast, **MAC** does not need such an auxiliary construction because, due to its static nature, labels are not terms but rather type-level entities and therefore known before execution. In this light, our erasure function can safely erase any sensitive information found in labeled terms according to their type. Our small-step semantics satisfies type-preservation, i.e., reduction does not change types of terms, therefore labels are unaffected by execution—freeing us from the need to use a special small-step relation like \longrightarrow_{ℓ} .

Masking Sensitive Exceptions. In previous work, labeled exceptions are erased by erasing their content according to their label, but always preserving their exceptional state (Stefan et al., 2012b). In contrast, we mask sensitive exceptions in erased programs. More specifically, erasing sensitive exceptions always results in erased unexceptional values—in other words, there are no sensitive exceptions in erased programs. Note that the simulation between terms and their erased counterparts guarantees that this rewriting is *sound*. In particular sensitive exception handling routines, the only routines which can distinguish exceptional from unexceptional sensitive values, gets also erased and do not occur in erased programs.

Memory. It is known that dealing with dynamic allocation of memory makes it challenging to prove noninterference (e.g., (Banerjee and Naumann, 2005; Hedin and Sands, 2006)). One manner to tackle this technicality is by establishing a bijection between public memory addresses of the two executions we want to relate and considering equality of public terms up to such notion (Banerjee and Naumann, 2005). Instead, and similar to other work (Heule et al., 2015; Stefan et al., 2011a), we compartmentalize the memory into isolated labeled segments, one for each label of the lattice. This way, allocation in one segment does not affect the others. The fact that GHC’s memory is non-split, does not compromise our security guarantees, because references are part of **MAC**’s internals and they cannot be inspected or deallocated explicitly. However, this memory model assumes infinite memory, since reference allocation never fails. This assumption is not realistic for actual systems, where physical resources such as memory are finite.¹⁷ We conjecture that this gap between **MAC** and the model presented here, i.e., memory exhaustion, constitutes a covert channel that can be used to leak secrets with the same bandwidth as the termination covert channel (Askarov et al., 2008). In the conference version of this work (Vassena and Russo, 2016), we have explored an alternative way to prove *single-step* simulation for terms *new* and *write* consists in extending the semantics of memory operations to node \bullet , i.e., by defining $|\bullet| = \bullet$ and $\bullet[\bullet \mapsto t] = \bullet$. Thanks to two-steps erasure, we can prove simulation as we did here, without recurring to a non-standard memory semantics. A non split-memory model requires some care when proving noninterference, and in fact, we have identified problems with the proofs in manuscripts and articles related to **LIO** (Stefan et al., 2011b, 2012b). We refer interested readers to Appendix B of our conference version (Vassena and Russo, 2016) for details.

7.5. Progress-Insensitive Noninterference

The sequential calculus that we have presented satisfies *progress-insensitive noninterference*. The proof of this result is based on two fundamental properties: *single-step simulation* and *determinancy* of the small step semantics. In the following, we assume well-typed terms.

¹⁷Unrestricted access to shared resources often constitutes a covert channel in concurrent systems. The resources can be either hardware (e.g., physical memory, caches (Stefan et al., 2013) and multi-core processors) or software, such as components of the run-time system of a high-level language. These include the scheduler (Russo and Sabelfeld, 2006a,b) and the garbage collector (Pedersen and Askarov, 2017) or the language evaluation strategy such as lazy evaluation (Vassena et al., 2017; Buiras and Russo, 2013).

Proposition 1 (Single-step Simulation). *If $c_1 \longrightarrow c_2$ then $\varepsilon_{\ell_A}(c_1) \longrightarrow \varepsilon_{\ell_A}(c_2)$.*

Proof (Sketch). By induction on the reduction steps and typing judgment. Sensitive computations are simulated by transition $\langle \Sigma, \bullet \rangle \longrightarrow \langle \Sigma, \bullet \rangle$, obtained by lifting rule [HOLE] with [PURE]. Non-sensitive computations are simulated by the same rule that performs the non-erased transition, except when it involves some *sensitive write* operations, e.g., in rules [NEW, WRITE₁, WRITE₂, JOIN, JOIN_χ], which are simulated by rules [NEW_•, WRITE_{•1}, WRITE_{•2}, JOIN_•].

Proposition 2 (Determinancy). *If $c_1 \longrightarrow c_2$ and $c_1 \longrightarrow c_3$ then $c_2 \equiv c_3$.*

Proof. By standard structural induction on the reductions¹⁸.

Before stating progress-insensitive noninterference, we define low-equivalence for configurations.

Definition 1 (ℓ_A -equivalence). *Two configurations c_1 and c_2 are indistinguishable from an attacker at security level ℓ_A , written $c_1 \approx_{\ell_A} c_2$, if and only if $\varepsilon_{\ell_A}(c_1) \equiv \varepsilon_{\ell_A}(c_2)$.*

Using Proposition 1 and 2, we show that our semantics preserves ℓ_A -equivalence.

Proposition 3 (\approx_{ℓ_A} Preservation). *If $c_1 \approx_{\ell_A} c_2$, $c_1 \longrightarrow c'_1$, and $c_2 \longrightarrow c'_2$, then $c'_1 \approx_{\ell_A} c'_2$.*

By repeatedly applying Proposition 3, we prove progress-insensitive noninterference.

Theorem 1 (PINI). *If $c_1 \approx_{\ell_A} c_2$, $c_1 \Downarrow c'_1$ and $c_2 \Downarrow c'_2$, then $c'_1 \approx_{\ell_A} c'_2$.*

8. Concurrency

Every day, millions of users around the world use concurrent applications, such as email, chat rooms, social networks, e-commerce platforms etc. These services are normally designed concurrently so that multithreaded servers can handle a large number of user requests simultaneously by running multiple instances of the same application. **MAC** features *concurrency* and *synchronization variables*, which shows that the secure-by-construction programming model that we propose is possible even in a concurrent setting. The extension is non-trivial: the possibility to run simultaneous *MAC* ℓ computations provides attackers with new means to bypass security checks.

¹⁸Symbol \equiv denotes equivalence up to alpha equivalence in the calculus with named variables. In our mechanized proofs we use Bruijn indexes and we obtain syntactic equality.

$$fork :: \ell_L \sqsubseteq \ell_H \Rightarrow MAC \ell_H () \rightarrow MAC \ell_L ()$$

Figure 29: API for concurrency.

8.1. Termination Attack

In Section 7, we have proved that the sequential calculus satisfies *progress-insensitive* noninterference, a security condition that is too weak for concurrent systems. The key observation is the fact that a *non-terminating* sensitive computation at security level ℓ_H embedded in a non-sensitive one at security level ℓ_L via *join*, will suppress public side-effects that follows *join*¹⁹. Since the embedded computation is sensitive, the suppressed public events may depend on a secret, therefore revealing a bit of secret information. To illustrate this point, we present the attack in Figure 27. We assume that there exists a function $print^{MAC}$ which prints an integer on a public channel. Observe how function *leak* may suppress subsequent public events with infinite loops.

Unfortunately concurrency magnifies the bandwidth of the termination covert channel to be linear in the size (of bits) of secrets (Stefan et al., 2012a)²⁰, which permits to leak any secret systematically and efficiently. If a thread runs *leak 0 secret*, the code publishes 0 *only if* the first bit of *secret* is 0; otherwise it loops (see function *loop*) and it does not produce any public effect—see Figure 28. Similarly, a thread running *leak 1 secret* will leak the second bit of *secret*, while a thread running *leak 2 secret* will leak the third bit of it and so on. An attacker might then leak the whole secret by spawning as many threads as bits in the secret, i.e., $|secret|$, where each thread runs the one-bit attack described above and n matches the bit being leaked (e.g., $n = 0$ for the first bit, $n = 1$ for the second one, etc.).

```

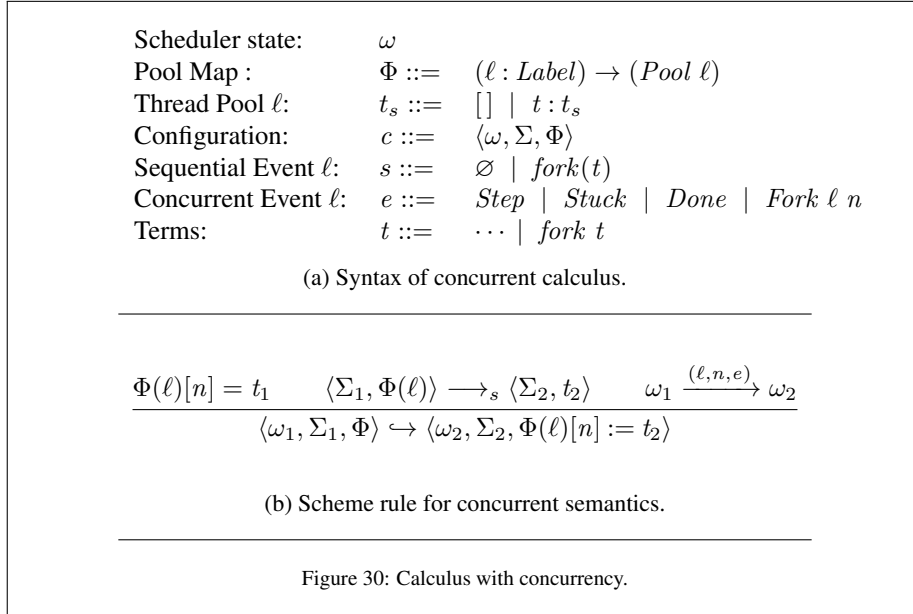
leak :: Int → Labeled H Secret → MAC L ()
leak n secret = do
  joinMAC (do bits ← unlabel secret
           when (bits !! n) loop
           return True)
  printMAC n

```

Figure 27: Termination leak.

¹⁹If the physical execution time of a program depends on the value of the secret, then an attacker with an arbitrary precise stopwatch can deduce information about the secret by timing the program. This covert channel is known as the *external timing covert channel* (Bortz and Boneh, 2007; Felten and Schneider, 2000). This article does not address the external timing covert channel, which is a harder problem and for which mitigation techniques exist (Askarov et al., 2010; Zhang et al., 2011, 2012).

²⁰Furthermore, the presence of threads introduce the *internal timing covert channel* (Smith and Volpano, 1998), a channel that gets exploited when, depending on secrets, the timing behavior of threads affect the order of events performed on public-shared resources. Since the same countermeasure closes both the internal timing and termination covert channels, we focus on the latter.



To securely support concurrency, **MAC** forces programmers to decouple *MAC* computations with sensitive labels from those performing observable side-effects—an approach also taken in LIO (Stefan et al., 2012a). As a result, non-terminating computations based on secrets cannot affect the outcome of public events.

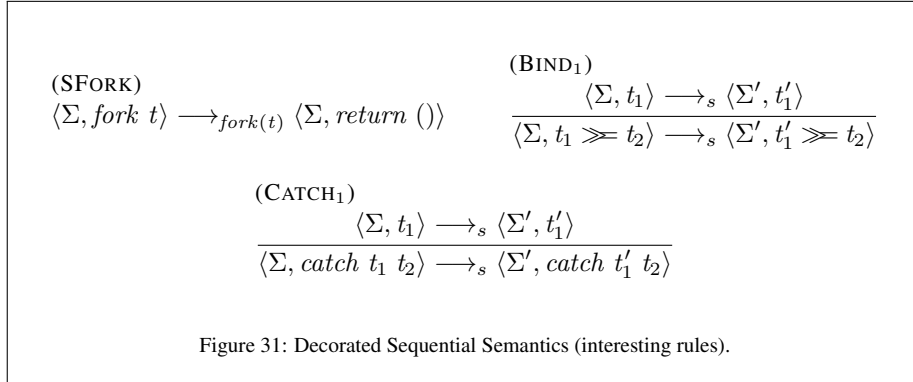
magnify :: Labeled *H* Secret \rightarrow MAC *L* ()
magnify secret =
 for [0 .. |secret|]
 ($\lambda n \rightarrow fork (leak\ n\ secret)$)

Figure 28: Attack magnification

To achieve this behavior, **MAC** replaces *join* by *fork*—see Figure 29. Informally, it is secure to spawn sensitive computations (of type *MAC* ℓ_H ()) from non-sensitive ones (of type *MAC* ℓ_L ()) because that decision depends on data at level ℓ_L , which is no more sensitive ($\ell_L \sqsubseteq \ell_H$). From now on, we call *sensitive (non-sensitive) threads* those executing *MAC* computations with a label non-observable (observable) to the attacker. In the two-point lattice, for example, threads running *MAC* *H* () computations are sensitive, while those running *MAC* *L* () are observable by the attacker.

8.2. Calculus

Figure 30 extends the calculus from Section 3 with concurrency. It introduces *global configurations* of the form $\langle \omega, \Sigma, \Phi \rangle$ composed by an *abstract* scheduler state ω , a store Σ and a pool map Φ , see Figure 30a. Threads are secure computations of type *MAC* ℓ () and are organized in isolated thread pools according to their security label. A pool t_s in the category *Pool* ℓ contains threads at security level ℓ and is accessed



exclusively through the pool map. We use the same notation for thread pools and pool maps that we have defined to manipulate and extend stores and memories. Term *fork* t spawns thread t and replaces *join* in the calculus. Without *join*, constructor *Labeled* _{χ} becomes redundant and is also removed. Our calculus includes also synchronization primitives (Russo, 2015), we refer to Appendix B for details.

Relation $c_1 \hookrightarrow c_2$ denotes that concurrent configurations c_1 steps to c_2 . Figure 30b shows the scheme rule for $c_1 \hookrightarrow c_2$ and highlights the top-level common aspects to all the rules, which we detail later on. The relation $\omega_1 \xrightarrow{(\ell, n, e)} \omega_2$ represents a transition in the scheduler, that depending on the initial state ω_1 , decides to run thread identified by (ℓ, n) , which is retrieved from the configuration $(\Phi(\ell)[n])$ and executed. Concurrent events inform the scheduler about the evolution of the global configuration, so that it can realize concrete scheduling policies and update its state accordingly. Event *Step* denotes a single sequential step, event *Fork* ℓ n informs the scheduler that the current thread has forked a new thread identified by (ℓ, n) , event *Done* is generated when a thread has terminated and event *Stuck* denotes that a thread is *stuck*, e.g., on a synchronization variable. Note that the scheduled thread determines, with its execution and with sequential event s , triggered by the decorated sequential step, i.e., $\langle \Sigma, t_1 \rangle \longrightarrow_s \langle \Sigma, t_2 \rangle$, which concurrent event e should be passed to the scheduler. Lastly, the final configuration is composed by the updated scheduler state, i.e., ω_2 , the updated memory, i.e., Σ_2 and the pool map updated with the executed thread, i.e., $\Phi(\ell)[n] := t_2$.

Decorated Semantics. Figure 31 shows the interesting rules of the decorated semantics. Rule [SFORK] is the only rule that explicitly generates event *fork*(t) and context rules [BIND₁, CATCH₁] propagate the same event generated by the premise step. All the other rules generate the empty event \emptyset . Note that, without context rules we could have given the semantics of *fork* in the concurrent semantics directly.

Concurrent Semantics. Figure 32 shows the actual semantics of the concurrent calculus, where each rule generates the appropriate event for the scheduler depending on the state of the thread scheduled and the sequential event. Concurrent rule [STEP] sends

$$\begin{array}{c}
\text{(STEP)} \\
\frac{\omega \xrightarrow{(\ell, n, \text{Step})} \omega' \quad \langle \Sigma, \Phi(\ell)[n] \rangle \longrightarrow_{\emptyset} \langle \Sigma', t' \rangle}{\langle \omega, \Sigma, \Phi \rangle \hookrightarrow \langle \omega', \Sigma', \Phi(\ell)[n] := t' \rangle} \\
\\
\text{(CFORK)} \\
\frac{\omega \xrightarrow{(\ell_L, n, \text{Fork } \ell_H m)} \omega' \quad |\Phi(\ell_H)| = m \quad \langle \Sigma, \Phi(\ell_L)[n] \rangle \xrightarrow{\text{fork}(t::\text{MAC } \ell_H ())} \langle \Sigma, t' \rangle \quad \Phi' = \Phi(\ell_H)[m \mapsto t]}{\langle \omega, \Sigma, \Phi \rangle \hookrightarrow \langle \omega', \Sigma, \Phi'(\ell_L)[n] := t' \rangle} \\
\\
\begin{array}{cc}
\text{(DONE)} & \text{(STUCK)} \\
\frac{\omega \xrightarrow{(\ell, n, \text{Done})} \omega' \quad \Phi(\ell)[n] = v}{\langle \omega, \Sigma, \Phi \rangle \hookrightarrow \langle \omega', \Sigma, \Phi \rangle} & \frac{\omega \xrightarrow{(\ell, n, \text{Stuck})} \omega' \quad \langle \Sigma, \Phi(\ell)[n] \rangle \not\rightarrow}{\langle \omega, \Sigma, \Phi \rangle \hookrightarrow \langle \omega', \Sigma, \Phi \rangle}
\end{array}
\end{array}$$

Figure 32: Concurrent Semantics

event *Step* to the scheduler, because the thread generates sequential event \emptyset , and then updates the store and the thread pool accordingly. Rule [CFORK] generates concurrent event *Fork* $\ell_H m$, because the scheduled thread, identified by label ℓ_L and number n , spawns a child thread with type $t::\text{MAC } \ell_H ()$, generating event $\text{fork}(t::\text{MAC } \ell_H ())$. Observe that the spawned thread is uniquely identified by the label ℓ_H and number m and placed in pool $\Phi(\ell_H)$ in the *free* position $m = |\Phi(\ell_H)|$. The extended pool map Φ' is lastly updated with the parent thread. In rule [DONE], $\Phi(\ell)[n] = v$ denotes that the scheduled thread is a value, i.e. the computation has terminated, then the rule sends event *Done* to the scheduler and leaves the store and pool map unchanged—terminated threads remain in pool map Φ . In rule [STUCK], the notation $\Phi(\ell)[n] \not\rightarrow$ denotes that the thread is *stuck*, i.e., it is not a value nor a redex. The scheduler is then informed by event *Stuck* and the store Σ and pool map Φ are left unchanged.

8.3. Round-robin scheduler

Figure 33 shows a round-robin scheduler with time-slot of one step, as an example of a scheduler that can be securely employed in our concurrent calculus. The state of the scheduler is a queue that tracks the identifiers of *alive* threads in the global configuration. A thread is uniquely identified by a pair consisting of a label, i.e., its security level, and a thread identifier, i.e., its position in the corresponding thread pool. The queue is concretely represented by a list of thread identifiers, whose first element identifies the next thread in the schedule. After executing one step (event *Step*), the current thread has used up its time slot and is enqueued. If the scheduled thread cannot execute (event *Stuck*), it is skipped and enqueued as well. When the current thread has terminated (event *Done*), the thread is not alive anymore and hence removed from the

$$\omega ::= (\ell, n) : \omega \mid []$$

$$(\ell, n) : \omega \xrightarrow{(\ell, n, Step)}_{RR} \omega \uparrow [(\ell, n)]$$

$$(\ell, n) : \omega \xrightarrow{(\ell, n, Stuck)}_{RR} \omega \uparrow [(\ell, n)] \quad (\ell, n) : \omega \xrightarrow{(\ell, n, Done)}_{RR} \omega$$

$$(\ell_L, n_1) : \omega \xrightarrow{(\ell_L, n_1, Fork \ell_H n_2)}_{RR} \omega \uparrow [(\ell_H, n_2), (\ell_L, n_1)]$$

Figure 33: Round-robin scheduler.

$$fmap :: (a \rightarrow b) \rightarrow Labeled \ell a \rightarrow Labeled \ell b$$

$$(\langle * \rangle) :: Labeled \ell (a \rightarrow b) \rightarrow Labeled \ell a \rightarrow Labeled \ell b$$

$$relabel :: \ell_L \sqsubseteq \ell_H \Rightarrow Labeled \ell_L a \rightarrow Labeled \ell_H a$$

Figure 34: API of Flexible labeled values.

queue. Message $(\ell_L, n_1, Fork \ell_H n_2)$ informs the scheduler that thread (ℓ_L, n_1) has spawned thread (ℓ_H, n_2) , which is then enqueued with the current thread.

9. Flexible Labeled Values

In this section we extend the API of labeled values with new operations that allow to perform *pure* (side-effect free) computations with labeled data—see Figure 34. Observe that these primitives operate on labeled data without using *label* and *unlabel*, thus avoiding incurring in the *no read-up* and *no write-down* restrictions and *irrespective of their security level*. For instance, a non-sensitive computation at security level ℓ_L can operate on sensitive labeled data at security level ℓ_H using *fmap*, without forking threads in a concurrent setting, thus introducing *flexibility* when data is processed by pure functions. We remark that, depending on the evaluation strategy of the host language (i.e. call-by-value or call-by-name), a naive implementation of these primitives is vulnerable to leaks via non-termination—we elaborate on this point later, in Section 9.3. Section 9.1 gives a broad description of these primitives, Section 9.2 shows their flexibility with an example, and Section 9.3 formalizes them in our calculus.

9.1. Functors and Relabeling

Intuitively, a functor is a container-like data structure which provides a method called *fmap* that applies (maps) a function over its contents, while preserving its structure. Lists are the most canonical example of a functor data-structure. In this case, *fmap* corresponds to the function *map*, which applies a function to each element of a

list, e.g. $fmap (+1) [1, 2, 3] \equiv [2, 3, 4]$. A functor structure for labeled values allows to manipulate sensitive data without the need to explicitly extract it—see Figure 34. For instance, $fmap (+1) d$, where $d :: Labeled\ H\ Int$ stores the number 42, produces the number 43 as a sensitive labeled value.

To aggregate data at possibly different security levels **MAC** provides primitives *relabel* and $(\langle * \rangle)$. Primitive *relabel* upgrades the security level of a labeled value, which is useful to “lift” data to an upper bound of all the data involved in a computation prior to combining them. Infix operator $(\langle * \rangle)$ supports function application within a labeled value, i.e. it allows to feed functions wrapped in a labeled value ($Labeled\ \ell\ (a \rightarrow b)$) with arguments also wrapped ($Labeled\ \ell\ a$), where aggregated results get wrapped as well ($Labeled\ \ell\ b$).

Discussion. In functional programming, operator $(\langle * \rangle)$ is part of the *applicative functors* (Mcbride and Paterson, 2008) interface, which in combination with *fmap*, is used to map functions over functors. Note that if labeled values were full-fledged applicative functors, our API would also include the primitive $pure :: a \rightarrow Labeled\ \ell\ a$. This primitive brings arbitrary values into labeled values, which might break the security principles enforced by **MAC**. Instead of *pure*, **MAC** centralizes the creation of labeled values in the primitive *label*. Observe that, by using *pure*, a programmer could write a computation $m :: MAC\ H\ (Labeled\ L\ a)$ where the *created* labeled information is sensitive rather than public. We argue that this situation ignores the no-write down principle, which might bring confusion among users of the library. More importantly, freely creating labeled values is not compatible with the security notion of *cleareance*, where secure computations have an upper bound on the kind of sensitive data they can observe and generate. This notion becomes useful to address certain covert channels (Zeldovich et al., 2006) as well as poison-pill attacks (Hritcu et al., 2013). While **MAC** does not yet currently support *cleareance*, it is an interesting direction for future work.

9.2. Examples

The functor API of labeled values, i.e., *fmap*, is a handy tool that functional programmers use to code simple concise functions elegantly. In Figure 35, the 1-line function *isShort* checks whether the password is weak because it

```
isShort :: Labeled H String → Labeled H Bool
isShort = fmap (λpwd → |pwd| ≤ 5)
```

Figure 35: A *pure* computation on a password.

is too short. In the anonymous function, *pwd* is the *unlabeled* password, and the expression $|pwd| \leq 5$ checks if the password contains less than 5 characters. Observe that what the function computes is an attribute of the password, therefore it should be considered *sensitive*. The API of *fmap* ensures that by preserving the label of the labeled argument, i.e., $Labeled\ H\ String$, in the resulting labeled value, i.e., $Labeled\ H\ Bool$. Compare the program in Figure 35 with the homonym program in Figure 36 written without *fmap*, but using *join* instead. Firstly, note that the imperative program has a different signature: it must necessarily involve *MAC* computation in order to perform

```

isShort :: Labeled H String → MAC L (Labeled H Bool)
isShort lpwd = do
  join (do
    pwd ← unlabel lpwd
    return (|pwd| ≤ 5))

```

Figure 36: Without *fmap pure* sensitive computations have an *impure* type.

$isWeak :: Labeled\ L' (String \rightarrow Bool).$


(a) Third party API.

```

f :: Labeled H String → Labeled H Bool
f pwd = (relabel isWeak) ⟨*⟩ pwd

```

(b) Embedding mistrusted code.

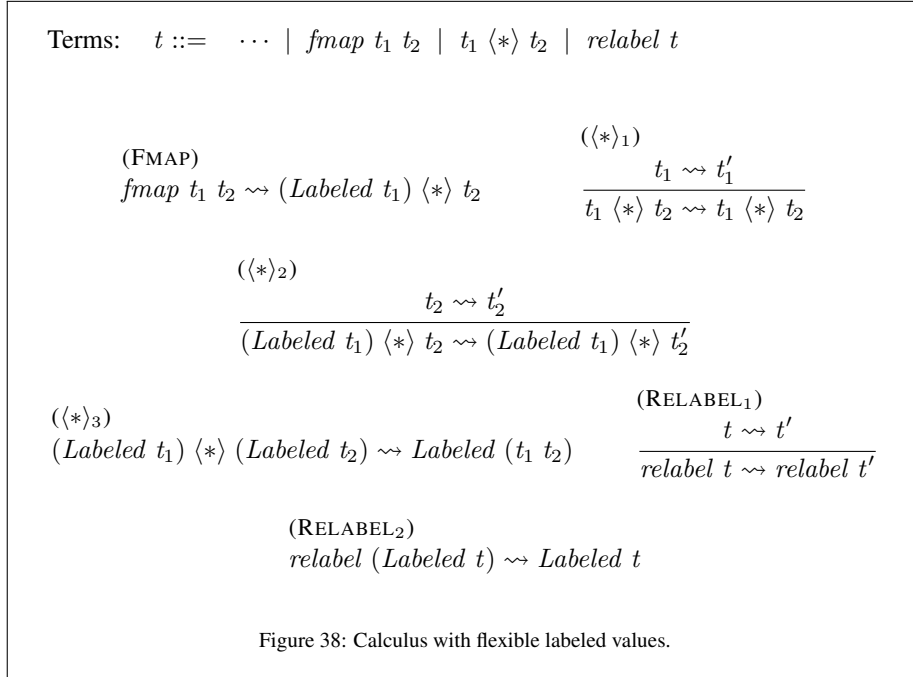


(c) 3-Points Lattice.

Figure 37: Combining heterogeneously labeled data.

unlabel. Since the password *lpwd* is sensitive, i.e., it has type *Labeled H String*, only a sensitive computation can unlabel it. Then, the program employs *join* to convert the sensitive computation into a sensitive labeled value, which then gets wrapped in a non-sensitive computation, i.e., *MAC L (Labeled H Bool)*. In a concurrent setting, where *join* is not available, the whole program must be completely restructured, because threads have type *MAC H ()* and may not return any other result in a non-sensitive computation.

The strength of a password is often estimated by combining several syntactic aspects, such as its length or the presence and number of special characters and digits. Suppose now that some third-party API function provides such syntactic checks in the form of a **MAC** labeled *pure* function *isWeak*, see Figure 37a. The type system guarantees that the function is secure, because it has type *String → Bool*, however the third party has labeled it with its own label *L'*, because it wants to strictly control who can use it and under what terms. In order to keep the code of our password-checker isolated from that of the third party, while still providing functionality to the user, we incorporate the new label *L'* into the system and modify the lattice as shown in Figure 37c. The lattice reflects our mistrust over the third-party code by making *L* and *L'* incomparable elements. Thanks to **MAC**'s security guarantees, we can safely run third-party mistrusted code, i.e., *isWeak*, with the user's secret password, as shown in Figure 37b. In particular *relabel* upgrades the function to *isWeak* to security level *H* (observe that $L' \sqsubseteq H$ in the lattice), and then applies the function to the password (*pwd*) using the applicative functor operator, i.e., $\langle * \rangle$, which protects the final result with label *H*.



9.3. Calculus

In Figure 38, we extend our calculus with the primitives for flexible manipulation of labeled values, discussed in the previous section. Firstly we add terms $fmap\ t_1\ t_2$, $t_1 \langle * \rangle t_2$ and $relabel\ t$, whose types correspond to those given in Figure 34. Primitive $fmap$ is implemented in terms of $\langle * \rangle$ in rule [FMAP], where the function is simply lifted to labeled value (every applicative functor is also a functor). Rules $[\langle * \rangle_1, \langle * \rangle_2]$ evaluate the first and second argument to a labeled value respectively, which are then combined by rule $[\langle * \rangle_3]$, which applies the function to the argument and wraps the result in a labeled value. Rule [RELABEL₁] evaluates its argument to weak-head normal form and rule [RELABEL₂] upgrades its label. Observe that since labels are types $relabel$ leaves the content of $Labeled$ unchanged. We remark that these primitives are secure both in the concurrent and sequential calculus, where their semantics must be adjusted to handle exceptional values as well, i.e., constructor $Labeled_\chi$, which is not present in the concurrent calculus. We refer to Appendix A for more details.

Discussion. The API of flexible labeled values shown in Figure 34 might seem insecure at first sight. In particular, it might be counter-intuitive that a *public* computation might be able to manipulate a *secret* with an *arbitrary* function without introducing potential leaks. Figure 39 shows an attack that attempts to leak via *non-termination* the n -th bit of a secret. Function $leak$ applies function $loopOn$ on the secret using $fmap$ and then performs a *non-sensitive* side-effect, i.e., $publish\ n$, which outputs the number n on a public channel. Interestingly, depending on the evaluation strategy of

```

leak :: Int → Labeled H Secret → MAC L ()
leak n secret = let result = fmap loopOn secret in publish n
  where loopOn = λbits → if (bits !! n) then loop else bits

```

Figure 39: Function *leak* attempts to leak the n -th bit of *secret*.

the language, the attack might succeed. Specifically, under a *call-by-value* evaluation strategy, function *loopOn* passed to *fmap* is eagerly applied to the secret, which might introduce a loop depending on the value of the n -th bit of the secret suppressing the subsequent *public* action *publish n*. Under a *call-by-name* evaluation strategy, however, function *loopOn* does not get immediately evaluated since *result* is not needed for computing *publish n*. Therefore, *publish n* gets executed independently of the value of the secret, i.e., no termination leaks are introduced. Instead, *loopOn* gets evaluated when and only if *result* is *unlabeled* and its content inspected—something that is possible only in a computation at security level at least as sensitive as *H* because of the *no-read up* policy, where it is secure to do so. We remark that it is possible to close this termination channel under a call-by-value semantics by defining *Labeled* with an explicit suspension, e.g. **data** *Labeled* $\ell a = \text{Labeled } (() \rightarrow a)$, and corresponding forcing operation, so that *fmap* behaves lazily as desired.

10. Soundness of Concurrent Calculus

The concurrent calculus that we have presented satisfies *progress sensitive noninterference*. Section 10.1 extends the erasure function for the concurrent calculus and for flexible labeled values. To obtain a parametric proof of noninterference, we assume certain properties about the scheduler. Specifically, our proof is valid for deterministic schedulers which fulfill progress and noninterference themselves, i.e., schedulers cannot leverage sensitive information in threads to determine what to schedule next. Section 10.2 formalizes the requirements for such *suitable* schedulers. In Section 10.3 we prove a scheduler-parametric progress-sensitive noninterference theorem for our calculus and we constructively obtain a proof that **MAC** is secure with a round-robin scheduler by simply instantiating our main theorem.

10.1. Erasure Function

Figure 40 shows the erasure function for the concurrent calculus. A concurrent configuration $\langle \omega, \Sigma, \Phi \rangle$ is erased by erasing each component, where the erasure of the scheduler state ω is scheduler specific (Figure 40a). Similarly to store Σ , pool map Φ is erased pointwise, i.e., $\varepsilon_{\ell_A}(\Phi) = \lambda \ell. \varepsilon_{\ell_A}(\Phi(\ell))$, and sensitive thread pools are rewritten to \bullet and erased homomorphically otherwise, just like memories (see Figure 40b). Observe that primitive *fork* performs a *write* effect because it adds a new thread to a thread pool, therefore we employ our *two-steps erasure* technique, just like we did for memory primitives. Specifically, the erasure function replaces *fork* with *fork_•*, whenever it spawns a *sensitive* thread, which would write to a *sensitive* thread pool ($\ell_H \not\sqsubseteq \ell_A$), see

$$\varepsilon_{\ell_A}(\langle \omega, \Sigma, \Phi \rangle) = \langle \varepsilon_{\ell_A}(\omega), \varepsilon_{\ell_A}(\Sigma), \varepsilon_{\ell_A}(\Phi) \rangle$$

(a) Erasure for concurrent configuration.

$$\varepsilon_{\ell_A}(t_s :: Pool \ell_H) = \begin{cases} \bullet & \text{if } \ell_H \not\sqsubseteq \ell_A \\ map \varepsilon_{\ell_A} t_s & \text{otherwise} \end{cases}$$

(b) Erasure for thread pool.

$$\varepsilon_{\ell_A}(fork t) = \begin{cases} fork_{\bullet} \varepsilon_{\ell_A}(t :: MAC \ell_H ()) & \text{if } \ell_H \not\sqsubseteq \ell_A \\ fork \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

(c) Erasure of fork.

$$\varepsilon_{\ell_A}(fork(t :: MAC \ell_H ())) = \begin{cases} fork_{\bullet}(\varepsilon_{\ell_A}(t)) & \text{if } \ell_H \not\sqsubseteq \ell_A \\ fork(\varepsilon_{\ell_A}(t)) & \text{otherwise} \end{cases}$$

(d) Erasure for *sequential* fork event.

$$\varepsilon_{\ell_A}(Fork \ell_H n) = \begin{cases} Step & \text{if } \ell_H \not\sqsubseteq \ell_A \\ Fork \ell_H n & \text{otherwise} \end{cases}$$

(e) Erasure for *concurrent* fork event.

Figure 40: Erasure function for concurrent calculus.

Figure 40c. Sequential fork-events are erased similarly in order to ensure simulation, i.e., the erasure function rewrites $fork(t)$ to $fork_{\bullet}(\varepsilon_{\ell_A}(t))$ when t is sensitive—see Figure 40d. Sequential event \emptyset is not affected by the erasure function. The erasure function masks spawning sensitive threads from the scheduler as well by erasing concurrent events accordingly (Figure 40e). In this case it rewrites event $Fork \ell_H n$ to $Step$ whenever $\ell_H \not\sqsubseteq \ell_A$ —the other events are not affected by the erasure function. In the sequential calculus $fork_{\bullet}$ is reduced by rule [SFORK $_{\bullet}$], defined in Figure 41, which simulates the decorated reduction of $fork$. A new concurrent rule [CFORK $_{\bullet}$]

Seq. Effect: $s ::= \dots \mid \text{fork}_\bullet(t)$
 Terms: $t ::= \dots \mid \text{fork}_\bullet t$

$$\text{(SFORK}_\bullet\text{)} \\ \langle \Sigma, \text{fork}_\bullet t \rangle \longrightarrow_{\text{fork}_\bullet(t)} \langle \Sigma, \text{return } () \rangle$$

$$\text{(CFORK}_\bullet\text{)} \\ \frac{\omega \xrightarrow{(\ell_L, n, \text{Step})} \omega' \quad \langle \Sigma, \Phi(\ell_L)[n] \rangle \longrightarrow_{\text{fork}_\bullet(t)} \langle \Sigma, t' \rangle}{\langle \omega, \Sigma, \Phi \rangle \hookrightarrow \langle \omega', \Sigma, \Phi(\ell_L)[n] := t' \rangle}$$

Figure 41: Sequential and concurrent semantics of fork_\bullet .

$$\varepsilon_{\ell_A}(\text{fmap } t_1 t_2 :: \text{Labeled } \ell_H \tau) = \begin{cases} \text{fmap}_\bullet \varepsilon_{\ell_A}(t_1) \varepsilon_{\ell_A}(t_2) & \text{if } \ell_H \not\sqsubseteq \ell_A \\ \text{fmap } \varepsilon_{\ell_A}(t_1) \varepsilon_{\ell_A}(t_2) & \text{otherwise} \end{cases}$$

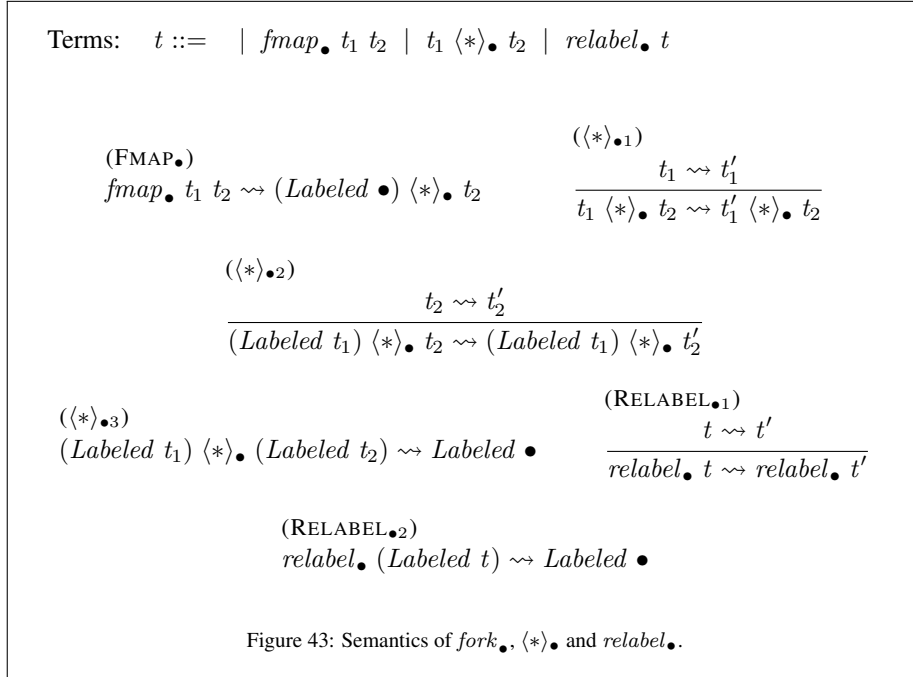
$$\varepsilon_{\ell_A}(t_1 \langle * \rangle t_2 :: \text{Labeled } \ell_H \tau) = \begin{cases} \varepsilon_{\ell_A}(t_1) \langle * \rangle_\bullet \varepsilon_{\ell_A}(t_2) & \text{if } \ell_H \not\sqsubseteq \ell_A \\ \varepsilon_{\ell_A}(t_1) \langle * \rangle \varepsilon_{\ell_A}(t_2) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\text{relabel } t :: \text{Labeled } \ell_H \tau) = \begin{cases} \text{relabel}_\bullet \varepsilon_{\ell_A}(t) & \text{if } \ell_H \not\sqsubseteq \ell_A \\ \text{relabel } \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

Figure 42: Erasure of flexible labeled values.

detects the sequential event $\text{fork}_\bullet(t)$ and *skips* spawning the thread, i.e., it does not insert it in the thread pool, and sends concurrent event *Step* to the scheduler, therefore simulating precisely rule [CFORK] when a non-sensitive thread of type $\text{MAC } \ell_L ()$ forks a sensitive thread $\text{MAC } \ell_H ()$.

Context-Aware Erasure Function. A common challenge when reasoning about security of IFC libraries is that the sensitivity of a term may depend on context where they are used. Consider for instance the primitive *relabel*, which upgrades the security level of a labeled term. A public number, e.g., $\text{Labeled } 42 :: \text{Labeled } L \text{ Int}$, should be treated as secret when in the context of relabeling, e.g., $\text{relabel } (\text{Labeled } 42) :: \text{Labeled } H \text{ Int}$. Doing otherwise, i.e., erasing the term homomorphically, breaks simulation because sensitive data produced by *relabel* remains *after* erasure. For example $\text{relabel } (\text{Labeled } 42)$ is homomorphically erased to $\text{relabel } \varepsilon_L(\text{Labeled } 42 :: \text{Labeled } L \text{ Int})$ which reduces on the **orange** path in Figure 22 to $\text{Labeled } 42 \neq$



$Labeled \bullet$, obtained on the **cyan** path by $\varepsilon_L(Labeled\ 42 :: Labeled\ H\ Int)$, thus breaking commutativity of rule [RELABEL₂].

Then, one might be tempted to stretch the definition of the erasure function to accommodate for the problematic cases shown above. Unfortunately, this approach does not work, because it will necessary break simulation for other cases. We support this statement by showing that this is the case for any arbitrary erasure function that is suitable for $relabel\ t :: Labeled\ H\ \tau$, where $t :: Labeled\ L\ \tau$. Observe that we need a different behavior for our erasure function for public labeled values when embedded in $relabel$, which we will capture in a different *auxiliary* erasure function ε'_L . Suppose we defined $\varepsilon_L(relabel\ t :: Labeled\ H\ \tau) = relabel\ \varepsilon'_L(t :: Labeled\ L\ \tau)$, for some suitable ε'_L that exhibits the desired behavior, e.g., $\varepsilon'_L(Labeled\ 42 :: Labeled\ L\ Int) = Labeled\ \bullet$. Alas, while this definition respects simulation for step [RELABEL₂], introducing a *different* erasure function in a *context-sensitive* way is fatal for simulation of beta reductions. More precisely, the original erasure function is *no longer homomorphic over substitution*, i.e., $\varepsilon_{\ell_A}([x / t_1] t_2) \neq [x / \varepsilon_{\ell_A}(t_1)] \varepsilon_{\ell_A}(t_2)$ —an essential property of the erasure function (Li and Zdancewic, 2010; Russo et al., 2008; Stefan et al., 2011b, 2012b; Heule et al., 2015), without which step [BETA] does not commute anymore. Essentially, function ε_{ℓ_A} is oblivious to the context in which some term will be substituted inside the body of a function, thus breaking simulation. As a counterexample, consider term $(\lambda x.relabel\ x)\ t$, which is erased *homomorphically*, that is $(\lambda x.relabel\ x)\ \varepsilon_L(t)$, and then beta-reduces on the **orange** path to $relabel\ \varepsilon_L(t)$. On the **cyan** path term $(\lambda x.relabel\ x)\ t$ beta-reduces to $relabel\ t$ and then is *context-*

sensitively erased to *relabel* $\varepsilon'_L(t)$. Observe that *relabel* $\varepsilon_L(t) \not\equiv \text{relabel } \varepsilon'_L(t)$ in general because ε'_L captures a different behavior than that exposed by ε_L , specifically for public labeled values, e.g., when $t = \text{Labeled } 42 :: \text{Labeled } L \text{ Int}$. To the best of our knowledge, this work is the first to point out this issue. Furthermore, we identify problematic cases in the formalization of previous work on **LIO** (Stefan et al., 2011b, 2012a) which lead to breaking the one-step simulation—see details in Appendix A of (Vassena and Russo, 2016). By using the *two-step erasure* technique, we can craft a sound erasure function that is *homomorphic over substitution* and is *context-aware*. The erasure function replaces *relabel* with *relabel* $_{\bullet}$, rule [RELABEL $_{\bullet 1}$] simulates rule [RELABEL $_1$] and rule [RELABEL $_{\bullet 2}$] performs *context-sensitive* erasure by producing *Labeled* \bullet , see Figure 43. Even though the actual erasure is done by rule [RELABEL $_{\bullet 2}$], we still have to erase the *argument* of *relabel*, or else the erasure function would not be *homomorphic over substitution*. Simulation of the context rule [RELABEL $_{\bullet}$] follows then by inductive hypothesis.

Primitive $\langle * \rangle$ raises a similar problem. To illustrate this point, consider the term $(\text{Labeled } t_1) \langle * \rangle (\text{Labeled } t_2)$ of type *Labeled* $H \text{ Int}$, which reduces to *Labeled* $(t_1 t_2)$ according to rule [$\langle * \rangle_3$].²¹ Following the **orange** path and applying the erasure function *homomorphically*, we get that $\varepsilon_L(\text{Labeled } t_1) \langle * \rangle \varepsilon_L(\text{Labeled } t_2)$, that is $(\text{Labeled } \bullet) \langle * \rangle (\text{Labeled } \bullet)$ which reduces to *Labeled* $(\bullet \bullet) \not\equiv \text{Labeled } \bullet$, obtained instead by first reducing the term and then erasing following the **cyan** path. Observe that rule [$\langle * \rangle_3$] produces a function application within a *Labeled* constructor, therefore it cannot possibly commute for sensitive labeled values, which always rewrite the content of a labeled value to \bullet . We then prove simulation using *two-steps erasure* again. Specifically, the erasure function replaces $\langle * \rangle$ with $\langle * \rangle_{\bullet}$, see Figure 42, and erasure is then performed by means of its semantics rules [$\langle * \rangle_{\bullet 1}$, $\langle * \rangle_{\bullet 2}$, $\langle * \rangle_{\bullet 3}$], listed in Figure 43, which simulate rules [$\langle * \rangle_1$, $\langle * \rangle_2$, $\langle * \rangle_3$] respectively. Observe that [$\langle * \rangle_{\bullet 3}$] ignores the content of the labeled values and simply yields *Labeled* \bullet to enforce the simulation property. Since *fmap* is defined in terms of $\langle * \rangle$, we likewise replace it with new node *fmap* $_{\bullet}$ and give its semantics in terms of $\langle * \rangle_{\bullet}$, see rule [FMAP $_{\bullet}$]. We remark that *fmap* $_{\bullet}$, $\langle * \rangle_{\bullet}$ and *relabel* $_{\bullet}$ and their semantics rules are introduced in the calculus as a device to prove *simulation* (they only occur in *erased* programs), they are not part of the surface syntax nor **MAC**.

10.2. Scheduler Requirements

We take advantage of the level of abstraction of our concurrent semantics and make our proof parametric in the scheduler state and its semantics. For this reason, we study what are the sufficient requirements of a scheduler to guarantee PSNI in our calculus. We evaluate our characterization of schedulers by formalizing a round-robin sched-

²¹In our conference version (Vassena et al., 2016), rule [$\langle * \rangle_3$] raises a problem also for public labeled values, because the erasure function is not homomorphic over function application, in particular $\varepsilon_L(t_1 t_2 :: \text{MAC } H \tau) = \bullet \not\equiv \varepsilon_L(t_1) \varepsilon_L(t_2)$. To avoid this problem, we replace function application with substitution, i.e. $(\text{Labeled } (\lambda x.t_1)) \langle * \rangle (\text{Labeled } t_2) \rightsquigarrow \text{Labeled } (t_1 [x / t_2])$, at the price of having a non-standard stricter semantics for $\langle * \rangle$. The erasure function presented here is homomorphic over function application and the semantics of $\langle * \rangle$ is standard.

uler, similar to that used by GHC’s run-time system (Marlow, 2013), and show that it satisfies the requirements listed in this section.

Our proof is valid for schedulers which are (i) deterministic, (ii) fulfill a restricted variant of single-step simulation from Figure 22, i.e., schedulers may not leverage on sensitive information to determine what observable thread should be scheduled next, (iii) do not leak secret information when scheduling a sensitive threads and (iv) guarantee progress of observable threads, i.e., execution of observable threads cannot be indefinitely deferred by sensitive ones. In the following, we use labels ℓ_L and ℓ_H to denote a security level that is visible resp. invisible to the attacker, i.e., $\ell_L \sqsubseteq \ell_A$ and $\ell_H \not\sqsubseteq \ell_A$. Furthermore, we call a scheduler step that runs a *non-sensitive* thread, e.g., $\omega_1 \xrightarrow{(\ell_L, n, e)} \omega_2$, *public* or *low* step. Similarly we refer to a run of a *sensitive* thread, e.g., $\omega_1 \xrightarrow{(\ell_H, n, e)} \omega_2$, as *secret* or *high* step. We formally characterize schedulers for which our security guarantees apply.

Requirement 1.

- i) **Determinacy:** if $\omega_1 \xrightarrow{(\ell, n, e)} \omega_2$ and $\omega_1 \xrightarrow{(\ell', n', e)} \omega'_2$, then $\ell \equiv \ell'$, $n \equiv n'$ and $\omega_2 \equiv \omega'_2$.
- ii) **Restricted Simulation:** if $\omega_1 \xrightarrow{(\ell_L, n, e)} \omega_2$ then $\varepsilon_{\ell_A}(\omega_1) \xrightarrow{(\ell_L, n, \varepsilon_{\ell_A}(e))} \varepsilon_{\ell_A}(\omega_2)$.
- iii) **No Observable Effect:** if $\omega_1 \xrightarrow{(\ell_H, n, e)} \omega_2$ then $\omega_1 \approx_{\ell_A} \omega_2$.
- iv) **Progress:** If $\omega_1 \xrightarrow{(\ell_L, n, e)} \omega'_1$ and $\omega_1 \approx_{\ell_A} \omega_2$ then ω_2 will schedule thread (ℓ_L, n) eventually.

Observe that determinacy of the scheduler is essential for determinacy of the concurrent semantics—after all, the scheduler state is part of the concurrent configuration. As it is expected from the concurrent calculus, we assume that the abstract scheduler satisfies a variant of the single-step simulation restricted to low steps²². “No observable effect”, i.e., Requirement (iii), ensures that high steps do not leak sensitive information in the scheduler state—we extend ℓ_A -equivalence to scheduler states, that is $\omega_1 \approx_{\ell_A} \omega_2$ if and only if $\varepsilon_{\ell_A}(\omega_1) \equiv \varepsilon_{\ell_A}(\omega_2)$. Observe that the erasure function of the scheduler state is scheduler specific, and thus we leave it unspecified. Requirement (iv) avoids revealing sensitive data by observing progress of non-sensitive threads via public events. Intuitively, a concurrent program might reveal sensitive information by forcing a sensitive thread to induce starvation of a non-sensitive thread, thus potentially suppressing subsequent public events. The formal definition of *eventually* is technically interesting. Since we wish to make our proof modular, our model is parametric in the scheduler, which is considered in isolation from the thread pool. In this situation, we cannot predict how long the high threads are going to run, because the scheduler is decoupled from the thread pool. We overcome this technicality by indexing the ℓ_A -equivalence relation between scheduler states. We then use the indexes to encode a

²²Different to our conference version (Vassena and Russo, 2016), we do not require lock-step simulation for high scheduler steps, i.e., when $\ell_H \not\sqsubseteq \ell_A$, for which is instead sufficient to show indistinguishability. This choice gives the same security guarantees and simplifies the formalization of a non-interfering scheduler.

single-step progress principle, i.e., Requirement 2 (explained below), and to exclude starvation, by making the progress principle a well-founded induction principle, i.e., Requirement 3 (explained below).

Definition 2 (Annotated Scheduler ℓ_A -equivalence). *Two states are (i, j) - ℓ_A -equivalent, written $\omega_1 \approx_{\ell_A}^{(i,j)} \omega_2$ if and only if $\omega_1 \approx_{\ell_A} \omega_2$ and i and j are upper bounds over the number of sensitive threads scheduled before the next common non-sensitive thread in ω_1 and ω_2 , respectively.*

The relation $\omega_1 \approx_{\ell_A}^{(i,j)} \omega_2$ captures an alignment measure of two ℓ_A -equivalent states and how close they are to schedule the next common non-sensitive thread. Informally, our noninterference proof excludes *starvation* of observable threads, that can leak information to the attacker, by ensuring that two ℓ_A -equivalent schedulers will eventually align and schedule the same non-sensitive thread, regardless of how the global configuration evolves. Specifically, our calculus requires that the indexes in $\omega_1 \approx_{\ell_A}^{(i,j)} \omega_2$ strictly decreases after every reduction. We capture the interplay between the (i, j) - ℓ_A -equivalent relationship and the evolution of schedulers by establishing unwinding-like conditions (Goguen and Meseguer, 1984).


Requirement 2 (Progress). *Given $\omega_1 \xrightarrow{(\ell_L, n, e)} \omega'_1$, and $\omega_1 \approx_{\ell_A}^{(i,j)} \omega_2$ then:*

- *If $j = 0$, then $\forall e' \exists \omega'_2: \omega_2 \xrightarrow{(\ell_L, n, e')} \omega'_2$.*
- *If $j > 0$, then there exists ℓ_H, n' such that $\forall e' \exists \omega'_2: \omega_2 \xrightarrow{(\ell_H, n', e')} \omega'_2$*

If a scheduler runs a public thread, then a (i, j) - ℓ_A -equivalent scheduler runs *at most* j secret threads before the same public thread. In particular, if $j = 0$ then the two schedules align and the threads generate ℓ_A -equivalent events²³, otherwise a secret thread is run ($j > 0$). In the second case the scheduler cannot predict what event will be triggered by thread (ℓ_H, n') , therefore, as a conservative approximation, the step may involve *any* possible event e' , which in addition determines the final state ω'_2 . Conceptually, by repeatedly applying Requirement 2, Requirement (iii) and by transitivity of \approx_{ℓ_A} , we could build the chain of high steps that precedes the common low-step. However, this recursion scheme is not well founded in general, because it does not exclude starvation, e.g., for *non-preemptive* schedulers (Heule et al., 2015). The following requirement guarantees instead that such chain is finite, i.e., that public threads cannot starve indefinitely due to secret threads.

Requirement 3 (No Starvation). *Given $\omega_1 \xrightarrow{(\ell_L, n, e)} \omega'_1$, $\omega_2 \xrightarrow{(\ell_H, n', e')} \omega'_2$, such that $\omega_1 \approx_{\ell_A}^{(i,j)} \omega_2$, then there exist j' such that $j' < j$ and $\omega'_1 \approx_{\ell_A}^{(i,j')} \omega'_2$.*

²³In our conference version (Vassena and Russo, 2016), the requirement expects the same event e in the other step, which is too strict. Intuitively an event *Fork* $\ell_H n$ contains a bit of secret information, namely the number n of secret threads, which could differ in the other run. It follows from *restricted simulation*, i.e., Requirement 1.ii, that the two events are in fact ℓ_A -equivalent, i.e., $e \approx_{\ell_A} e'$, defined as $\varepsilon_{\ell_A}(e) \equiv \varepsilon_{\ell_A}(e')$. Note that ℓ_A -equivalence captures this scenario precisely: *Fork* $\ell_H n \approx_{\ell_A}$ *Fork* $\ell_H n'$, because $\varepsilon_{\ell_A}(\text{Fork } \ell_H n) \equiv \varepsilon_{\ell_A}(\text{Fork } \ell_H n') \equiv \text{Step}$.



$$\begin{array}{c}
 \boxed{\boxed{\approx_{\ell_A}^{(0,0)} \boxed{\quad}}} \quad \frac{\omega_1 \approx_{\ell_A}^{(i,j)} \omega_2}{(\ell_L, n) : \omega_1 \approx_{\ell_A}^{(0,0)} (\ell_L, n) : \omega_2} \quad \frac{\omega_1 \approx_{\ell_A}^{(i,j)} \omega_2}{(\ell_H, n) : \omega_1 \approx_{\ell_A}^{(i+1,j)} \omega_2} \\
 \\
 \frac{\omega_1 \approx_{\ell_A}^{(i,j)} \omega_2}{\omega_1 \approx_{\ell_A}^{(i,j+1)} (\ell_H, n) : \omega_2}
 \end{array}$$

Figure 45: Annotated ℓ_A -equivalence (Round-robin).

Intuitively, the combination of Requirement 2 and 3 ensures that the two schedules will align *eventually*²⁴. Figure 44 highlights this intuition. The colored scheduler steps denote running either a secret (red for ℓ_H) or a public (blue for ℓ_L) thread respectively and the dashed line links ℓ_A -equivalent states. Given two initial scheduler states such that $\omega_0 \approx_{\ell_A}^{(i,j)} \omega$, where ω runs a public thread, *progress*, i.e., Requirement 2, guarantees that ω_0 steps to ω_1 , running a secret thread. By Requirement 3, it follows that $\omega_1 \approx_{\ell_A}^{(i,j')}$ ω , where $j' < j$. After repeating this mechanism at most j times (j is strictly smaller after each step), we obtain $\omega_j \approx_{\ell_A}^{(i,0)} \omega$, from which it follows that ω_j runs the same thread, stepping to ω'_j . We conclude that $\omega'_j \approx_{\ell_A} \omega'$ by low-simulation and determinism, i.e., Requirements (i) and (ii).

Definition 3 (Non-interfering Scheduler). *A scheduler is non-interfering if it satisfies Requirements 1, 2, and 3.*

Round Robin. We show that round-robin fulfills all the requirements and hence is an eligible candidate scheduler for our calculus. Firstly, it is immediately evident from the reductions that round-robin is *deterministic*, i.e., it fulfills scheduler requirement (i). We define the erasure function to filter out the identifiers of threads non observable to the attacker, i.e., $\varepsilon_{\ell_A}(s) = \text{filter}(\lambda(\ell, n) \rightarrow \ell \sqsubseteq \ell_A) s$. By induction on

²⁴In our conference version (Vassena and Russo, 2016), Requirements 2 and 3 are combined, but technically we need to split these two requirements. Progress of the concurrent configuration requires the former, while the latter ensures a well-founded inductive principle.

the scheduler reduction, it follows that round-robin satisfies *restricted simulation*, *no observable effect*, i.e., scheduler requirements (ii) and (iii). Before proving *progress* Figure 45 defines annotated ℓ_A -equivalence. In particular, if $\omega_1 \approx_{\ell_A}^{(0,0)} \omega_2$ for non-empty states ω_1 and ω_2 , then round-robin will schedule the same *low* thread in the next reduction. Lastly round-robin is *starvation-free* because it has a finite time-slot and is preemptive.

Proposition 4 (*RR non-interfering*). *Round-robin is non-interfering.*

10.3. Progress-sensitive Noninterference

The proof of progress-sensitive noninterference relies on lemmas similar to those listed in Requirement 1. In the following, we write $c_1 \hookrightarrow_{(\ell, n)} c_2$ to denote that configuration c_1 steps to c_2 executing thread (ℓ, n) and we use ℓ_L and ℓ_H to denote $\ell_L \sqsubseteq \ell_A$ and $\ell_H \not\sqsubseteq \ell_A$ respectively. As usual, we write \hookrightarrow^* for the reflexive transitive closure of \hookrightarrow . We write $c_1 \approx_{\ell_A} c_2$ if and only if $\varepsilon_{\ell_A}(c_1) \equiv \varepsilon_{\ell_A}(c_2)$, to denote ℓ_A -equivalence between configurations and we lift scheduler annotations, i.e., $c_1 \approx_{\ell_A}^{(i,j)} c_2$ if and only if $c_1 \approx_{\ell_A} c_2$ and $\omega_1 \approx_{\ell_A}^{(i,j)} \omega_2$.

Proposition 5.

- i) Determinancy: *if $c_1 \hookrightarrow c_2$ and $c_1 \hookrightarrow c_3$ then $c_2 \equiv c_3$.*
- ii) Restricted Simulation: *if $c_1 \hookrightarrow_{(\ell_L, n)} c_2$ then $\varepsilon_{\ell_A}(c_1) \hookrightarrow_{(\ell_L, n)} \varepsilon_{\ell_A}(c_2)$.*
- iii) No Observable Effect: *if $c_1 \hookrightarrow_{(\ell_H, n)} c_2$ then $c_1 \approx_{\ell_A} c_2$.*

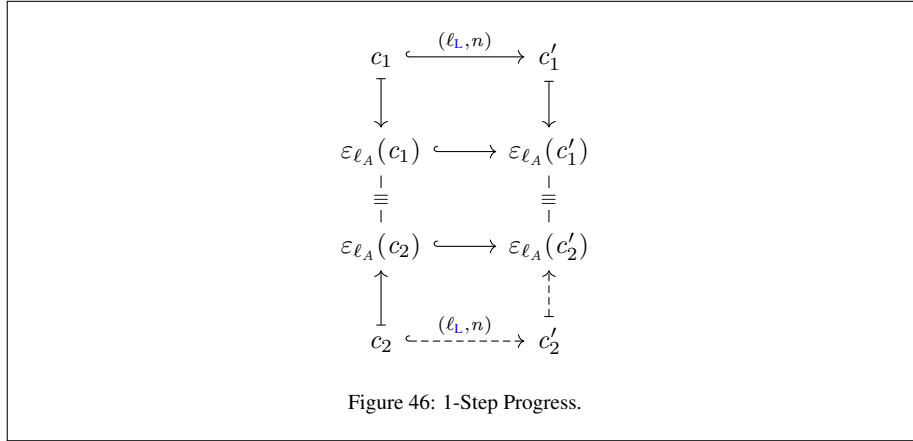
Using Proposition 5, we show that the concurrent semantics preserves ℓ_A -equivalence.

Proposition 6 (\approx_{ℓ_A} Preservation). *If $c_1 \approx_{\ell_A} c_2$ and $c_1 \hookrightarrow_{(\ell, n)} c'_1$, then*

- *If $\ell \not\sqsubseteq \ell_A$, then $c'_1 \approx_{\ell_A} c_2$.*
- *If $\ell \sqsubseteq \ell_A$ and $c_2 \hookrightarrow_{(\ell, n)} c'_2$, then $c'_1 \approx_{\ell_A} c'_2$.*

Progress sensitive noninterference requires to prove that ℓ_A -equivalence is preserved between two ℓ_A -equivalent configurations, even if only one steps. When a secret thread steps, the theorem follows easily by Proposition 6 and transitivity. The interesting case of the proof consists in showing *progress* of a public thread, which is simulated by the execution of multiple high threads followed by the same public thread, which corresponds to the diagram in Figure 44. Intuitively we prove *progress* by firstly simulating the secret threads that precede the public thread in the schedule (*scheduler progress*), then by simulating the common public thread under erasure (*restricted simulation*) and lastly *reconstructing* from the erased step the original step in the other public thread. Before proving this proposition, we have to restrict configurations c_1 and c_2 to be *valid*—we explain why we need this assumption later on.

Definition 4 (Valid Configuration). *A concurrent configuration c is valid if and only if it does not contain any invalid memory reference, node \bullet and terms $\text{new}\bullet$, $\text{write}\bullet$, $\text{fork}\bullet$, $\text{fmap}\bullet$, $\langle *\rangle\bullet$, $\text{relabel}\bullet$.*



Assuming valid configurations, we can prove *1-Step simulation*, i.e., the reconstruction of the other public step.

Proposition 7 (1-Step Progress). *If $c_1 \approx_{\ell_A}^{(i,0)} c_2$, $c_1 \hookrightarrow_{(\ell_L, n)} c_1'$ and c_2 is valid, then there exists c_2' such that $c_2 \hookrightarrow_{(\ell_L, n)} c_2'$.*

The diagram in Figure 46 shows our proof technique. Since the initial configurations are ℓ_A -equivalent, i.e., $c_1 \approx_{\ell_A}^{(i,0)} c_2$, then the erased initial configurations are equivalent, i.e., $\varepsilon_{\ell_A}(c_1) \equiv \varepsilon_{\ell_A}(c_2)$. Furthermore, since the schedulers in c_1 and c_2 are *aligned* (the second index in the annotated ℓ_A -equivalence is 0), the fact that the first scheduler runs thread (ℓ_L, n) , implies that the second runs it as well (Proposition 2). Given $c_1 \hookrightarrow_{(\ell_L, n)} c_1'$ we obtain the erased reduction step $\varepsilon_{\ell_A}(c_1) \hookrightarrow_{(\ell_L, m)} \varepsilon_{\ell_A}(c_1')$, by *restricted simulation* and we then *reconstruct* c_2' and the other step $c_2 \hookrightarrow_{(\ell_L, n)} c_2'$ from the step $\varepsilon_{\ell_A}(c_1) \hookrightarrow_{(\ell_L, m)} \varepsilon_{\ell_A}(c_1')$, $\varepsilon_{\ell_A}(c_1) \equiv \varepsilon_{\ell_A}(c_2)$ and the assumption that c_1 and c_2 are valid.

Validity. We explain by means of an example why we need to assume that the configurations c_1 and c_2 are *valid*. The fact that non-sensitive threads can write to sensitive resources, such as memories, complicates the reconstruction of a non-erased reduction step from an erased one, because, intuitively, too much information has been erased. For instance, since the erasure function rewrites secret memories and addresses to \bullet , we need to assume that the other program is in a “consistent state” in order to replay sensitive write memory-operations. Concretely, consider a public thread performing a secret write, i.e., $\langle \Sigma, \text{write } (\text{Ref } n) t \rangle \longrightarrow \langle \Sigma(\ell_H)[n] := t, \text{return } () \rangle$. A low-equivalent program will be $\langle \Sigma', \text{write } (\text{Ref } n') t' \rangle$, for some store Σ' , address n' and term t' such that $\Sigma \approx_{\ell_A} \Sigma'$, $\text{Ref } n \approx_{\ell_A} \text{Ref } n'$ and $t \approx_{\ell_A} t'$. Unfortunately, there is no guarantee that n' is a valid address in memory $\Sigma'(\ell_H)$. Observe that the erasure function maps is non-injective: it maps *both* valid and invalid references to $\text{Ref } \bullet$, therefore knowing that n is defined in $\Sigma(\ell_H)$ does not guarantee that n' is valid in $\Sigma'(\ell_H)$. Before proving *progress*, we show that our semantics preserves *validity*.

Proposition 8 (Valid Preservation). *If c_1 is valid and $c_1 \hookrightarrow c_2$ then c_2 is valid.*

Proposition 9 (Progress). *If $c_1 \approx_{\ell_A}^{(i,j)} c_2$, $c_1 \hookrightarrow_{(\ell_L, n)} c'_1$, and c_1, c_2 are valid configurations, then there exists c'_2 and c''_2 such that $c_2 \hookrightarrow^* c'_2 \hookrightarrow_{(\ell_L, n)} c''_2$.*

Proof (Sketch). The proof is driven by scheduler *progress*, i.e. Requirement 2, which determines what thread is scheduled next.

- ($j > 0$) The scheduler runs a secret thread, which is executed leading to the next intermediate configuration c'_2 , i.e. $c_2 \hookrightarrow_{(\ell_H, n')} c'_2$. By *no starvation*, i.e., Requirement 3, and *no observable effect*, i.e. Proposition 5.iii, it follows that $c_1 \approx_{\ell_A}^{(i, j')} c'_2$ for some $j' < j$ and we then apply induction.
- ($j = 0$) The scheduler runs public thread (ℓ_L, n) and the proposition follows from Proposition 7.

By combining *progress*, i.e., Proposition 9 and ℓ_A -equivalence *preservation*, i.e., Proposition 6, we prove PSNI.

Theorem 2 (Progress-sensitive noninterference). *Given valid global configurations c_1, c'_1, c_2 , and a non-interfering scheduler, if $c_1 \approx_{\ell_A} c_2$ and $c_1 \hookrightarrow c'_1$, then there exists c'_2 such that $c_2 \hookrightarrow^* c'_2$ and $c_2 \approx_{\ell_A} c'_2$.*

We conclude with a corollary that asserts that **MAC** satisfies PSNI.

Corollary 1. *MAC satisfies PSNI.*

Proof. By applying Theorem 2 and Proposition 4.

11. Related work

Mechanized Proofs. Russo presents the library **MAC** as a functional pearl and relies on its simplicity to convince readers about its correctness (Russo, 2015). This work bridges the gap on **MAC**'s lack of formal guarantees and exhibits interesting insights on the proofs of its soundness. **LIO** is a library structurally similar to **MAC** but dynamically enforcing IFC (Stefan et al., 2011b). The core calculus of **LIO**, i.e., side-effect free computations together with exception handling, has been modeled in the Coq proof assistant (Stefan et al., 2012b). Different from our work, these mechanized proofs do not simplify the treatment of sensitive exceptions by masking them in erased programs. In parallel to (Stefan et al., 2012b), Breeze (Hritcu et al., 2013) is a pure programming language that explores the design space of IFC and exceptions, which is accompanied with mechanized proofs in Coq. Bichhawat et al. develop an intra-procedural analysis for Javascript bytecode, which prevents implicit leaks in presence of exceptions and unstructured control flow constructs (Bichhawat et al., 2014).

Parametricity. Parametric polymorphism prevents a polymorphic function from inspecting its argument. In a similar manner, a non-interferent program cannot change its observable behaviour depending on the secret. Researchers have explored further this deep and subtle connection by obtaining a translation from DCC (Abadi et al., 1999) to System F in order to leverage on parametricity (Tse and Zdancewic, 2004). Shikuma and Igarashi (Shikuma and Igarashi, 2006) points out an error on such translation and gives a counterexample of a leaked translation. Recently, Bowman and Ahmed Bowman and Ahmed (2015) provide a sound translation from DCC into System F.

Concurrency. Considering IFC for a general scheduler could lead to refinements attacks. In this light, Russo and Sabelfeld provide termination-insensitive noninterference for a wide-class of deterministic schedulers (Russo and Sabelfeld, 2006a). Barthe et al. adopt this idea for Java-like bytecode (Barthe et al., 2009). Although we also consider deterministic schedulers, our security guarantees are stronger by considering leaks of information via abnormal termination. Heule et al. describe how to retrofit IFC in a programming language with sandboxes (Heule et al., 2015). Similar to our work, their soundness proofs are parametric on deterministic schedulers and provide progress-sensitive noninterference with informal arguments regarding thread progress—in this work, we spell out formal requirements on schedulers capable to guarantee thread progress. A series of work for π -calculus consider non-deterministic schedulers while providing progress-sensitive noninterference (Honda et al., 2000; Kobayashi, 2005; Honda and Yoshida, 2007; Pottier, 2002). Mantel and Sudbrock propose a novel scheduler-independent trace-based information-flow control property for multi-threaded programs and identify the class of *robust scheduler*, which satisfy that condition (Mantel and Sudbrock, 2010). While there are some similarities between the requirements of those robust schedulers and those discussed here in Section 10.2, that work assumes terminating threads, while our progress-sensitive noninterference theorem does not.

Security Libraries. Li and Zdancewic’s seminal work (Li and Zdancewic, 2006) shows how the structure *arrows* can provide IFC as a library in Haskell. Tsai et al. extend that work to support concurrency and data with heterogeneous labels (Tsai et al., 2007). Russo et al. implement the security library **SecLib** using a simpler structure than arrows (Russo et al., 2008), i.e. monads—rather than labeled values, this work introduces a monad which statically label side-effect free values. The security library **LIO** (Stefan et al., 2011a, 2012a) enforces IFC for both sequential and concurrent settings dynamically. **LIO** presents operations similar to *fmap* and *(*)* for labeled values with differences in the returning type due to **LIO**’s checks for clearance—this work provides a foundation to analyze the security implications of such primitives. Mechanized proofs for **LIO** are given only for its core sequential calculus (Stefan et al., 2011a). Inspired by **SecLib** and **LIO**’s designs, **MAC** leverages Haskell’s type system to enforce IFC (Russo, 2015) statically. Unlike **LIO**, data-dependent security policies cannot be expressed in **MAC**, due to its static nature. This limitation is addressed by **HLIO**, which provides a hybrid approach by means of some advanced Haskell’s type-system features: IFC is statically enforced while allowing the programmers to defer selected security checks until run-time (Buiras et al., 2015). Several works have also inves-

tigated the use of dependent types to precisely capture the nature of data-dependent security policies (Lourenço and Caires, 2015; Murray et al., 2016; Nanevski et al., 2011; Morgenstern and Licata, 2010).

Our work studies the security implications of extending **LIO**, **MAC**, and **HLIO** with a rich structure for labeled values. Devriese and Piessens provide a monad transformer to extend imperative-like APIs with support for IFC in Haskell (Devriese and Piessens, 2011). Jaskelioff and Russo implements a library which dynamically enforces IFC using secure multi-execution (SME) (Jaskelioff and Russo, 2011)—a technique that runs programs multiple times (once per security level) and varies the semantics of inputs and outputs to protect confidentiality. Rather than running multiple copies of a program, Schmitz et al. provide a library with *faceted values* (Schmitz et al., 2016), where values present different behavior according to the privilege of the observer. Different from the work above, we present a fully-fledged mechanized proof for our sequential and concurrent calculus which includes references, synchronization variables, and exceptions.

IFC Tools. IFC research has produced compilers capable of preserving confidentiality of data: Jif (Myers et al., 2001) and Paragon (Broberg et al., 2013) (based on Java), and FlowCaml (Simonet, 2003) (based on Caml). The SPARK language presents a IFC analysis which has been extended to guarantee progress-sensitive non-inference (Rafnsson et al., 2016). JSFlow (Hedin et al., 2014) is one of the state-of-the-art IFC system for the web (based on JavaScript). These tools preserve confidentiality in a fine-grained fashion where every piece of data is explicitly label. Specifically, there is no abstract data type to label data, so our results cannot directly apply to them.

Operating Systems. **MAC** borrows ideas from Mandatory Access Control (MAC) (Bell and La Padula, 1976; Biba, 1977) and phrases them into a programming language setting. Although originated in the 70s, there are modern manifestations of this idea (Zeldovich et al., 2006; Krohn et al., 2007; Murray et al., 2013), applied to diverse scenarios, like the web (Stefan et al., 2014; Bauer et al., 2015) and mobile devices (Jia et al., 2013; Bugiel et al., 2013). Due to its complexity, it is not surprising that OS-based MAC systems lack accompanying soundness guarantees or mechanized proofs—**seL4** being the exception (Murray et al., 2013). The level of abstractions handled by **MAC** and OSes are quite different, thus making uncertain how our insights could help to formalize OS-based MAC systems. MAC systems (Bell and La Padula, 1976) assign a label with an entire OS process—settling a single policy for all the data handled by it. In principle, it would be possible to extend such MAC-like systems to include a notion of labeled values with the functor structure as well as the relabeling primitive proposed by this work. For instance, COWL (Stefan et al., 2014) presents the notion of *labeled blob* and *labeled XHR* which is isomorphic to the notion of labeled values, thus making possible to apply our results. Furthermore, because many MAC-like system often ignore termination leaks (Efsthopoulos et al., 2005; Zeldovich et al., 2006), there is no need to use call-by-name evaluation to obtain security guarantees.

12. Conclusion

We present a *full-fledged* formalization of **MAC** in Agda, where noninterference is proven by term erasure. To the best of our knowledge, this is the first work of its kind for IFC libraries in Haskell, both for completeness and number of features included in the model. Thanks to our mechanized proofs, we identify challenges arising from erasing terms depending on the context where they appear and propose *two-steps erasure*—an effective technique to systematically deal with such cases. We present an extension of **MAC** that provides labeled values with an applicative functor-like structure and a relabeling operation, enabling convenient and expressive manipulation of labeled values using side effect-free code and saving programmers from introducing unnecessary sub-computations, e.g., forking a thread. We have proved this extension secure both in sequential and concurrent settings, using *two-steps erasure*. This work bridges the gap between existing IFC libraries (which focus on side-effecting code) and the usual Haskell programming model (which favors pure code), with a view to making IFC in Haskell more practical. Our mechanized proofs also make explicit sufficient scheduler requirements to guarantee PSNI—something that has been only treated informally before (Stefan et al., 2012a; Heule et al., 2015). As a result, our security proofs for the concurrent calculus are valid for a wide-range of deterministic schedulers. It is our hope that the insights gained by this work will help to formally verify other IFC programming languages.

References

- Abadi, M., Banerjee, A., Heintze, N., Riecke, J., Jan. 1999. A Core Calculus of Dependency. In: Proc. ACM Symp. on Principles of Programming Languages. pp. 147–160.
- Askarov, A., Hunt, S., Sabelfeld, A., Sands, D., 2008. Termination-insensitive non-interference leaks more than just a bit. In: Proc. of the European Symposium on Research in Computer Security (ESORICS '08). Springer-Verlag.
- Askarov, A., Zhang, D., Myers, A. C., 2010. Predictive black-box mitigation of timing channels. In: Proc. of the 17th ACM conference on Computer and Communications Security. ACM.
- Banerjee, A., Naumann, D. A., March 2005. Stack-based access control and secure information flow. Journal Functional Programming 15, 131–177.
- Barthe, G., Rezk, T., Russo, A., Sabelfeld, A., Aug. 2009. Security of multithreaded programs by compilation. Special issue of ACM Transactions on Information and System Security (TISSEC).
- Bauer, L., Cai, S., Jia, L., Passaro, T., Stroucken, M., Tian, Y., 2015. Run-time monitoring and formal analysis of information flows in Chromium. In: Proc. of the Annual Network & Distributed System Security Symposium. Internet Society.

- Bell, D. E., La Padula, L., 1976. Secure computer system: Unified exposition and multics interpretation. Tech. Rep. MTR-2997, Rev. 1, MITRE Corporation, Bedford, MA.
- Biba, K. J., 1977. Integrity considerations for secure computer systems. ESD-TR-76-372.
- Bichhawat, A., Rajani, V., Garg, D., Hammer, C., 2014. Information Flow Control in WebKit’s JavaScript Bytecode. Springer Berlin Heidelberg.
- Bortz, A., Boneh, D., 2007. Exposing private information by timing web applications. In: Proc. of the 16th World Wide Web. ACM.
- Bowman, W. J., Ahmed, A., 2015. Noninterference for free. In: Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015, Vancouver, BC, Canada, September 1-3, 2015. pp. 101–113.
- Broberg, N., van Delft, B., Sands, D., 2013. Paragon for practical programming with information-flow control. In: APLAS. Vol. 8301 of Lecture Notes in Computer Science. Springer, pp. 217–232.
- Bugiel, S., Heuser, S., Sadeghi, A.-R., 2013. Flexible and fine-grained mandatory access control on android for diverse security and privacy policies. In: Proc. of the 22Nd USENIX Conference on Security. SEC’13. USENIX Association.
- Buiras, P., Russo, A., 2013. Lazy programs leak secrets. In: Proc. Nordic Conference in Secure IT Systems (NORDSEC). Springer-Verlag.
- Buiras, P., Stefan, D., Russo, A., 2014. On dynamic flow-sensitive floating-label systems. In: Proceedings of the 2014 IEEE 27th Computer Security Foundations Symposium. CSF ’14. IEEE Computer Society, Washington, DC, USA, pp. 65–79. URL <http://dx.doi.org/10.1109/CSF.2014.13>
- Buiras, P., Vytiniotis, D., Russo, A., 2015. HLIO: Mixing static and dynamic typing for information-flow control in Haskell. In: Proc. of the ACM SIGPLAN International Conference on Functional Programming (ICFP ’15). ACM.
- Denning, D. E., Denning, P. J., Jul. 1977. Certification of programs for secure information flow. Communications of the ACM 20 (7), 504–513.
- Devriese, D., Piessens, F., 2011. Information flow enforcement in monadic libraries. In: Proc. of the ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI ’11). ACM.
- Efstathopoulos, P., Krohn, M., VanDeBogart, S., Frey, C., Ziegler, D., Kohler, E., Mazières, D., Kaashoek, F., Morris, R., 2005. Labels and event processes in the asbestos operating system. In: Proc. of the twentieth ACM symp. on Operating systems principles. SOSP ’05. ACM.
- Felten, E. W., Schneider, M. A., 2000. Timing attacks on Web privacy. In: Proc. of the 7th ACM conference on Computer and communications security. CCS ’00. ACM.

- Goguen, J., Meseguer, J., 1982. Security policies and security models. In: Proc of IEEE Symposium on Security and Privacy. IEEE Computer Society.
- Goguen, J. A., Meseguer, J., April 1984. Unwinding and inference control. In: Security and Privacy, 1984 IEEE Symposium on.
- Hedin, D., Birgisson, A., Bello, L., Sabelfeld, A., Mar. 2014. JSFlow: Tracking information flow in JavaScript and its APIs. In: Proc. of the ACM Symposium on Applied Computing (SAC '14). ACM.
- Hedin, D., Sands, D., 2006. Noninterference in the presence of non-opaque pointers. In: Proc. of the 19th IEEE Computer Security Foundations Workshop. IEEE Computer Society Press.
- Heule, S., Stefan, D., Yang, E. Z., Mitchell, J. C., Russo, A., April 2015. IFC inside: Retrofitting languages with dynamic information flow control. In: Conference on Principles of Security and Trust (POST). Springer.
- Honda, K., Vasconcelos, V. T., Yoshida, N., 2000. Secure information flow as typed process behaviour. In: Proc. of the 9th European Symposium on Programming Languages and Systems. Springer-Verlag.
- Honda, K., Yoshida, N., Oct. 2007. A uniform type structure for secure information flow. ACM Trans. Program. Lang. Syst.
- Hritcu, C., Greenberg, M., Karel, B., Peirce, B. C., Morrisett, G., 2013. All your IFC exceptions are belong to us. In: Proc. of the IEEE Symposium on Security and Privacy. IEEE Computer Society.
- Hughes, J., 1984. Why functional programming matters. *The Computer Journal* 32.
- Jaskelioff, M., Russo, A., Jun. 2011. Secure multi-execution in Haskell. In: Proc. Andrei Ershov International Conference on Perspectives of System Informatics. LNCS. Springer-Verlag.
- Jia, L., Aljuraidan, J., Fragkaki, E., Bauer, L., Stroucken, M., Fukushima, K., Kiyomoto, S., Miyake, Y., Sep. 2013. Run-time enforcement of information-flow properties on Android (extended abstract). In: Computer Security—ESORICS 2013: 18th European Symposium on Research in Computer Security. Springer.
- Kobayashi, N., Dec. 2005. Type-based information flow analysis for the π -calculus. *Acta Inf.* 42 (4).
- Krohn, M., Yip, A., Brodsky, M., Cliffer, N., Kaashoek, M. F., Kohler, E., Morris, R., 2007. Information flow control for standard OS abstractions. In: Proceedings of ACM SIGOPS Symposium on Operating Systems Principles. SOSP '07. ACM.
- Li, P., Zdancewic, S., 2006. Encoding information flow in Haskell. In: Proc. of the IEEE Workshop on Computer Security Foundations. IEEE Computer Society.

- Li, P., Zdancewic, S., 2010. Arrows for secure information flow. *Theoretical Computer Science* 411 (19), 1974–1994.
- Lourenço, L., Caires, L., Jan. 2015. Dependent information flow types. *SIGPLAN Not.* 50 (1).
- Mantel, H., Sudbrock, H., 2010. *Flexible Scheduler-Independent Security*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 116–133.
URL https://doi.org/10.1007/978-3-642-15497-3_8
- Marlow, S., July 2013. *Parallel and concurrent programming in Haskell*. O’Reilly.
- Mcbride, C., Paterson, R., 2008. Applicative programming with effects. *Journal of Functional Programming*.
- Meurer, S., Wismler, R., 2012. Apefs: An infrastructure for permission-based filtering of android apps. In: Schmidt, A., Russello, G., Krontiris, I., Lian, S. (Eds.), *Security and Privacy in Mobile Information and Communication Systems*. Vol. 107. Springer Berlin Heidelberg.
- Moggi, E., 1991. Notions of computation and monads. *Information and Computation* 93 (1), 55–92.
- Morgenstern, J., Licata, D. R., 2010. Security-typed programming within dependently typed programming. In: *Proc. of the 15th ACM SIGPLAN International Conference on Functional Programming*. ACM.
- Murray, T., Matichuk, D., Brassil, M., Gammie, P., Bourke, T., Seefried, S., Lewis, C., Gao, X., Klein, G., 2013. seL4: From general purpose to a proof of information flow enforcement. *2012 IEEE Symposium on Security and Privacy* 0.
- Murray, T., Sison, R., Pierzchalski, E., Rizkallah, C., Jun. 2016. Compositional verification and refinement of concurrent value-dependent noninterference. In: *IEEE Computer Security Foundations Symposium*. Lisbon, Portugal, pp. 417–431.
- Myers, A. C., Jan. 1999. JFlow: Practical mostly-static information flow control. In: *Proc. ACM Symp. on Principles of Programming Languages*. pp. 228–241.
- Myers, A. C., Zheng, L., Zdancewic, S., Chong, S., Nystrom, N., 2001. Jif: Java Information Flow, <http://www.cs.cornell.edu/jif>.
- Nanevski, A., Banerjee, A., Garg, D., 2011. Verification of information flow and access control policies with dependent types. In: *Proceedings of the 2011 IEEE Symposium on Security and Privacy*. SP ’11. IEEE Computer Society.
- Pedersen, M. V., Askarov, A., 2017. From trash to treasure: Timing-sensitive garbage collection. In: *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*. pp. 693–709.
URL <https://doi.org/10.1109/SP.2017.64>

- Pottier, F., 2002. A simple view of type-secure information flow in the π -calculus. In: In Proc. of the 15th IEEE Computer Security Foundations Workshop. pp. 320–330.
- Pottier, F., Simonet, V., Jan. 2002. Information Flow Inference for ML. In: Proc. ACM Symp. on Principles of Programming Languages. pp. 319–330.
- Rafnsson, W., Garg, D., Sabelfeld, A., 2016. Progress-sensitive security for SPARK. In: Engineering Secure Software and Systems - 8th International Symposium, ESSoS 2016, London, UK, April 6-8, 2016. Proceedings.
- Roy, I., Porter, D. E., Bond, M. D., McKinley, K. S., Witchel, E., 2009. Laminar: Practical fine-grained decentralized information flow control. In: Proceedings of the 2009 ACM SIGPLAN Conference on Programming Language Design and Implementation. PLDI '09. ACM.
- Russo, A., 2015. Functional pearl: Two can keep a secret, if one of them uses Haskell. In: Proc. of the 20th ACM SIGPLAN International Conference on Functional Programming. ICFP 2015. ACM.
- Russo, A., Claessen, K., Hughes, J., Sep. 2008. A library for light-weight information-flow security in Haskell. In: Proc. ACM SIGPLAN symposium on Haskell (HASKELL '08). ACM.
- Russo, A., Sabelfeld, A., Jul. 2006a. Securing interaction between threads and the scheduler. In: Proc. IEEE Computer Sec. Foundations Workshop. pp. 177–189.
- Russo, A., Sabelfeld, A., Jun. 2006b. Security for Multithreaded Programs under Cooperative Scheduling. In: Proc. Andrei Ershov International Conference on Perspectives of System Informatics. LNCS. Springer-Verlag.
- Sabelfeld, A., Myers, A. C., Jan. 2003a. Language-based information-flow security. IEEE Journal on Selected Areas in Communications 21 (1), 5–19.
- Sabelfeld, A., Myers, A. C., Jan. 2003b. Language-Based Information-Flow Security. IEEE J. Selected Areas in Communications 21 (1), 5–19.
- Schmitz, T., Rhodes, D., Austin, T. H., Knowles, K., Flanagan, C., 2016. Faceted dynamic information flow via control and data monads. In: Piessens, F., Vigan, L. (Eds.), POST. Vol. 9635 of Lecture Notes in Computer Science. Springer.
- Shikuma, N., Igarashi, A., 2006. Proving noninterference by a fully complete translation to the simply typed lambda-calculus. In: Advances in Computer Science - ASIAN 2006. Secure Software and Related Issues, 11th Asian Computing Science Conference, Tokyo, Japan, December 6-8, 2006, Revised Selected Papers. pp. 301–315.
- Simonet, V., 2003. The Flow Caml system, software release at <http://crystal.inria.fr/~simonet/soft/flowcaml/>.

- Smith, G., Volpano, D., 1998. Secure information flow in a multi-threaded imperative language. In: Proc. ACM symposium on Principles of Programming Languages (POPL '98).
- Stefan, D., Buiras, P., Yang, E. Z., Levy, A., Terei, D., Russo, A., Mazières, D., 2013. Eliminating cache-based timing attacks with instruction-based scheduling. In: Proc. European Symp. on Research in Computer Security. LNCS. Springer-Verlag.
- Stefan, D., Russo, A., Buiras, P., Levy, A., Mitchell, J. C., Mazières, D., 2012a. Addressing covert termination and timing channels in concurrent information flow systems. In: Proc. of the ACM SIGPLAN International Conference on Functional Programming (ICFP '12). ACM.
- Stefan, D., Russo, A., Mitchell, J. C., Mazières, D., 2011a. Flexible dynamic information flow control in Haskell. In: Proc. of the ACM SIGPLAN Haskell symposium (HASKELL '11).
- Stefan, D., Russo, A., Mitchell, J. C., Mazières, D., 2011b. Flexible dynamic information flow control in Haskell. In: Proceedings of the 4th ACM symposium on Haskell. ACM, New York, NY, USA, pp. 95–106.
- Stefan, D., Russo, A., Mitchell, J. C., Mazières, D., 2012b. Flexible dynamic information flow control in the presence of exceptions. Arxiv preprint arXiv:1207.1457, to appear in Journal of Functional Programming, Cambridge University Press.
- Stefan, D., Yang, E. Z., Marchenko, P., Russo, A., Herman, D., Karp, B., Mazières, D., Oct. 2014. Protecting users by confining JavaScript with COWL. In: 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14). USENIX Association.
- Tsai, T. C., Russo, A., Hughes, J., Jul. 2007. A library for secure multi-threaded information flow in Haskell. In: Proc. IEEE Computer Security Foundations Symposium (CSF '07).
- Tse, S., Zdancewic, S., 2004. Translating dependency into parametricity. In: Proc. of the Ninth ACM SIGPLAN International Conference on Functional Programming. ACM.
- Vassena, M., Breitner, J., Russo, A., 2017. Securing concurrent lazy programs against information leakage. In: Proc. IEEE Computer Sec. Foundations Symposium. (CSF '17).
- Vassena, M., Buiras, P., Waye, L., Russo, A., Sep. 2016. Flexible manipulation of labeled values for information-flow control libraries. In: Proceedings of the 12th European Symposium On Research In Computer Security. Springer.
- Vassena, M., Russo, A., 2016. On formalizing information-flow control libraries. In: Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security. PLAS '16. ACM, New York, NY, USA, pp. 15–28.
URL <http://doi.acm.org/10.1145/2993600.2993608>

- Wadler, P., Blott, S., 1989. How to make ad-hoc polymorphism less ad hoc. In: Proceedings of the 16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL '89. ACM, New York, NY, USA, pp. 60–76.
URL <http://doi.acm.org/10.1145/75277.75283>
- Zeldovich, N., Boyd-Wickizer, S., Kohler, E., Mazières, D., 2006. Making information flow explicit in HiStar. In: Proc. of the 7th USENIX Symp. on Operating Systems Design and Implementation. USENIX.
- Zhang, D., Askarov, A., Myers, A. C., 2011. Predictive mitigation of timing channels in interactive systems. In: Proc. of the 18th ACM conference on Computer and Communications Security. ACM.
- Zhang, D., Askarov, A., Myers, A. C., 2012. Language-based Control and Mitigation of Timing Channels. In: Proc. ACM Conference on Programming Language Design and Implementation. ACM.

$$\begin{array}{l}
\text{(RELABEL}_\chi\text{)} \qquad \qquad \qquad \text{(RELABEL}_{\bullet,\chi}\text{)} \\
\text{relabel } (\text{Labeled}_\chi t) \rightsquigarrow \text{Labeled}_\chi t \qquad \text{relabel}_{\bullet} (\text{Labeled}_\chi t) \rightsquigarrow \text{Labeled } \bullet \\
\\
\text{((*)}_{\chi 1}\text{)} \\
\frac{t_2 \rightsquigarrow t'_2}{(\text{Labeled}_\chi t_1) (*) t_2 \rightsquigarrow (\text{Labeled}_\chi t_1) (*) t'_2} \\
\\
\text{((*)}_{\chi 2}\text{)} \\
(\text{Labeled}_\chi t_1) (*) (\text{Labeled } t_2) \rightsquigarrow \text{Labeled}_\chi t_1 \\
\\
\text{((*)}_{\chi 3}\text{)} \\
(\text{Labeled}_\chi t_1) (*) (\text{Labeled}_\chi t_2) \rightsquigarrow \text{Labeled}_\chi t_1 \\
\\
\text{((*)}_{\chi 4}\text{)} \\
(\text{Labeled } t_1) (*) (\text{Labeled}_\chi t_2) \rightsquigarrow \text{Labeled}_\chi t_2
\end{array}$$

Figure A.47: Semantics of flexible labeled values with exceptions.

Appendix A. Flexible Labeled Values in Sequential Calculus

In this section, we extend the semantics of flexible labeled values described in Section 9 for the sequential setting, where labeled values have an additional constructor, namely Labeled_χ . This constructor is used to prevent *sensitive* exceptions from leaking into a *non-sensitive* context, when embedding a secret computation in a public one using *join*. The semantics of the primitives *relabel* and $\langle * \rangle$ handle exceptional values, by propagating the exceptions, which is exactly what happens in rule [RELABEL $_\chi$]—see Figure A.47. Rule [RELABEL $_{\bullet,\chi}$] simulates rule [RELABEL $_\chi$] in the sequential setting, specifically when a public exceptional labeled value gets relabeled with a sensitive label (note that the resulting erased, i.e., $\text{Labeled } \bullet$, is sensitive and contains no exception). Rules [$\langle * \rangle_{\chi 1}$, $\langle * \rangle_{\chi 2}$, $\langle * \rangle_{\chi 3}$, $\langle * \rangle_{\chi 4}$] yield (propagate) the first exception observed when $\langle * \rangle$ is applied to exceptional values. In particular, rule [$\langle * \rangle_{\chi 3}$] applies when both arguments are exceptions and returns the first one triggered during evaluation, i.e., the left one. Rules [$\langle * \rangle_{\chi 1}$, $\langle * \rangle_{\chi 2}$, $\langle * \rangle_{\chi 3}$] are somewhat unusual. In particular, even though our language is *non-strict*, the rules give a *strict* semantics to $\langle * \rangle$ —note that they reduce unnecessarily the second term, even though it is not used in the final result. It would have been more natural, in this context, to replace them by a single rule $\text{Labeled}_\chi t_1 \langle * \rangle t_2 \rightsquigarrow \text{Labeled}_\chi t_1$, that does not evaluate the second term. The two alternative semantics are equivalent, except for abnormal *non-terminating* terms, that we denote with \perp . With *strict* semantics, the term $(\text{Labeled}_\chi t_1) \langle * \rangle \perp$ results in \perp , because it loops due to rule [$\langle * \rangle_{\chi 1}$], instead it terminates with a *non-strict* semantics, i.e., $(\text{Labeled}_\chi t_1) \langle * \rangle \perp \rightsquigarrow (\text{Labeled}_\chi t_1)$. We remark that the two semantics are equivalent for *terminating* programs and therefore security is not at stake:

```

data  $MVar\ \ell\ \tau$ 
 $newMVar :: \ell_L \sqsubseteq \ell_H \Rightarrow MAC\ \ell_L\ (MVar\ \ell_H\ \tau)$ 
 $takeMVar :: MVar\ \ell\ \tau \rightarrow MAC\ \ell\ \tau$ 
 $putMVar :: MVar\ \ell\ \tau \rightarrow \tau \rightarrow MAC\ \ell\ ()$ 

```

Figure B.48: API of synchronization primitives.

the sequential calculus is already *termination insensitive*. Technically, we give a strict definition of $\langle * \rangle$, because erasing sensitive exceptions are replaced by non-exceptional values, i.e., $\varepsilon_L(\text{Labeled}_\chi\ t :: \text{Labeled}\ H\ \tau) = \text{Labeled}\ \bullet$. Therefore, we could not prove simulation for a non-strict applicative functor, since, crucially, it is sensitive to exceptions. While these behaviour could be simulated by an erasure function that preserves sensitive exceptions, i.e., $\varepsilon_L(\text{Labeled}_\chi\ \tau :: \text{Labeled}\ H\ \tau) = \text{Labeled}_\chi\ \bullet$, it is an open question how to prove *single-step simulation* for *join*, specifically for rules [JOIN, JOIN $_\chi$].

Appendix B. Synchronization Primitives

In this section we extend our calculus with synchronization primitives, an essential feature for concurrent programs. Using synchronized mutable variables (*MVar*) users can implement simple inter-thread communication mechanisms such as binary semaphores and channels.

The type $MVar\ \ell\ \tau$ denotes a labeled mutable location that is either empty or full and contains a term of type τ of security level τ . Figure B.48 shows the API of basic synchronization primitives, based on *MVar*. Specifically, function *newMVar* creates an empty *MVar*. Function *takeMVar* empty a full *MVar* and returns its content or *blocks* otherwise. Function *putMVar* fills an empty *MVar* or *blocks* otherwise. Primitive *newMVar* performs a *write* operation, therefore its type is restricted to comply with the *no write-down* policy, just like the type of *new* for memory. Interestingly, and unlike memory primitives *read* and *write*, the type of *takeMVar* and *putMVar* accepts only one security level. Intuitively, that is the case because *MVar*'s primitives perform *both* read and write side-effects, therefore both *no read-up* and *no write-down* security policies apply. For instance, to execute *putMVar*, it is necessary to observe (read) if the *MVar* is empty. We show how those security policy guide our design and lead us to give the API shown in Figure B.48 as the only secure option. Assume that primitive *takeMVar* had a completely unrestricted type, i.e., $\forall \ell_1\ \ell_2 . MVar\ \ell_1\ \tau \rightarrow MAC\ \ell_2\ \tau$. Since *takeMVar* returns the content of the *MVar*—a *read* effect that is secure only if ℓ_2 is at least as sensitive as ℓ_1 , i.e., $\ell_1 \sqsubseteq \ell_2$. Observe however that *takeMVar* empties the *MVar* as well, after returning its content—a *write* effect that is secure only if ℓ_1 is at least as sensitive as ℓ_2 , i.e., $\ell_2 \sqsubseteq \ell_1$. By the antisymmetry of the security lattice, it follows that the interaction between computations and synchronization variables is secure only when they have the same security level, i.e., $\ell_1 \equiv \ell_2$. The same principle applies for *putMVar*.

Memory ℓ	$t_s ::= []$		$c : t_s$
Cell	$c ::= \otimes$		$\llbracket t \rrbracket$
Types:	$\tau ::= \dots$		$MVar \ell \tau$
Values:	$v ::= \dots$		$MVar n$
Terms:	$t ::= \dots$		$newMVar$ $takeMVar t$ $putMVar t_1 t_2$

(NEWMVAR)

$$\frac{|\Sigma(\ell)| = n}{\langle \Sigma, newMVar \rangle \longrightarrow \langle \Sigma(\ell)[n] := \otimes, return (MVar n) \rangle}$$

(PUTMVAR₁)

$$\frac{t_1 \rightsquigarrow t'_1}{\langle \Sigma, putMVar t_1 t_2 \rangle \longrightarrow \langle \Sigma, putMVar t'_1 t_2 \rangle}$$

(PUTMVAR₂)

$$\frac{\Sigma(\ell)[n] = \otimes}{\langle \Sigma, putMVar (MVar n) t \rangle \longrightarrow \langle \Sigma(\ell)[n] := \llbracket t \rrbracket, return () \rangle}$$

(TAKEMVAR₁)

$$\frac{t \rightsquigarrow t'}{\langle \Sigma, takeMVar t \rangle \longrightarrow \langle \Sigma, takeMVar t' \rangle}$$

(TAKEMVAR₂)

$$\frac{\Sigma(\ell)[n] = \llbracket t \rrbracket}{\langle \Sigma, takeMVar (MVar n) \rangle \longrightarrow \langle \Sigma(\ell)[n] := \otimes, return t \rangle}$$

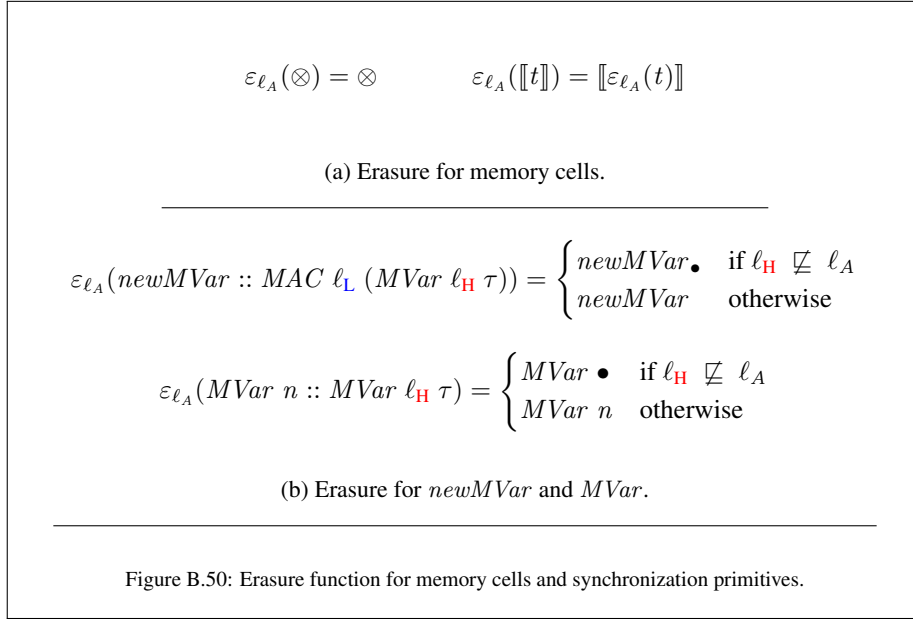
Figure B.49: MAC with synchronization primitives.

Appendix B.1. Calculus

Figure B.49 extends the concurrent calculus with synchronization primitives. A synchronization variable is represented as a value $MVar n :: MVar \ell \tau$ where n is an address²⁵, pointing to the n -th cell of the ℓ -memory, which contains a term of type τ . We adjust our memory model to work with synchronization variables²⁶. We introduce a new syntactic category, memory cell c , which can be either *empty*, i.e., \otimes , or *full* with some term t , i.e., $\llbracket t \rrbracket$. Rule [NEWMVAR] evaluates term $newMVar$ by adding an empty memory cell to the ℓ -labeled memory, i.e., $\Sigma(\ell)[n] := \otimes$ and returning a reference to it, i.e., $MVar n$. Rule [PUTMVAR₁] evaluates the reference and rule [PUTMVAR₂]

²⁵In **MAC** a $MVar$ is just a wrapper around unlabeled synchronization variables from the standard library. Here we denote synchronization variables as an index, just like we did for memory references.

²⁶We model mutable references as a special case of synchronization variables that are always full.



fills the empty cell it refers to with the term, i.e., $\Sigma(\ell)[n] := \llbracket t \rrbracket$ and returns unit. Rule [TAKEMVAR₁] evaluates the reference and rule [TAKEMVAR₂] returns the content of the non-empty cell it refers to, i.e., $\Sigma(\ell)[n] = \llbracket t \rrbracket$ for *some* term t , and empties it, i.e., $\Sigma(\ell)[n] := \otimes$. Observe that the premise of rules [PUTMVAR₂] and [TAKEMVAR₂] accounts for the *blocking* behavior of the synchronization primitives by making the configuration *stuck*. In particular, primitive *putMVar* *blocks* if the cell is non-empty, i.e., $\langle \Sigma, \text{putMVar } (MVar \ n) \ t \rangle \not\rightarrow$ if $\Sigma(\ell)[n] \neq \otimes$ and similarly *takeMVar* *blocks* if the cell is empty, i.e., $\langle \Sigma, \text{takeMVar } (MVar \ n) \rangle \not\rightarrow$ if $\Sigma(\ell)[n] \equiv \otimes$.

Appendix B.2. Erasure Function

Proving that synchronization primitives are secure is straightforward in our setting. The primitives are clearly *deterministic* and showing *single-step simulation* is simpler than for references because primitives *putMVar* and *takeMVar* work within the same security level. Memory cells are erased homomorphically (Figure B.50a). Applying the *two-steps erasure* technique, the erasure function replaces term *newMVar* with *newMVar*_•, when it creates a sensitive synchronization variable—see Figure B.50b. The erasure function rewrites the address of a synchronization reference to • if it points to a sensitive memory. Figure B.51 shows rule [NEWMVAR_•], which reduces term *newMVar*_•, returns a dummy reference, i.e., *MVar* •, and *skips* the write effect, leaving the store Σ unchanged. Observe that we do not need to replace *takeMVar* with a special term *takeMVar*_•, because the primitive can only write to a memory at the same security level as the computation, therefore either they are both sensitive and the computation rewritten to • or both non-sensitive and erased homomorphically.

Terms: $t ::= \dots \mid newMVar_{\bullet}$

$$\begin{array}{c} \text{(NEWMVAR}_{\bullet}\text{)} \\ \langle \Sigma, newMVar_{\bullet} \rangle \longrightarrow \langle \Sigma, return (MVar \bullet) \rangle \end{array}$$

Figure B.51: Semantics of $newMVar_{\bullet}$.

Appendix C. Typing Rules

Figure C.52 gives the typing rules for the extended calculus, i.e., term \bullet and other \bullet -annotated terms used when applying two-steps erasure. The special term \bullet can assume any type thanks to the typing rule [HOLE]. The typing rule of each \bullet -annotated term corresponds exactly to the typing rule of the unannotated terms. As a consequence of these typing rules, the erasure function is type-preserving, i.e., if $\Gamma \vdash t : \tau$ then $\varepsilon_{\ell_A}(\Gamma) \vdash \varepsilon_{\ell_A}(t) : \tau$.

$$\begin{array}{c}
\text{(HOLE)} \\
\Gamma \vdash \bullet : \tau
\end{array}
\qquad
\begin{array}{c}
\text{(NEW}\bullet\text{)} \\
\frac{\ell_L \sqsubseteq \ell_H \quad \Gamma \vdash t : \tau}{\Gamma \vdash \text{new}\bullet t : \text{MAC } \ell_L (\text{Ref } \ell_H \tau)}
\end{array}$$

$$\begin{array}{c}
\text{(WRITE}\bullet\text{)} \\
\frac{\ell_L \sqsubseteq \ell_H \quad \Gamma \vdash t_1 : \tau \quad \Gamma \vdash t_2 : \text{Ref } \ell_H \tau}{\Gamma \vdash \text{write}\bullet t_1 t_2 : \text{MAC } \ell_L ()}
\end{array}$$

$$\begin{array}{c}
\text{(JOIN}\bullet\text{)} \\
\frac{\ell_L \sqsubseteq \ell_H \quad \Gamma \vdash t : \text{MAC } \ell_H \tau}{\Gamma \vdash \text{join}\bullet t : \text{MAC } \ell_L (\text{Labeled } \ell_H \tau)}
\end{array}
\qquad
\begin{array}{c}
\text{(FORK}\bullet\text{)} \\
\frac{\ell_L \sqsubseteq \ell_H \quad \Gamma \vdash t : \text{MAC } \ell_H ()}{\Gamma \vdash \text{fork}\bullet t : \text{MAC } \ell_L ()}
\end{array}$$

$$\begin{array}{c}
\text{(FMAP}\bullet\text{)} \\
\frac{\Gamma \vdash t_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash t_2 : (\text{Labeled } \ell \tau_1)}{\Gamma \vdash \text{fmap}\bullet t_1 t_2 : \text{Labeled } \ell \tau_2}
\end{array}$$

$$\begin{array}{c}
\text{(\langle*\rangle}\bullet\text{)} \\
\frac{\Gamma \vdash t_1 : \text{Labeled } \ell (\tau_1 \rightarrow \tau_2) \quad \Gamma \vdash t_2 : (\text{Labeled } \ell \tau_1)}{\Gamma \vdash t_1 \langle*\rangle\bullet t_2 : \text{Labeled } \ell \tau_2}
\end{array}$$

$$\begin{array}{c}
\text{(RELABEL}\bullet\text{)} \\
\frac{\ell_L \sqsubseteq \ell_H \quad \Gamma \vdash t : (\text{Labeled } \ell_L \tau)}{\Gamma \vdash \text{relabel}\bullet t : \text{Labeled } \ell_H \tau}
\end{array}$$

$$\begin{array}{c}
\text{(NEWMVAR}\bullet\text{)} \\
\frac{\ell_L \sqsubseteq \ell_H}{\Gamma \vdash \text{newMVar}\bullet : \text{MAC } \ell_L (\text{MVar } \ell_H \tau)}
\end{array}$$

Figure C.52: Typing rules for the extended calculus.