A Dependently Typed Library for Static Information-Flow Control in Idris*

Simon Gregersen, Søren Eller Thomsen, and Aslan Askarov

Aarhus University, Aarhus, Denmark
{gregersen,sethomsen,askarov}@cs.au.dk

Abstract. Safely integrating third-party code in applications while protecting the confidentiality of information is a long-standing problem. Pure functional programming languages, like Haskell, make it possible to enforce lightweight information-flow control through libraries like **MAC** by Russo. This work presents DepSec, a **MAC** inspired, dependently typed library for static information-flow control in Idris. We showcase how adding dependent types increases the expressiveness of state-of-the-art static information-flow control libraries and how DepSec matches a special-purpose dependent information-flow type system on a key example. Finally, we show novel and powerful means of specifying statically enforced declassification policies using dependent types.

Keywords: Information-Flow Control · Dependent Types · Idris.

1 Introduction

Modern software applications are increasingly built using libraries and code from multiple third parties. At the same time, protecting confidentiality of information manipulated by such applications is a growing, yet long-standing problem. Third-party libraries could in general have been written by anyone and they are usually run with the same privileges as the main application. While powerful, such privileges open up for abuse.

Traditionally, access control [7] and encryption have been the main means for preventing data dissemination and leakage, however, such mechanisms fall short when third-party code needs access to sensitive information to provide its functionality. The key observation is that these mechanisms only place restrictions on the access to information but not its propagation. Once information is accessed, the accessor is free to improperly transmit or leak the information in some form, either by intention or error.

Language-based Information-Flow Control [35] is a promising technique for enforcing information security. Traditional enforcement techniques analyze how information at different security levels flows within a program ensuring that information flows only to appropriate places, suppressing illegal flows. To achieve this, most information-flow control tools require the design of new languages, compilers, or interpreters (e.g. [12, 16, 21, 22, 25, 28, 38]). Despite a large, growing body of work on language-based information-flow security, there has been little adoption of the proposed techniques. For information-flow policies to be enforced in such systems, the whole system has to be

* This is an extended version of a paper of the same title presented at POST 2019.

written in new languages – an inherently expensive and time-consuming process for large software systems. Moreover, in practice, it might very well be that only small parts of an application are governed by information-flow policies.

Pure functional programming languages, like Haskell, have something to offer with respect to information security as they strictly separate side-effect free and side-effectful code. This makes it possible to enforce lightweight information-flow control through libraries [11, 19, 33, 34, 41] by constructing an embedded domain-specific security sub-language. Such libraries enforce a secure-by-construction programming model as any program written against the library interface is not capable of leaking secrets. This construction forces the programmer to write security-critical code in the sub-language but otherwise allows them to freely interact and integrate with non-security critical code written in the full language. In particular, static enforcement libraries like **MAC** [33] are appealing as no run-time checks are needed and code that exhibits illegal flows is rejected by the type checker at compile-time. Naturally, the expressiveness of Haskell's type system sets the limitation on which programs can be deemed secure and which information flow policies can be guaranteed.

Dependent type theories [23, 30] are implemented in many programming languages such as Coq [13], Agda [31], Idris [8], and F* [43]. Programming languages that implement such theories allow types to dependent on values. This enables programmers to give programs a very precise type and increased confidence in its correctness.

In this paper, we show that dependent types provide a direct and natural way of expressing precise data-dependent security policies. Dependent types can be used to represent rich security policies in environments like databases and data-centric web applications where, for example, new classes of users and new kinds of data are encountered at run-time and the security level depends on the manipulated data itself [22]. Such dependencies are not expressible in less expressive systems like **MAC**. Among other things, with dependent types, we can construct functions where the security level of the output depends on an argument:

```
getPassword : (u : Username) -> Labeled u String
```

Given a user name `u`, `getPassword` retrieves the corresponding password and classifies it at the security level of `u`. As such, we can express much more precise security policies that can depend on the manipulated data.

Idris is a general-purpose functional programming language with full-spectrum dependent types, that is, there is no restrictions on which values may appear in types. The language is strongly influenced by Haskell and has, among others, inherited its strict encapsulation of side-effects. Idris essentially asks the question: "What if Haskell had full dependent types?" [9]. This work, essentially, asks:

> "What if **MAC** had full dependent types?"

We address this question using Idris because of its positioning as a general-purpose language rather than a proof assistant. All ideas should be portable to equally expressive systems with full dependent types and strict monadic encapsulation of side-effects.

In summary, the contributions of this paper are as follows.

- We present DepSec, a **MAC** inspired statically enforced dependently typed information-flow control library for Idris.

- We show how adding dependent types strictly increases the expressiveness of state-of-the-art static information-flow control libraries and how DepSec matches the expressiveness of a special-purpose dependent information-flow type system on a key example.
- We show how DepSec enables and aids the construction of policy-parameterized functions that abstract over the security policy.
- We show novel and powerful means to specify statically-ensured declassification using dependent types for a wide variety of policies.
- We show progress-insensitive noninterference [1] for the core library in a sequential setting.

Outline The rest of the paper proceeds through a presentation of the DepSec library (Section 2); a conference manager case study (Section 3) and the introduction of policy-parameterized functions (Section 4) both showcasing the expressiveness of DepSec; means to specify statically-ensured declassification policies (Section 5); soundness of the core library (Section 6); and related work (Section 7).

All code snippets presented in the following are extracts from the source code. All source code is implemented in Idris 1.3.1. and available at

`https://github.com/simongregersen/DepSec`.

The source code of the core library is also available in Appendix C.

1.1 Assumptions and threat model

In the rest of this paper, we require that code is divided up into trusted code, written by someone we trust, and untrusted code, written by a potential attacker. The trusted computing base (TCB) has no restrictions, but untrusted code does not have access to modules providing input/output behavior, the data constructors of the domain specific language and a few specific functions related to declassification. In Idris, this means that we specifically do not allow access to `IO` functions and `unsafePerformIO`. In DepSec, constructors and functions marked with a `TCB` comment are inaccessible to untrusted code. Throughout the paper we will emphasize when this is the case.

We require that all definitions made by untrusted code are total, that is, defined for all possible inputs and are guaranteed to terminate. This is necessary if we want to trust proofs given by untrusted code. Otherwise, it could construct an element of the empty type from which it could prove anything:

```
empty : Void
empty = empty
```

In Idris, this can be checked using the `--total` compiler flag. Furthermore, we do not consider concurrency nor any internal or termination covert channels.

2 The DepSec library

In information-flow control, labels are used to model the sensitivity of data. Such labels usually form a security lattice [14] where the induced partial ordering $\sqsubseteq$ specifies

allowed flows of information and hence the security policy. For example, $\ell_1 \sqsubseteq \ell_2$ specifies that data with label ℓ_1 is allowed to flow to entities with label ℓ_2. In DepSec, labels are represented by values that form a verified join semilattice implemented as Idris interfaces[1]. That is, we require proofs of the lattice properties when defining an instance of `JoinSemilattice`.

```
interface JoinSemilattice a where
  join : a -> a -> a
  associative :
    (x, y, z : a) -> x `join` (y `join` z) = (x `join` y) `join` z
  commutative : (x, y : a)    -> x `join` y = y `join` x
  idempotent  : (x : a)       -> x `join` x = x
```

Dependent function types (often referred to as Π types) in Idris can express such requirements. If `A` is a type and `B` is a type indexed by a value of type `A` then `(x : A) -> B` is the type of functions that map arguments `x` of type `A` to values of type `B x`.

A lattice induces a partial ordering, which gives a direct way to express flow constraints. We introduce a verified partial ordering together with an implementation of this for `JoinSemilattice`. That is, to define an instance of the `Poset` interface we require a concrete instance of an associated data type `leq` as well as proofs of necessary algebraic properties of `leq`.

```
interface Poset a where
  leq : a -> a -> Type
  reflexive     : (x : a) -> x `leq` x
  antisymmetric : (x, y : a) -> x `leq` y -> y `leq` x -> x = y
  transitive    : (x, y, z : a) -> x `leq` y -> y `leq` z -> x `leq` z

implementation JoinSemilattice a => Poset a where
  leq x y = (x `join` y = y)
  ...
```

This definition allows for generic functions to impose as few restrictions as possible on the user while being able to exploit the algebraic structure in proofs, as will become evident in Section 3 and 4. For the sake of the following case studies, we also have a definition of a `BoundedJoinSemilattice` requiring a least element `Bottom` of an instance of `JoinSemilattice` and a proof of the element being the unit.

The Core API Figure 1 presents the type signature of DepSec's core API. Notice that names beginning with a lower case letter that appear as a parameter or index in a type declaration will be automatically bound as an implicit argument in Idris, and the `auto` annotation on implicit arguments means that Idris will attempt to fill in the implicit argument by searching the calling context for an appropriate value.

Abstract data type `Labeled` ℓ a denotes a value of type a with sensitivity level ℓ. We say that `Labeled` ℓ a is *indexed* by ℓ and *parameterized* by a. Abstract data type `DIO` ℓ a denotes a secure computation that handles values with sensitivity level ℓ and

[1] Interfaces in Idris are similar to type classes in Haskell.

```
data Labeled : label -> Type -> Type where
  MkLabeled : valueType -> Labeled label valueType -- TCB

data DIO : l -> Type -> Type where
  MkDIO : IO valueType -> DIO l valueType -- TCB

Monad (DIO l) where
  ...

label : Poset label => {l : label} -> a -> Labeled l a

unlabel : Poset label => {l, l' : label}
       -> {auto flow : l `leq` l'}
       -> Labeled l a
       -> DIO l' a

plug : Poset label => {l, l' : label}
    -> DIO l' a
    -> {auto flow : l `leq` l'}
    -> DIO l (Labeled l' a)

run : DIO l a -> IO a -- TCB

lift : IO a -> DIO l a -- TCB
```

Fig. 1. Type signature of the core DepSec API.

results in a value of type a. It is internally represented as a wrapper around the regular `IO` monad that, similar to the one in Haskell, can be thought of as a state monad where the state is the entire world. Notice that both data constructors `MkLabeled` and `MkDIO` are not available to untrusted code as this would allow pattern matching and uncontrolled unwrapping of protected entities. As a consequence, we introduce functions `label` and `unlabel` for labeling and unlabeling values. Like Rajani and Garg [32], but unlike **MAC**, the type signature of `label` imposes no lattice constraints on the computation context. This does not leak information as, if $l \sqsubseteq l'$ and a computation c has type `DIO` l' (`Labeled` l V) for any type V, then there is no way for the labeled return value of c to escape the computation context with label l'.

As in **MAC**, the API contains a function `plug` that safely integrates sensitive computations into less sensitive ones. This avoids the need for nested computations and *label creep*, that is, the raising of the current label to a point where the computation can no longer perform useful tasks [33, 46]. Finally, we also have functions `run` and `lift` that are only available to trusted code for unwrapping of the `DIO` ℓ monad and lifting of the `IO` monad into the `DIO` ℓ monad.

Labeled resources Data type `Labeled` ℓ a is used to denote a labeled Idris value with type a. This is an example of a *labeled resource* [33]. By itself, the core library does not allow untrusted code to perform any side effects but we can safely incorporate, for

example, file access and mutable references as other labeled resources. Figure 2 presents type signatures for files indexed by security levels used for secure file handling while mutable references are available in Appendix C. Abstract data type `SecFile` ℓ denotes a secure file with sensitivity level ℓ. As for `Labeled` ℓ a, the data constructor `MkSecFile` is not available to untrusted code.

The function `readFile` takes as input a secure file `SecFile l'` and returns a computation with sensitivity level `l` that returns a labeled value with sensitivity level `l'`. Notice that the `l` $\sqsubseteq$ `l'` flow constraint is required to enforce the *no read-up* policy [7]. That is, the result of the computation returned by `readFile` only involves data with sensitivity at most `l`. The function `writeFile` takes as input a secure file `SecFile l''` and a labeled value of sensitivity level `l'`, and it returns a computation with sensitivity level `l` that returns a labeled value with sensitivity level `l''`. Notice that both the `l` $\sqsubseteq$ `l'` and `l'` $\sqsubseteq$ `l''` flow constraints are required, essentially enforcing the *no write-down* policy [7], that is, the file never receives data more sensitive than its sensitivity level.

Finally, notice that the standard library functions for reading and writing files in Idris used to implement the functions in Figure 2 do not raise exceptions. Rather, both functions return an instance of the sum type `Either`. We stay consistent with Idris' choice for this instead of adding exception handling as done in **MAC**.

```
data SecFile : {label : Type} -> (l : label) -> Type where
  MkSecFile : (path : String) -> SecFile l -- TCB

readFile : Poset label => {l,l' : label}
        -> {auto flow : l `leq` l'}
        -> SecFile l'
        -> DIO l (Labeled l' (Either FileError String))

writeFile : Poset label => {l,l',l'' : label}
         -> {auto flow : l `leq` l'} -> {auto flow' : l' `leq` l''}
         -> SecFile l''
         -> Labeled l' String
         -> DIO l (Labeled l'' (Either FileError ()))
```

Fig. 2. Type signatures for secure file handling.

3 Case study: Conference manager system

This case study showcases the expressiveness of DepSec by reimplementing a conference manager system with a fine-grained data-dependent security policy introduced by Lourenço and Caires [22]. Lourenço and Caires base their development on a minimal λ-calculus with references and collections and they show how secure operations on relevant scenarios can be modelled and analysed using *dependent information flow types*. Our reimplementation demonstrates how DepSec matches the expressiveness of such a special-purpose built dependent type system on a key example.

In this scenario, a user is either a regular user, an author user, or a program committee (PC) member. The conference manager contains information about the users, their submissions, and submission reviews. This data is stored in lists of references to records, and the goal is to statically ensure, by typing, the confidentiality of the data stored in the conference manager system. As such, the security policy is:

- A registered user's information is not observable by other users.
- The content of a paper can be seen by its authors as well as its reviewers.
- Comments to the PC of a submission's review can only be seen by other members that are also reviewers of that submission.
- The only authors that are allowed to see the grade and the review of the submission are those that authored that submission.

To achieve this security policy, Lourenço and Caires make use of indexed security labels [21]. The security level U is partitioned into a number of security compartments such that $U(uid)$ represents the compartment of the registered user with id uid. Similarly, the security level A is indexed such that $A(uid, sid)$ stands for the compartment of data belonging to author uid and their submission sid, and PC is indexed such that $PC(uid, sid)$ stands for data belonging to the PC member with user id uid assigned to review the submission with id sid. Furthermore, levels $\top$ and $\bot$ are introduced such that, for example, $U(\bot) \sqsubseteq U(uid) \sqsubseteq U(\top)$. Now, the security lattice is defined using two equations:

$$\forall uid, sid.\ U(uid) \sqsubseteq A(uid, sid) \tag{1}$$

$$\forall uid1, uid2, sid.\ A(uid1, sid) \sqsubseteq PC(uid2, sid) \tag{2}$$

Lourenço and Caires are able to type a list of submissions with a dependent sum type that assigns the content of the paper the security level $A(uid, sid)$, where uid and sid are fields of the record. For example, if a concrete submission with identifier 2 was made by the user with identifier 1, the content of the paper gets classified at security level $A(1, 2)$. In consequence, $A(1, 2) \sqsubseteq PC(n, 2)$ for any $uid\ n$ and the content of the paper is only observable by its assigned reviewers. Similar types are given for the list of user information and the list of submission reviews, enforcing the security policy described in the above.

To express this policy in DepSec, we introduce abstract data types `Id` and `Compartment` (cf. Figure 3) followed by an implementation of the `BoundedJoinSemilattice` interface that satisfies equations (1) and (2).

```
data Id : Type where
  Top : Id
  Nat : Nat -> Id
  Bot : Id
```

```
data Compartment : Type where
  U  : Id -> Compartment
  A  : Id -> Id -> Compartment
  PC : Id -> Id -> Compartment
```

Fig. 3. Abstract data types for the conference manager sample security lattice.

Using the above, the required dependent sum types can easily be encoded with DepSec in Idris as presented in Figure 4. With these typings in place, implementing the

```
record User where
  constructor MkUser
  uid   : Id
  name  : Labeled (U uid) String
  univ  : Labeled (U uid) String
  email : Labeled (U uid) String

record Submission where
  constructor MkSubmission
  uid   : Id
  sid   : Id
  title : Labeled (A uid sid) String
  abs   : Labeled (A uid sid) String
  paper : Labeled (A uid sid) String

record Review where
  constructor MkReview
  uid     : Id
  sid     : Id
  PC_only : Labeled (PC uid sid) String
  review  : Labeled (A Top sid) String
  grade   : Labeled (A Top sid) Integer
```

Fig. 4. Conference manager types encoded with DepSec.

examples from Lourenço and Caires [22] is straightforward. For example, the function `viewAuthorPapers` takes as input a list of submissions and a user identifier `uid1` from which it returns a computation that returns a list of submissions authored by the user with identifier `uid1`. Notice that `uid` denotes the automatically generated record projection function that retrieves the field `uid` of the record, and that `(x: A ** B)` is notation for a dependent pair (often referred to as a Σ type) where `A` and `B` are types and `B` may depend on `x`.

```
viewAuthorPapers : Submissions
                -> (uid1 : Id)
                -> DIO Bottom (List (sub : Submission ** uid1 = (uid sub)))
```

The `addCommentSubmission` operation is used by the PC members to add comments to the submissions. The function takes as input a list of reviews, a user identifier of a PC member, a submission identifier, and a comment with label `A uid1 sid1`. It returns a computation that updates the `PC_only` field in the review of the paper with identifier `sid1`.

```
addCommentSubmission : Reviews -> (uid1 : Id) -> (sid1 : Id)
                    -> Labeled (A uid1 sid1) String
                    -> DIO Bottom ()
```

Notice that to implement this specific type signature, up-classification is necessary to assign the comment with type `Labeled (A uid1 sid1) String` to a field with type `Labeled (PC uid sid1) String`. This can be achieved soundly with the `relabel` primitive introduced by Vassena et al. [46] as `A uid1 sid1` $\sqsubseteq$ `PC uid sid1`. We include this primitive in Appendix C. Several other examples are available in the accompanying source code. The entire case study amounts to about 300 lines of code where half of the lines implement and verify the lattice.

4 Policy-parameterized functions

A consequence of using a dependently typed language, and the design of DepSec, is that functions can be defined such that they abstract over the security policy while retaining precise security levels. This makes it possible to reuse code across different applications and write other libraries on top of DepSec. We can exploit the existence of a lattice `join`, the induced poset, and their verified algebraic properties to write such functions.

```
readTwoFiles : BoundedJoinSemilattice label
            => {l, l' : label}
            -> SecFile l
            -> SecFile l'
            -> DIO Bottom (Labeled (join l l') (Either FileError String))
readTwoFiles file1 file2 {l} {l'} =
  do file1' <- readFile {flow = leq_bot_x l} file1
     file2' <- readFile {flow = leq_bot_x l'} file2
     let dio : DIO (join l l') (Either FileError String)
       = do c1 <- unlabel {flow = join_x_xy l l'} file1'
            c2 <- unlabel {flow = join_y_xy l l'} file2'
            pure $ case (c1, c2) of
                        (Right c1', Right c2') => Right $ c1' ++ c2'
                        (Left e1, _) => Left e1
                        (_, Left e2) => Left e2
     plug {flow = leq_bot_x (join l l')} dio
```

Fig. 5. Reading two files to a string labeled with the join of the labels of the files.

Figure 5 presents the function `readTwoFiles` that is parameterized by a bounded join semilattice. It takes two secure files with labels `l` and `l'` as input and returns a computation that concatenates the contents of the two files labeled with the join of `l` and `l'`. To implement this, we make use of the `unlabel` and `readFile` primitives from Figure 1 and 2, respectively. This computation unlabels the contents of the files and returns the concatenation of the contents if no file error occurred. Notice that `pure` is the Idris function for monadic return, corresponding to the `return` function in Haskell. Finally, this computation is plugged into the surrounding computation. Notice how the usage of `readFile` and `unlabel` introduces several proof obligations, namely $\bot \sqsubseteq$ `l`, `l'`, `l` $\sqcup$ `l'` and `l`, `l'` $\sqsubseteq$ `l` $\sqcup$ `l'`. When working on a concrete lattice these obligations are usually fulfilled by Idris' automatic proof search but, currently, such proofs need to be given manually in the general case. All obligations follow immediately from the algebraic properties of the bounded semilattice and are given in three auxiliary lemmas `leq_bot_x`, `join_x_xy`, and `join_y_xy` available in Appendix C (amounting to 10 lines of code).

Writing functions operating on a fixed number of resources is limiting. However, the function in Figure 5 can easily be generalized to a function working on an arbitrary data structure containing files with different labels from an arbitrary lattice. Similar to

the approach taken by Buiras et al. [11] that hide the label of a labeled value using a data type definition, we hide the label of a secure file with a dependent pair

```
GenFile : Type -> Type
GenFile label = (l : label ** SecFile l)
```

that abstracts away the concrete sensitivity level of the file. Moreover, we introduce a specialized join function

```
joinOfFiles : BoundedJoinSemilattice label
           => List (GenFile label)
           -> label
```

that folds the `join` function over a list of file sensitivity labels. Now, it is possible to implement a function that takes as input a list of files, reads the files, and returns a computation that concatenates all their contents (if no file error occurred) where the return value is labeled with the join of all their sensitivity labels.

```
readFiles : BoundedJoinSemilattice a
         => (files: (List (GenFile a)))
         -> DIO Bottom (Labeled (joinOfFiles files)
                                (Either (List FileError) String))
```

When implementing this, one has to satisfy non-trivial proof obligations as, for example, that $l \sqsubseteq$ `joinOfFiles(files)` for all secure files $f \in$ `files` where the label of f is l. While provable (in 40 lines of code in our development), if equality is decidable for elements of the concrete lattice we can postpone such proof obligations to a point in time where it can be solved by reflexivity of equality. By defining a decidable lattice order

```
decLeq : JoinSemilattice a => DecEq a => (x, y : a) -> Dec (x `leq` y)
decLeq x y = decEq (x `join` y) y
```

we can get such a proof "for free" by inserting a dynamic check of whether the flow is allowed. With this, a `readFiles'` function with the exact same functionality as the original `readFiles` function can be implemented with minimum effort. In the below, `prf` is the proof that the label `l` of `file` may flow to `joinOfFiles files` .

```
readFiles' : BoundedJoinSemilattice a => DecEq a
          => (files: (List (GenFile a)))
          -> DIO Bottom (Labeled (joinOfFiles files)
                                 (Either (List FileError) String))
readFiles' files =
  ...
  case decLeq l (joinOfFiles files) of
    Yes prf => ...
    No _ => ...
```

The downside of this is the introduction of a negative case, the `No`-case, that needs handling even though it will never occur if `joinOfFiles` is implemented correctly.

In combination with `GenFile`, `decLeq` can be used to implement several other interesting examples. For instance, a function that reads all files with a sensitivity label below a certain label to a string labeled with that label. The accompanying source code showcases multiple such examples that exploit decidable equality.

5 Declassification

Realistic applications often release some secret information as part of their intended behavior; this action is known as *declassification*.

In DepSec, trusted code may declassify secret information without adhering to any security policy as trusted code has access to both the `DIO` ℓ a and `Labeled` ℓ a data constructors. However, only giving trusted code the power of declassification is limiting as we want to allow the use of third-party code as much as possible. The main challenge we address is how to grant untrusted code the right amount of power such that declassification is only possible in the intended way.

Sabelfeld and Sands [37] identify four dimensions of declassification: *what*, *who*, *where*, and *when*. In this section, we present novel and powerful means for static declassification with respect to three of the four dimensions and illustrate these with several examples. To statically enforce different declassification policies we take the approach of Sabelfeld and Myers [36] and use escape hatches, a special kind of functions. In particular, we introduce the notion of a *hatch builder*; a function that creates an escape hatch for a particular resource and which can only be used when a certain condition is met. Such an escape hatch can therefore be used freely by untrusted code.

5.1 The *what* dimension

Declassification policies related to the *what* dimension place restrictions on exactly "what" and "how much" information is released. It is in general difficult to statically predict how data to be declassified is manipulated or changed by programs [34] but exploiting dependent types can get us one step closer.

To control what information is released, we introduce the notion of a *predicate hatch builder* only available to trusted code for producing hatches for untrusted code.

```
predicateHatchBuilder : Poset lt => {l, l' : lt} -> {D, E : Type}
                     -> (d : D)
                     -> (P : D -> E -> Type)
                     -> (d : D ** Labeled l (e : E ** P d e)
                                 -> Labeled l' E) -- TCB
```

Intuitively, the hatch builder takes as input a data structure `d` of type `D` followed by a predicate `P` upon `d` and something of type `E`. It returns a dependent pair of the initial data structure and a declassification function from sensitivity level `l` to `l'`. To actually declassify a labeled value `e` of type `E` one has to provide a proof that `P d e` holds. Notice that this proof may be constructed in the context of the sensitivity level `l` that we are declassifying from.

The reason for parameterizing the predicate `P` by a data structure of type `D` is to allow declassification to be restricted to a specific context or data structure. This is used in the following example of an auction system, in which only the highest bid of a specific list of bids can be declassified.

Example Consider a two point lattice where `L` $\sqsubseteq$ `H`, `H` $\not\sqsubseteq$ `L` and an auction system where participants place bids secretly. All bids are labeled `H` and are put into a data structure `BidLog`. In the end, we want only the winning bid to be released and hence declassified to label `L`. To achieve this, we define a declassification predicate `HighestBid`.

```
HighestBid : BidLog -> Bid -> Type
HighestBid = \log, b => (Elem (label b) log, MaxBid b log)
```

Informally, given a log `log` of labeled bids and a bid `b`, the predicate states that the bid is in the log, `Elem (label b) log`, and that it is the maximum bid, `MaxBid b log`. We apply `predicateHatchBuilder` to a log of bids and the `HighestBid` predicate to obtain a specialized escape hatch of type `BidHatch` that enforces the declassification policy defined by the predicate.

```
BidHatch : Type
BidHatch = (log : BidLog ** Labeled H (b : Bid ** HighestBid log b)
                           -> Labeled L Bid)
```

This hatch can be used freely by untrusted code when implementing the auction system. By constructing a function

```
getMaxBid : (r : BidLog) -> DIO H (b : Bid ** HighestBid r b)
```

untrusted code can plug the resulting computation into an `L` context and declassify the result value using the argument `hatch` function.

```
auction : BidHatch -> DIO L (Labeled L Bid)
auction ([] ** _) = pure $ label ("no bids", 0)
auction (r :: rs ** hatch) =
  do max <- plug (getMaxBid (r :: rs))
    let max' : Labeled L Bid = hatch max
    ...
```

To show the `HighestBid` predicate (which in our implementation comprises 40 lines of code), untrusted code will need a generalized `unlabel` function that establishes the relationship between `label` and the output of `unlabel`. The only difference is its return type: a computation that returns a value and a proof that when labeling this value we will get back the initial input. This definition poses no risk to soundness as the proof is protected by the computation sensitivity level.

```
unlabel' : Poset lt => {l,l': lt}
        -> {auto flow: l `leq` l'}
        -> (labeled: Labeled l a)
        -> DIO l' (c : a ** label c = labeled)
```

Limiting hatch usage Notice how escape hatches, generally, can be used an indefinite number of times. The `Control.ST` library [10] provides facilities for creating, reading, writing, and destroying state in the type of Idris functions and, especially, allows tracking of state change in a function type. This allows us to limit the number of

times a hatch can be used. Based on a concept of resources, a dependent type `STrans` tracks how resources change when a function is invoked. Specifically, a value of type `STrans m returnType in_res out_res` represents a sequence of actions that manipulate state where `m` is an underlying computation context in which the actions will be executed, `returnType` is the return type of the sequence, `in_res` is the required list of resources available before executing the sequence, and `out_res` is the list of resources available after executing the sequence.

To represent state transitions more directly, `ST` is a type level function that computes an appropriate `STrans` type given a underlying computation context, a result type, and a list of *actions*, which describe transitions on resources. Actions can take multiple forms but the one we will make use of is of the form `lbl ::: ty_in :-> ty_out` that expresses that the resource `lbl` begins in state `ty_in` and ends in state `ty_out`. By instantiating `ST` with `DIO l` as the underlying computation context:

```
DIO' : l -> (ty : Type) -> List (Action ty) -> Type
DIO' l = ST (DIO l)
```

and use it together with a resource `Attempts`, we can create a function `limit` that applies its first argument `f` to its second argument `arg` with `Attempts (S n)` as its initial required state and `Attempts n` as the output state.

```
limit : (f : a -> b) -> (arg : a)
     -> DIO' l b [attempts ::: Attempts (S n) :-> Attempts n]
```

That is, we encode that the function consumes "an attempt." With the `limit` function it is possible to create functions where users are forced, by typing, to specify how many times it is used.

As an example, consider a variant of an example by Russo et al. [34] where we construct a specialized hatch `passwordHatch` that declassifies the boolean comparison of a secret number with an arbitrary number.

```
passwordHatch : (labeled : Labeled H Int)
             -> (guess : Int)
             -> DIO' l Bool [attempts ::: Attempts (S n) :-> Attempts n]
passwordHatch (MkLabeled v) = limit (\g => g == v)
```

To use this hatch, untrusted code is forced to specify how many times it is used.

```
pwCheck : Labeled H Int
       -> DIO' L () [attempts ::: Attempts (3 + n) :-> Attempts n]
pwCheck pw =
  do x1 <- passwordHatch pw 1
     x2 <- passwordHatch pw 2
     x3 <- passwordHatch pw 3
     x4 <- passwordHatch pw 4 -- type error!
     ...
```

5.2 The *who* and *when* dimensions

To handle declassification policies related to *who* may declassify information and *when* declassification may happen we introduce the notion of a *token hatch builder* only available to trusted code for producing hatches for untrusted code to use.

```
tokenHatchBuilder : Poset labelType => {l, l' : labelType} -> {E, S : Type}
                 -> (Q : S -> Type)
                 -> (s : S ** Q s) -> Labeled l E -> Labeled l' E -- TCB
```

The hatch builder takes as input a predicate `Q` on something of type `S` and returns a declassification function from sensitivity level `l` to `l'` given that the user can prove the existence of some `s` such that `Q s` holds. As such, by limiting when and how untrusted can obtain a value that satisfy predicate `Q`, we can construct several interesting declassification policies.

The rest of this section discusses how predicate hatches can be used for time-based and authority-based control of declassification; the use of the latter is demonstrated on a case study.

Time-based hatches To illustrate the idea of token hatches for the *when* dimension of declassification, consider the following example. Let `Time` be an abstract data type with a data constructor only available to trusted code and `tick : DIO l Time` a function that returns the current system time wrapped in the `Time` data type such that this is the only way for untrusted code to construct anything of type `Time`. Notice that this does not expose an unrestricted timer API as untrusted code can not inspect the actual value.

Now, we instantiate the token hatch builder with a predicate that demands the existence of a `Time` token that is greater than some specific value.

```
TimeHatch : Time -> Type
TimeHatch t = (t' ** t <= t' = True) -> Labeled H Nat -> Labeled L Nat
```

As such, `TimeHatch t` can only be used after a specific point in time `t` has passed as only then untrusted code will be able to satisfy the predicate.

```
timer : Labeled H Nat -> TimeHatch t -> DIO L ()
timer secret {t} timeHatch =
  do time <- tick
     case decEq (t <= time) True of
       Yes prf =>
         let declassified : Labeled L Nat = timeHatch (time ** prf) secret
         ...
       No _ => ...
```

Authority-based hatches The *Decentralized Labeling Model* (DLM) [26] marks data with a set of principals who owns the information. While executing a program, the program is given *authority*, that is, it is authorized to act on behalf of some set of principals. Declassification simply makes a copy of the released data and marks it with the same set of principals but excludes the authorities.

Similarly to Russo et al. [34], we adapt this idea such that it works on a security lattice of `Principal`s, assign authorities with security levels from the lattice, and let authorities declassify information at that security level.

To model this, we define the abstract data type `Authority` with a data constructor available only to trusted code so that having an instance of `Authority s` corresponds to having the authority of the principal `s`. Notice how assignment of authorities to pieces of code consequently is a part of the trusted code. Now, we instantiate the token hatch builder with a predicate that demands the authority of `s` to declassify information at that level.

```
authHatch : { l, l' : Principal }
         -> (s ** (l = s, Authority s))
         -> Labeled l a -> Labeled l' a
authHatch {l} = tokenHatchBuilder (\s => (l = s, Authority s))
```

That is, `authHatch` makes it possible to declassify information at level `l` to `l'` given an instance of the `Authority l` data type.

Example Consider the scenario of an online dating service that has the distinguishing feature of allowing its users to specify the visibility of their profiles at a fine-grained level. To achieve this, the service allows users to provide a *discovery agent* that controls their visibility. Consider a user, Bob, whose implementation of the discovery agent takes as input his own profile and the profile of another user, say Alice. The agent returns a possibly side-effectful computation that returns an option type indicating whether Bob wants to be discovered by Alice. If that is the case, a profile is returned by the computation with the information about Bob that he wants Alice to be able to see. When Alice searches for candidate matches, her profile is run against the discovery agents of all candidates and the result is added to her browsing queue.

To implement this dating service, we define the record type `ProfileInfo A` that contains personal information related to principal `A`.

```
record ProfileInfo (A : Principal) where
  constructor MkProfileInfo
  name      : Labeled A String
  gender    : Labeled A String
  birthdate : Labeled A String
  ...
```

The interesting part of the dating service is the implementation of discovery agents. Figure 6 presents a sample discovery agent that matches all profiles with the opposite gender and only releases information about the name and gender. The discovery agent demands the authority of `A` and takes as input two profiles `a : ProfileInfo A` and `b : ProfileInfo B`. The resulting computation security level is `B` so to incorporate information from `a` into the result, declassification is needed. This is achieved by providing `authHatch` with the authority proof of `A`. The discovery agent `sampleDiscoverer` in Figure 6 unlabels `B`'s gender, declassifies and unlabels `A`'s gender and name, and compares the two genders. If the genders match, a profile with type `ProfileInfo B` only containing the name and gender of `A` is returned. Otherwise, `Nothing` is returned

indicating that `A` does not want to be discovered. Notice that `Refl` is the constructor for the built-in equality type in Idris and it is used to construct the proof of equality between principals required by the hatch.

```
sampleDiscoverer : {A, B : Principal}
                -> Authority A
                -> (a : ProfileInfo A)
                -> (b : ProfileInfo B)
                -> DIO B (Maybe (ProfileInfo B))
sampleDiscoverer {A} {B} auth a b =
  do bGender <- unlabel $ gender b
     aGender <- unlabel $ authHatch (A ** (Refl, auth)) (gender a)
     aName <- unlabel $ authHatch (A ** (Refl, auth)) (name a)
     case decEq bGender aGender of
       Yes _ => pure Nothing
       No _  => pure (Just (MkProfileInfo aName aGender "" "" ""))
```

Fig. 6. A discovery agent that matches with all profiles of the opposite gender and only releases the name and gender.

6 Soundness

Recent works [45, 46] present a mechanically-verified model of **MAC** and show progress-insensitive noninterference (PINI) for a sequential calculus. We use this work as a starting point and discuss necessary modification in the following. Notice that this work does not consider any declassification mechanisms and neither do we; we leave this as future work.

The proof relies on the *two-steps erasure* technique, an extension of the *term erasure* [20] technique that ensures that the same public output is produced if secrets are erased before or after program execution. The technique relies on a type-driven erasure function ε_{ℓ_A} on terms and configurations where ℓ_A denotes the attacker security level. A configuration consists of an ℓ-indexed compartmentalized store Σ and a term t. A configuration $\langle \Sigma, t \rangle$ is erased by erasing t and by erasing Σ pointwise, i.e. $\varepsilon_{\ell_A}(\Sigma) = \lambda \ell . \varepsilon_{\ell_A}(\Sigma(\ell))$. On terms, the function essentially rewrites data and computations above ℓ_A to a special $\bullet$ value. The full definition of the erasure function is available in Appendix A.5. From this definition, the definition of low-equivalence of configurations follows.

Definition 1. *Let c_1 and c_2 be configurations. c_1 and c_2 are said to be ℓ_A-equivalent, written $c_1 \approx_{\ell_A} c_2$, if and only if $\varepsilon_{\ell_A}(c_1) \equiv \varepsilon_{\ell_A}(c_2)$.*

After defining the erasure function, the noninterference theorem follows from showing a *single-step simulation* relationship between the erasure function and a small-step reduction relation: erasing sensitive data from a configuration and then taking a step is the same as first taking a step and then erasing sensitive data. This is the content of the following proposition.

Proposition 1. *If* $c_1 \approx_{\ell_A} c_2$, $c_1 \to c_1'$, *and* $c_2 \to c_2'$ *then* $c_1' \approx_{\ell_A} c_2'$.

The main theorem follows by repeated applications of Proposition 1.

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

Both the statement and the proof of noninterference for DepSec are mostly similar to the ones for **MAC** and available in Appendix B. Nevertheless, one has to be aware of a few subtleties.

First, one has to realize that even though dependent types in a language like Idris may depend on data, the data itself is not a part of a value of a dependent type. Recall the type `Vect n Nat` of vectors of length `n` with components of type `Nat` and consider the following program.

```
length : Vect n a -> Nat
length {n = n} xs = n
```

This example may lead one to believe that it is possible to extract data from a dependent type. This is *not* the case. Both `n` and `a` are implicit arguments to the `length` function that the compiler is able to infer. The actual type is

```
length : {n : Nat} -> {a : Type} -> Vect n a -> Nat
```

As a high-level dependently typed functional programming language, Idris is elaborated to a low-level core language based on dependent type theory [9]. In the elaboration process, such implicit arguments are made explicit when functions are defined and inferred when functions are invoked. This means that in the underlying core language, only explicit arguments are given. Our modeling given in Appendix A.1 reflects this fact soundly.

Second, to model the extended expressiveness of DepSec, we extend both the semantics and the type system with compile-time pure-term reduction and higher-order dependent types. These definitions are standard (defined for Idris by Brady [9]) and available in Appendix A.2 and A.3. Moreover, as types now become first-class terms, the definition of ε_{ℓ_A} has to be extended to cover the new kinds of terms. As before, primitive types are unaffected by the erasure function, but dependent and indexed types, such as the type `DIO`, have to be erased homomorphically, e.g., ε_{ℓ_A} (`DIO` ℓ τ : `Type`) $\triangleq$ `DIO` $\varepsilon_{\ell_A}(\ell)$ $\varepsilon_{\ell_A}(\tau)$. The intuition of why this is sensible comes from the observation that indexed dependent types considered as terms may contain values that will have to be erased. This is purely a technicality of the proof. If defined otherwise, the erasure function would not commute with capture-avoiding substitution on terms, $\varepsilon_{\ell_A}(t[v/x]) = \varepsilon_{\ell_A}(t)[\varepsilon_{\ell_A}(v)/x]$, which is vital for the remaining proof.

7 Related work

Security libraries The pioneering and formative work by Li and Zdancewic [19] shows how *arrows* [17], a generalization of monads, can provide information-flow control without runtime checks as a library in Haskell. Tsai et al. [44] further extend this work to handle side-effects, concurrency, and heterogeneous labels. Russo et al. [34] eliminate the need for arrows and implement the security library **SecLib** in Haskell based

solely on monads. Rather than labeled values, this work introduces a monad which statically label side-effect free values. Furthermore, it presents combinators to dynamically specify and enforce declassification policies that bear a resemblance to the policies that DepSec are able to enforce statically.

The security library **LIO** [40, 41] dynamically enforces information-flow control in both sequential and concurrent settings. Stefan et al. [39] extend the security guarantees of this work to also cover exceptions. Similar to this work, Stefan et al. [41] present a simple API for implementing secure conference reviewing systems in **LIO** with support for data-dependent security policies.

Inspired by the design of **SecLib** and **LIO**, Russo [33] introduces the security library **MAC**. The library statically enforces information-flow control in the presence of advanced features like exceptions, concurrency, and mutable data structures by exploiting Haskell's type system to impose flow constraints. Vassena and Russo [45], Vassena et al. [46] show progress-insensitive noninterference for **MAC** in a sequential setting and progress-sensitive noninterference in a concurrent setting, both using the two-steps erasure technique.

The flow constraints enforcing confidentiality of read and write operations in DepSec are identical to those of **MAC**. This means that the examples from **MAC** that do not involve concurrency can be ported directly to DepSec. To the best of our knowledge, data-dependent security policies like the one presented in Section 3 cannot be expressed and enforced in **MAC**, unlike **LIO** that allows such policies to be enforced dynamically. DepSec allows for such security policies to be enforced statically. Moreover, Russo [33] does not consider declassification. To address the static limitations of **MAC**, **HLIO** [11] takes a hybrid approach by exploiting advanced features in Haskell's type-system like singleton types and constraint polymorphism. Buiras et al. [11] are able to statically enforce information-flow control while allowing selected security checks to be deferred until run-time.

Dependent types for security Several works have considered the use of dependent types to capture the nature of data-dependent security policies. Zheng and Myers [50, 51] proposed the first dependent security type system for dealing with dynamic changes to runtime security labels in the context of Jif [28], a full-fledged IFC-aware compiler for Java programs, where similar to our work, operations on labels are modeled at the level of types. Zhang et al. [49] use dependent types in a similar fashion for the design of a hardware description language for timing-sensitive information-flow security.

A number of functional languages have been developed with dependent type systems and used to encode value-dependent information flow properties, e.g. Fine [42]. These approaches require the adoption of entirely new languages and compilers where DepSec is embedded in an already existing language. Morgenstern and Licata [24] encode an authorization and IFC-aware programming language in Agda. However, their encoding does not consider side-effects. Nanevski et al. [29] use dependent types to verify information flow and access control policies in an interactive manner.

Lourenço and Caires [22] introduce the notion of *dependent information-flow types* and propose a *fine-grained* type system; every value and function have an associated security level. Their approach is different to the *coarse-grained* approach taken in our work where only some computations and values have associated security labels. Rajani

and Garg [32] show that both approaches are equally expressive for static IFC techniques and Vassena et al. [47] show the same for dynamic IFC techniques.

Principles for Information Flow Bastys et al. [6] put forward a set of informal principles for information flow security definitions and enforcement mechanisms: *attacker-driven security, trust-aware enforcement, separation of policy annotations and code, language-independence, justified abstraction, and permissiveness*.

DepSec follows the principle of trust-aware enforcement, as we make clear the boundary between the trusted and untrusted components in the program. Additionally, the design of our declassification mechanism follows the principle of separation of policy annotations and code. The use of dependent types increases the permissiveness of our enforcement as we discuss throughout the paper. While our approach is not fully language-independent, we posit that the approach may be ported to other programming languages with general-purpose dependent types.

Declassification enforcement Our hatch builders are reminiscent of downgrading policies of Li and Zdancewic [18]. For example, similar to them, DepSec's declassification policies naturally express the idea of *delimited release* [35] that provides explicit characterization of the declassifying computation. Here, DepSec's policies can express a broad range of policies that can be expressed through predicates, an improvement over simple expression-based enforcement mechanisms for delimited release [4, 5, 35].

An interesting point in the design of declassification policies is *robust declassification* [48] that demands that untrusted components must not affect information release. *Qualified robustness* [2, 27] generalizes this notion by giving untrusted code a limited ability to affect information release through the introduction of an explicit endorsement operation. Our approach is orthogonal to both notions of robustness as the intent is to let the untrusted components declassify information but only under very controlled circumstances while adhering to the security policy.

8 Conclusion and future work

In this paper, we have presented DepSec – a library for statically enforced information-flow control in Idris. Through several case studies, we have showcased how the DepSec primitives increase the expressiveness of state-of-the-art information-flow control libraries and how DepSec matches the expressiveness of a special-purpose dependent information-flow type system on a key example. Moreover, the library allows programmers to implement policy-parameterized functions that abstract over the security policy while retaining precise security levels.

By taking ideas from the literature and by exploiting dependent types, we have shown powerful means of specifying statically enforced declassification policies related to *what*, *who*, and *when* information is released. Specifically, we have introduced the notion of predicate hatch builders and token hatch builders that rely on the fulfillment of predicates and possession of tokens for declassification. We have also shown how the `ST` monad [10] can be used to limit hatch usage statically.

Finally, we have discussed the necessary means to show progress-insensitive noninterference in a sequential setting for a dependently typed information-flow control library like DepSec.

Future work There are several avenues for further work. Integrity is vital in many security policies and is not considered in **MAC** nor DepSec. It will be interesting to take integrity and the presence of concurrency into the dependently typed setting and consider internal and termination covert channels as well. It also remains to prove our declassification mechanisms sound. Here, attacker-centric epistemic security conditions [3, 15] that intuitively express many declassification policies may be a good starting point.

Acknowledgements Thanks are due to Mathias Vorreiter Pedersen, Bas Spitters, Alejandro Russo, and Marco Vassena for their valuable insights and the anonymous reviewers for their comments on this paper. This work is partially supported by DFF project 6108-00363 from The Danish Council for Independent Research for the Natural Sciences (FNU), Aarhus University Research Foundation, and the Concordium Blockchain Research Center, Aarhus University, Denmark.

Bibliography

[1] Askarov, A., Hunt, S., Sabelfeld, A., Sands, D.: Termination-insensitive noninterference leaks more than just a bit. In: Jajodia, S., López, J. (eds.) Computer Security - ESORICS 2008, 13th European Symposium on Research in Computer Security, Málaga, Spain, October 6-8, 2008. Proceedings. Lecture Notes in Computer Science, vol. 5283, pp. 333–348. Springer (2008), https://doi.org/10.1007/978-3-540-88313-5_22

[2] Askarov, A., Myers, A.C.: Attacker control and impact for confidentiality and integrity. Logical Methods in Computer Science 7(3) (2011), https://doi.org/10.2168/LMCS-7(3:17)2011

[3] Askarov, A., Sabelfeld, A.: Gradual release: Unifying declassification, encryption and key release policies. In: 2007 IEEE Symposium on Security and Privacy (S&P 2007), 20-23 May 2007, Oakland, California, USA. pp. 207–221. IEEE Computer Society (2007), https://doi.org/10.1109/SP.2007.22

[4] Askarov, A., Sabelfeld, A.: Localized delimited release: combining the what and where dimensions of information release. In: Hicks, M.W. (ed.) Proceedings of the 2007 Workshop on Programming Languages and Analysis for Security, PLAS 2007, San Diego, California, USA, June 14, 2007. pp. 53–60. ACM (2007), https://doi.org/10.1145/1255329.1255339

[5] Askarov, A., Sabelfeld, A.: Tight enforcement of information-release policies for dynamic languages. In: Proceedings of the 22nd IEEE Computer Security Foundations Symposium, CSF 2009, Port Jefferson, New York, USA, July 8-10, 2009. pp. 43–59. IEEE Computer Society (2009), https://doi.org/10.1109/CSF.2009.22

[6] Bastys, I., Piessens, F., Sabelfeld, A.: Prudent design principles for information flow control. In: Proceedings of the 13th Workshop on Programming Languages and Analysis for Security. pp. 17–23. ACM (2018)

[7] Bell, D.E., La Padula, L.J.: Secure computer system: Unified exposition and multics interpretation. Tech. rep., MITRE CORP BEDFORD MA (1976)

[8] Brady, E.: IDRIS —: systems programming meets full dependent types. In: Jhala, R., Swierstra, W. (eds.) Proceedings of the 5th ACM Workshop Programming Languages meets Program Verification, PLPV 2011, Austin, TX, USA, January 29, 2011. pp. 43–54. ACM (2011), https://doi.org/10.1145/1929529.1929536

[9] Brady, E.: Idris, a general-purpose dependently typed programming language: Design and implementation. J. Funct. Program. 23(5), 552–593 (2013), https://doi.org/10.1017/S095679681300018X

[10] Brady, E.: State machines all the way down (Jan 2016), draft available at idris-lang.org/drafts/sms.pdf

[11] Buiras, P., Vytiniotis, D., Russo, A.: HLIO: mixing static and dynamic typing for information-flow control in haskell. In: Fisher, K., Reppy, J.H. (eds.) Proceedings of the 20th ACM SIGPLAN International Conference on Functional Program-

ming, ICFP 2015, Vancouver, BC, Canada, September 1-3, 2015. pp. 289–301. ACM (2015), https://doi.org/10.1145/2784731.2784758

[12] Chapman, R., Hilton, A.: Enforcing security and safety models with an information flow analysis tool. In: McCormick, J.W., Sward, R.E. (eds.) Proceedings of the 2004 Annual ACM SIGAda International Conference on Ada: The Engineering of Correct and Reliable Software for Real-Time & Distributed Systems using Ada and Related Technologies 2004, Atlanta, GA, USA, November 14-14, 2004. pp. 39–46. ACM (2004), https://doi.org/10.1145/1032297.1032305

[13] Coquand, T., Huet, G.P.: The calculus of constructions. Inf. Comput. 76(2/3), 95–120 (1988), https://doi.org/10.1016/0890-5401(88)90005-3

[14] Denning, D.E., Denning, P.J.: Certification of programs for secure information flow. Commun. ACM 20(7), 504–513 (1977), https://doi.org/10.1145/359636.359712

[15] Halpern, J.Y., O'Neill, K.R.: Secrecy in multiagent systems. ACM Trans. Inf. Syst. Secur. 12(1), 5:1–5:47 (2008), https://doi.org/10.1145/1410234.1410239

[16] Hedin, D., Birgisson, A., Bello, L., Sabelfeld, A.: Jsflow: tracking information flow in javascript and its apis. In: Cho, Y., Shin, S.Y., Kim, S., Hung, C., Hong, J. (eds.) Symposium on Applied Computing, SAC 2014, Gyeongju, Republic of Korea - March 24 - 28, 2014. pp. 1663–1671. ACM (2014), https://doi.org/10.1145/2554850.2554909

[17] Hughes, J.: Generalising monads to arrows. Sci. Comput. Program. 37(1-3), 67–111 (2000), https://doi.org/10.1016/S0167-6423(99)00023-4

[18] Li, P., Zdancewic, S.: Downgrading policies and relaxed noninterference. In: Palsberg, J., Abadi, M. (eds.) Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005. pp. 158–170. ACM (2005), https://doi.org/10.1145/1040305.1040319

[19] Li, P., Zdancewic, S.: Encoding information flow in haskell. In: 19th IEEE Computer Security Foundations Workshop, (CSFW-19 2006), 5-7 July 2006, Venice, Italy. p. 16. IEEE Computer Society (2006), https://doi.org/10.1109/CSFW.2006.13

[20] Li, P., Zdancewic, S.: Arrows for secure information flow. Theor. Comput. Sci. 411(19), 1974–1994 (2010), https://doi.org/10.1016/j.tcs.2010.01.025

[21] Liu, J., George, M.D., Vikram, K., Qi, X., Waye, L., Myers, A.C.: Fabric: a platform for secure distributed computation and storage. In: Matthews, J.N., Anderson, T.E. (eds.) Proceedings of the 22nd ACM Symposium on Operating Systems Principles 2009, SOSP 2009, Big Sky, Montana, USA, October 11-14, 2009. pp. 321–334. ACM (2009), https://doi.org/10.1145/1629575.1629606

[22] Lourenço, L., Caires, L.: Dependent information flow types. In: Rajamani, S.K., Walker, D. (eds.) Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015. pp. 317–328. ACM (2015), https://doi.org/10.1145/2676726.2676994

[23] Martin-Löf, P., Sambin, G.: Intuitionistic type theory, vol. 9. Bibliopolis Naples (1984)
[24] Morgenstern, J., Licata, D.R.: Security-typed programming within dependently typed programming. In: Hudak, P., Weirich, S. (eds.) Proceeding of the 15th ACM SIGPLAN international conference on Functional programming, ICFP 2010, Baltimore, Maryland, USA, September 27-29, 2010. pp. 169–180. ACM (2010), https://doi.org/10.1145/1863543.1863569
[25] Myers, A.C.: Jflow: Practical mostly-static information flow control. In: Appel, A.W., Aiken, A. (eds.) POPL '99, Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Antonio, TX, USA, January 20-22, 1999. pp. 228–241. ACM (1999), https://doi.org/10.1145/292540.292561
[26] Myers, A.C., Liskov, B.: Protecting privacy using the decentralized label model. ACM Trans. Softw. Eng. Methodol. 9(4), 410–442 (2000), https://doi.org/10.1145/363516.363526
[27] Myers, A.C., Sabelfeld, A., Zdancewic, S.: Enforcing robust declassification. In: 17th IEEE Computer Security Foundations Workshop, (CSFW-17 2004), 28-30 June 2004, Pacific Grove, CA, USA. pp. 172–186. IEEE Computer Society (2004), http://doi.ieeecomputersociety.org/10.1109/CSFW.2004.9
[28] Myers, A.C., Zheng, L., Zdancewic, S., Chong, S., Nystrom, N.: Jif 3.0: Java information flow (July 2006)
[29] Nanevski, A., Banerjee, A., Garg, D.: Verification of information flow and access control policies with dependent types. In: 32nd IEEE Symposium on Security and Privacy, S&P 2011, 22-25 May 2011, Berkeley, California, USA. pp. 165–179. IEEE Computer Society (2011), https://doi.org/10.1109/SP.2011.12
[30] Nordström, B., Petersson, K., Smith, J.M.: Programming in Martin-Löf's Type Theory: An Introduction. Clarendon Press, New York, NY, USA (1990)
[31] Norell, U.: Towards a practical programming language based on dependent type theory, vol. 32. Citeseer (2007)
[32] Rajani, V., Garg, D.: Types for information flow control: Labeling granularity and semantic models. In: 31st IEEE Computer Security Foundations Symposium, CSF 2018, Oxford, United Kingdom, July 9-12, 2018. pp. 233–246. IEEE Computer Society (2018), https://doi.org/10.1109/CSF.2018.00024
[33] Russo, A.: Functional pearl: two can keep a secret, if one of them uses haskell. In: Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015, Vancouver, BC, Canada, September 1-3, 2015. pp. 280–288 (2015), https://doi.org/10.1145/2784731.2784756
[34] Russo, A., Claessen, K., Hughes, J.: A library for light-weight information-flow security in haskell. In: Gill, A. (ed.) Proceedings of the 1st ACM SIGPLAN Symposium on Haskell, Haskell 2008, Victoria, BC, Canada, 25 September 2008. pp. 13–24. ACM (2008), https://doi.org/10.1145/1411286.1411289
[35] Sabelfeld, A., Myers, A.C.: Language-based information-flow security. IEEE Journal on Selected Areas in Communications 21(1), 5–19 (2003), https://doi.org/10.1109/JSAC.2002.806121
[36] Sabelfeld, A., Myers, A.C.: A model for delimited information release. In: Futatsugi, K., Mizoguchi, F., Yonezaki, N. (eds.) Software Security - Theories and

Systems, Second Mext-NSF-JSPS International Symposium, ISSS 2003, Tokyo, Japan, November 4-6, 2003, Revised Papers. Lecture Notes in Computer Science, vol. 3233, pp. 174–191. Springer (2003), https://doi.org/10.1007/978-3-540-37621-7_9

[37] Sabelfeld, A., Sands, D.: Dimensions and principles of declassification. In: 18th IEEE Computer Security Foundations Workshop, (CSFW-18 2005), 20-22 June 2005, Aix-en-Provence, France. pp. 255–269. IEEE Computer Society (2005), https://doi.org/10.1109/CSFW.2005.15

[38] Simonet, V.: Flow Caml in a nutshell. In: Hutton, G. (ed.) Proc. of the first APPSEM-II workshop. Nottingham, United Kingdom (Mar 2003)

[39] Stefan, D., Mazières, D., Mitchell, J.C., Russo, A.: Flexible dynamic information flow control in the presence of exceptions. J. Funct. Program. 27, e5 (2017), https://doi.org/10.1017/S0956796816000241

[40] Stefan, D., Russo, A., Buiras, P., Levy, A., Mitchell, J.C., Mazières, D.: Addressing covert termination and timing channels in concurrent information flow systems. In: Thiemann, P., Findler, R.B. (eds.) ACM SIGPLAN International Conference on Functional Programming, ICFP'12, Copenhagen, Denmark, September 9-15, 2012. pp. 201–214. ACM (2012), https://doi.org/10.1145/2364527.2364557

[41] Stefan, D., Russo, A., Mitchell, J.C., Mazières, D.: Flexible dynamic information flow control in haskell. In: Claessen, K. (ed.) Proceedings of the 4th ACM SIGPLAN Symposium on Haskell, Haskell 2011, Tokyo, Japan, 22 September 2011. pp. 95–106. ACM (2011), https://doi.org/10.1145/2034675.2034688

[42] Swamy, N., Chen, J., Chugh, R.: Enforcing stateful authorization and information flow policies in fine. In: Gordon, A.D. (ed.) Programming Languages and Systems, 19th European Symposium on Programming, ESOP 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6012, pp. 529–549. Springer (2010), https://doi.org/10.1007/978-3-642-11957-6_28

[43] Swamy, N., Chen, J., Fournet, C., Strub, P., Bhargavan, K., Yang, J.: Secure distributed programming with value-dependent types. In: Chakravarty, M.M.T., Hu, Z., Danvy, O. (eds.) Proceeding of the 16th ACM SIGPLAN international conference on Functional Programming, ICFP 2011, Tokyo, Japan, September 19-21, 2011. pp. 266–278. ACM (2011), https://doi.org/10.1145/2034773.2034811

[44] Tsai, T., Russo, A., Hughes, J.: A library for secure multi-threaded information flow in haskell. In: 20th IEEE Computer Security Foundations Symposium, CSF 2007, 6-8 July 2007, Venice, Italy. pp. 187–202. IEEE Computer Society (2007), https://doi.org/10.1109/CSF.2007.6

[45] Vassena, M., Russo, A.: On formalizing information-flow control libraries. In: Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security, PLAS@CCS 2016, Vienna, Austria, October 24, 2016. pp. 15–28 (2016), https://doi.org/10.1145/2993600.2993608

[46] Vassena, M., Russo, A., Buiras, P., Waye, L.: Mac a verified static information-flow control library. Journal of Logical and Algebraic Methods in Programming

95, 148 – 180 (2018), http://www.sciencedirect.com/science/article/pii/S235222081730069X

[47] Vassena, M., Russo, A., Garg, D., Rajani, V., Stefan, D.: From fine- to coarse-grained dynamic information flow control and back. PACMPL 3(POPL), 76:1–76:31 (2019), https://doi.org/10.1145/2694344.2694372

[48] Zdancewic, S., Myers, A.C.: Robust declassification. In: 14th IEEE Computer Security Foundations Workshop (CSFW-14 2001), 11-13 June 2001, Cape Breton, Nova Scotia, Canada. pp. 15–23. IEEE Computer Society (2001), https://doi.org/10.1109/CSFW.2001.930133

[49] Zhang, D., Wang, Y., Suh, G.E., Myers, A.C.: A hardware design language for timing-sensitive information-flow security. In: Özturk, Ö., Ebcioglu, K., Dwarkadas, S. (eds.) Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '15, Istanbul, Turkey, March 14-18, 2015. pp. 503–516. ACM (2015), https://doi.org/10.1145/2694344.2694372

[50] Zheng, L., Myers, A.C.: Dynamic security labels and noninterference (extended abstract). In: Dimitrakos, T., Martinelli, F. (eds.) Formal Aspects in Security and Trust: Second IFIP TC1 WG1.7 Workshop on Formal Aspects in Security and Trust (FAST), an event of the 18th IFIP World Computer Congress, August 22-27, 2004, Toulouse, France. IFIP, vol. 173, pp. 27–40. Springer (2004), https://doi.org/10.1007/0-387-24098-5_3

[51] Zheng, L., Myers, A.C.: Dynamic security labels and static information flow control. Int. J. Inf. Sec. 6(2-3), 67–84 (2007), https://doi.org/10.1007/s10207-007-0019-9

A The Calculus

This section formalizes DepSec as TT_{sec}, a dependently typed call-by-value λ-calculus extended with conditional expressions, references, unit, integer, and boolean values, as well as higher order dependent types and security primitives.

A.1 Syntax

Figure 7 shows the formal syntax of the pure calculus underlying TT_{sec} where meta variables t, c, b, and τ denote terms, constants, binders, and types, respectively. The syntax closely resembles the syntax of TT, the underlying calculus of Idris, but with the addition of a conditional construct and base types `Int`, `Bool`, and `()`. We extend

Terms, t	::=	c	(constant)
	|	x	(variable)
	|	b`.`t	(binding)
	|	$t\ t$	(application)
	|	`if` t `then` t `else` t	(conditional)
	|	τ	(type constructor)
Constants, c	::=	i	(integer literal)
	|	*bool*	(boolean literal)
	|	`()`	(unit literal)
Binders, b	::=	$\lambda x : t$	(abstraction)
	|	$\forall x : t$	(function space)
Types, τ	::=	`Type`	(type of types)
	|	`Int`	(integer type)
	|	`Bool`	(boolean type)
	|	`()`	(unit type)

Fig. 7. Syntax of the core calculus underlying TT_{sec}.

this standard calculus with the security primitives of DepSec. Figure 8 presents the extensions to Figure 7 that forms the formal syntax of TT_{sec}. We introduce the security monad $\texttt{DIO}_v\ t$ as well as type $\texttt{DIO}_\tau\ \ell\ t$ and monadic operators `pure` t, $t \gg\!= t$, and `plug` t. We introduce a labeled value $\texttt{Labeled}_v\ t$, a type $\texttt{Labeled}_\tau\ \ell\ t$, and labeling and unlabeling functions `label` t and `unlabel` t. As an example of a labeled resource we introduce references as values $\texttt{Ref}^n_v$ as well as means for allocating, reading, and writing to references.

A.2 Operational semantics

Definition 2 (Small-step pure semantics). *Let* Term *be the set of terms in* TT_{sec} *and let* $t_1, t_2 \in \mathsf{Term}$*. The relation*

$$t_1 \rightsquigarrow t_2 \subseteq \mathsf{Term} \times \mathsf{Term}$$

Terms, $t, \ell ::= \ldots$

$\mid \texttt{pure}\ t$	(return operator)
$\mid \texttt{DIO}_v\ t$	(computation)
$\mid \texttt{Labeled}_v\ t$	(labeled value)
$\mid \texttt{Ref}_v^{\ell}\ n$	(reference)
$\mid t \gg\!= t$	(bind)
$\mid \texttt{label}\ t$	(labeling)
$\mid \texttt{unlabel}\ t$	(unlabeling)
$\mid \texttt{plug}\ t$	(DIO plugging)
$\mid \texttt{new}^{\ell}\ t$	(new reference)
$\mid \texttt{read}\ t$	(read reference)
$\mid \texttt{write}\ t\ t$	(write reference)

Types, $\tau ::= \ldots$

$\mid \texttt{DIO}_\tau\ \ell\ t$
$\mid \texttt{Labeled}_\tau\ \ell\ t$
$\mid \texttt{Ref}_\tau\ \ell\ t$

Fig. 8. Syntax of TT_{sec}.

denotes the small-step operational semantics of the TT_{sec} calculus. The relation $t_1 \rightsquigarrow t_2$ denotes that t_1 reduces to t_2 in one reduction step according to the inference rules in Figure 9.

We explicitly distinguish pure-term evaluation from top-level monadic-term evaluation. The extended semantics is represented as the relation $c_1 \longrightarrow c_2$ introduced in Definition 4 which extends the pure semantics $\rightsquigarrow$ via Lift.

Definition 3 (Store). *Let* `Label` *be a set of security labels. The function*

$$\Sigma : \texttt{Label} \to \mathit{List}\ \mathsf{Term}$$

denotes a store *compartmentalized into isolated labeled segments, one for each label. We write $\Sigma(\ell)[n]$ to retrieve the nth cell in the ℓ-memory and $\Sigma(\ell)[n] ::= t$ for the store obtained by performing the update $\Sigma(\ell)[n \mapsto t]$.*

Definition 4 (Monadic-term semantics). *Let $\langle \Sigma, t\rangle$ be a* configuration *consisting of a store Σ and a term $t \in \mathsf{Term}$. Let Conf be the set of all such configurations. The relation*

$$c_1 \longrightarrow c_2 \subseteq \mathsf{Conf} \times \mathsf{Conf}$$

denotes the monadic-term evaluation according to the inference rules of Figure 9. $\langle \Sigma, t\rangle \longrightarrow^* \langle \Sigma', t'\rangle$ denotes the reflexive transitive closure of $\longrightarrow$, and we write $\langle \Sigma, t\rangle \Downarrow \langle \Sigma', v\rangle$ if and only if v is a value and $\langle \Sigma, t\rangle \longrightarrow^* \langle \Sigma', v\rangle$.

Note that we consider all non-reducible terms to be values and that constructors $\texttt{Labeled}_v$, $\texttt{DIO}_v$, and $\texttt{Ref}_v$ are not available to the user but only introduced in the semantics to model the run-time value produced by e.g. `label` and `pure`.

Core calculus

$$\textsc{App}_1\ \frac{t_1 \rightsquigarrow t_1'}{t_1\ t_2 \rightsquigarrow t_1'\ t_2} \qquad \textsc{App}_2\ \frac{t \rightsquigarrow t'}{v\ t \rightsquigarrow v\ t'} \qquad \textsc{Beta}\ \frac{}{(\lambda x.t)\ v \rightsquigarrow t[v/x]}$$

$$\textsc{If}_1\ \frac{t_1 \rightsquigarrow t_1'}{\texttt{if}\ t_1\ \texttt{then}\ t_2\ \texttt{else}\ t_3 \rightsquigarrow \texttt{if}\ t_1'\ \texttt{then}\ t_2\ \texttt{else}\ t_3} \qquad \textsc{If}_2\ \frac{}{\texttt{if}\ \texttt{True}\ \texttt{then}\ t_2\ \texttt{else}\ t_3 \rightsquigarrow t_2}$$

$$\textsc{If}_3\ \frac{}{\texttt{if}\ \texttt{False}\ \texttt{then}\ t_2\ \texttt{else}\ t_3 \rightsquigarrow t_3}$$

DepSec (pure)

$$\textsc{Bind}_1\ \frac{}{\texttt{DIO}_v\ t_1 \gg\!= t_2 \rightsquigarrow t_2\ t_1} \qquad \textsc{Pure}_1\ \frac{t \rightsquigarrow t'}{\texttt{pure}\ t \rightsquigarrow \texttt{pure}\ t'} \qquad \textsc{Pure}_2\ \frac{}{\texttt{pure}\ v \rightsquigarrow \texttt{DIO}_v\ v}$$

$$\textsc{Label}_1\ \frac{t \rightsquigarrow t'}{\texttt{label}\ t \rightsquigarrow \texttt{label}\ t'} \qquad \textsc{Label}_2\ \frac{}{\texttt{label}\ v \rightsquigarrow (\texttt{Labeled}_v\ v)} \qquad \textsc{Unlabel}_1\ \frac{t \rightsquigarrow t'}{\texttt{unlabel}\ t \rightsquigarrow \texttt{unlabel}\ t'}$$

$$\textsc{Unlabel}_2\ \frac{}{\texttt{unlabel}\ (\texttt{Labeled}_v\ v) \rightsquigarrow \texttt{pure}\ v} \qquad \textsc{New}_1\ \frac{t \rightsquigarrow t'}{\texttt{new}^\ell\ t \rightsquigarrow \texttt{new}^\ell\ t'}$$

$$\textsc{Write}_1\ \frac{t_1 \rightsquigarrow t_1'}{\texttt{write}\ t_1\ t_2 \rightsquigarrow \texttt{write}\ t_1'\ t_2} \qquad \textsc{Write}_2\ \frac{t_2 \rightsquigarrow t_2'}{\texttt{write}\ v\ t_2 \rightsquigarrow \texttt{write}\ v\ t_2'} \qquad \textsc{Read}_1\ \frac{t \rightsquigarrow t'}{\texttt{read}\ t \rightsquigarrow \texttt{read}\ t'}$$

$$\textsc{DIO}_1\ \frac{t \rightsquigarrow t'}{\texttt{DIO}_\tau\ t\ \tau \rightsquigarrow \texttt{DIO}_\tau\ t'\ \tau} \qquad \textsc{DIO}_2\ \frac{\tau \rightsquigarrow \tau'}{\texttt{DIO}_\tau\ v\ \tau \rightsquigarrow \texttt{DIO}_\tau\ v\ \tau'}$$

$$\textsc{Labeled}_1\ \frac{t \rightsquigarrow t'}{\texttt{Labeled}_\tau\ t\ \tau \rightsquigarrow \texttt{Labeled}_\tau\ t'\ \tau} \qquad \textsc{Labeled}_2\ \frac{\tau \rightsquigarrow \tau'}{\texttt{Labeled}_\tau\ v\ \tau \rightsquigarrow \texttt{Labeled}_\tau\ v\ \tau'}$$

$$\textsc{Ref}_1\ \frac{t \rightsquigarrow t'}{\texttt{Ref}_\tau\ t\ \tau \rightsquigarrow \texttt{Ref}_\tau\ t'\ \tau} \qquad \textsc{Ref}_2\ \frac{\tau \rightsquigarrow \tau'}{\texttt{Ref}_\tau\ v\ \tau \rightsquigarrow \texttt{Ref}_\tau\ v\ \tau'} \qquad \textsc{Forall}_1\ \frac{\tau \rightsquigarrow \tau'}{\forall x : \tau.t \rightsquigarrow \forall x : \tau'.t}$$

$$\textsc{Forall}_2\ \frac{t \rightsquigarrow t'}{\forall x : \tau.t \rightsquigarrow \forall x : \tau.t'}$$

DepSec (monadic)

$$\textsc{Lift}\ \frac{t \rightsquigarrow t'}{\langle \Sigma, t\rangle \longrightarrow \langle \Sigma, t'\rangle} \qquad \textsc{Bind}_2\ \frac{\langle \Sigma, t_1\rangle \longrightarrow \langle \Sigma', t_1'\rangle}{\langle \Sigma, t_1 \gg\!= t_2\rangle \longrightarrow \langle \Sigma', t_1' \gg\!= t_2\rangle}$$

$$\textsc{Plug}\ \frac{\langle \Sigma, t\rangle \Downarrow \langle \Sigma', \mathtt{DIO}_v\ t'\rangle}{\langle \Sigma, \mathtt{plug}\ t\rangle \longrightarrow \langle \Sigma', \mathtt{pure}\ (\mathtt{Labeled}_v\ t')\rangle}$$

$$\textsc{New}_2\ \frac{|\Sigma(\ell)| = n}{\langle \Sigma, \mathtt{new}^{\ell}\ (\mathtt{Labeled}_v\ v)\rangle \longrightarrow \langle \Sigma(\ell)[n] ::= v, \mathtt{pure}\ (\mathtt{Ref}_v^{\ell}\ n)\rangle}$$

$$\textsc{Write}_3\ \frac{}{\langle \Sigma, \mathtt{write}\ (\mathtt{Ref}_v^{\ell}\ n)\ (\mathtt{Labeled}_v\ v)\rangle \longrightarrow \langle \Sigma(\ell)[n] ::= v, \mathtt{pure}\ ()\rangle}$$

$$\textsc{Read}_2\ \frac{}{\langle \Sigma, \mathtt{read}\ (\mathtt{Ref}_v^{\ell}\ n)\rangle \longrightarrow \langle \Sigma, \mathtt{pure}\ (\mathtt{Labeled}_v\ \Sigma(\ell)[n])\rangle}$$

Fig. 9. Operational semantics of TT_{sec}.

A.3 Typing rules

Similar to TT, type checking and the dynamic semantics are defined mutually since evaluation relies on terms to be well-typed, and type checking relies on evaluation as equivalence of terms or types is determined by comparing their normal forms. Compile-time evaluation of TT_{sec} is defined by the pure reductions rules in Figure 9 relative to a context Γ. *Conversion* ($\simeq$) is the smallest equivalence relation closed under reduction, that is, if $\Gamma \vdash x \simeq y$ then x and y reduce to the same normal form.

The type inference rules for TT_{sec} is presented in Figure 10. These rules use the *cumulativity* ($\preceq$) relation defined in Figure 11. In TT, the type of types, `Type`, is parameterized by a universe level (constructing an infinite hierarchy of universes) to prevent Girard's paradox. As universe levels are transparent to the user, this is not relevant for our noninterference proof and we ignore this matter in the following. As for TT, we also conjecture that TT_{sec} respects usual properties such as type preservation and uniqueness of typing at compile-time.

Core calculus

$$\textsc{T-Type}\ \frac{}{\Gamma \vdash \mathtt{Type} : \mathtt{Type}} \qquad \textsc{T-Const}_1\ \frac{}{\Gamma \vdash i : \mathtt{Int}} \qquad \textsc{T-Const}_2\ \frac{}{\Gamma \vdash \mathit{bool} : \mathtt{Bool}} \qquad \textsc{T-Const}_3\ \frac{}{\Gamma \vdash () : ()}$$

$$\textsc{T-Const}_4\ \frac{}{\Gamma \vdash \texttt{Int} : \texttt{Type}} \qquad \textsc{T-Const}_5\ \frac{}{\Gamma \vdash \texttt{Bool} : \texttt{Type}} \qquad \textsc{T-Var}\ \frac{(s : S) \in \Gamma}{\Gamma \vdash s : S}$$

$$\textsc{T-App}\ \frac{\Gamma \vdash f : \forall x : S.T \qquad \Gamma \vdash s : S}{\Gamma \vdash f\ s : T[s/x]} \qquad \textsc{T-Lam}\ \frac{\Gamma; x : S \vdash e : T \qquad \Gamma \vdash \forall x : S.T : \texttt{Type}}{\Gamma \vdash \lambda x : S.e : \forall x : S.T}$$

$$\textsc{T-Forall}\ \frac{\Gamma; x : S \vdash T : \texttt{Type} \qquad \Gamma \vdash S : \texttt{Type}}{\Gamma \vdash \forall x : S.T : \texttt{Type}}$$

$$\textsc{T-IfThenElse}\ \frac{\Gamma \vdash t_1 : \texttt{Bool} \qquad \Gamma; t_1 \equiv \texttt{True} \vdash t_2 : S \qquad \Gamma; t_1 \equiv \texttt{False} \vdash t_3 : S}{\Gamma \vdash \texttt{if}\ t_1\ \texttt{then}\ t_2\ \texttt{else}\ t_3 : S}$$

$$\textsc{T-Conv}\ \frac{\Gamma \vdash x : A \qquad \Gamma \vdash A' : \texttt{Type} \qquad \Gamma \vdash A \preceq A'}{\Gamma \vdash x : A'}$$

DepSec

$$\textsc{T-DIO}\ \frac{\Gamma \vdash \ell : \texttt{Label} \qquad \Gamma \vdash t : \texttt{Type}}{\Gamma \vdash \texttt{DIO}_\tau\ \ell\ t : \texttt{Type}} \qquad \textsc{T-Labeled}\ \frac{\Gamma \vdash \ell : \texttt{Label} \qquad \Gamma \vdash t : \texttt{Type}}{\Gamma \vdash \texttt{Labeled}_\tau\ \ell\ t : \texttt{Type}}$$

$$\textsc{T-Ref}\ \frac{\Gamma \vdash \ell : \texttt{Label} \qquad \Gamma \vdash t : \texttt{Type}}{\Gamma \vdash \texttt{Ref}_\tau\ \ell\ t : \texttt{Type}} \qquad \textsc{T-Label}\ \frac{\Gamma \vdash \ell : \texttt{Label} \qquad \Gamma \vdash s : S}{\Gamma \vdash \texttt{label}\ s : \texttt{Labeled}_\tau\ \ell\ S}$$

$$\textsc{T-Unlabel}\ \frac{\ell_L \sqsubseteq \ell_H \qquad \Gamma \vdash s : \texttt{Labeled}_\tau\ \ell_L\ S}{\Gamma \vdash \texttt{unlabel}\ s : \texttt{DIO}_\tau\ \ell_H\ S} \qquad \textsc{T-Bind}\ \frac{\Gamma \vdash s : \texttt{DIO}_\tau\ \ell\ S \qquad \Gamma \vdash t : S \to \texttt{DIO}_\tau\ \ell\ T}{\Gamma \vdash s \gg\!= t : \texttt{DIO}_\tau\ \ell\ T}$$

$$\textsc{T-Pure}\ \frac{\Gamma \vdash s : S \qquad \Gamma \vdash \ell : \texttt{Label}}{\Gamma \vdash \texttt{pure}\ s : \texttt{DIO}_\tau\ \ell\ S} \qquad \textsc{T-Plug}\ \frac{\ell_L \sqsubseteq \ell_H \qquad \Gamma \vdash s : \texttt{DIO}_\tau\ \ell_H\ S}{\Gamma \vdash \texttt{plug}\ s : \texttt{DIO}_\tau\ \ell_L\ (\texttt{Labeled}_\tau\ \ell_H\ S)}$$

$$\textsc{T-NewRef}\ \frac{\ell_L \sqsubseteq \ell_M \sqsubseteq \ell_H \qquad \Gamma \vdash s : \texttt{Labeled}_\tau\ \ell_M\ S}{\Gamma \vdash \texttt{new}^{\ell_H}\ s : \texttt{DIO}_\tau\ \ell_L\ (\texttt{Ref}_\tau\ \ell_H\ S)}$$

$$\text{T-WriteRef}\quad \frac{\ell_L \sqsubseteq \ell_M \sqsubseteq \ell_H \qquad \Gamma \vdash s : \texttt{Ref}_\tau\ \ell_H\ S \qquad \Gamma \vdash t : \texttt{Labeled}_\tau\ \ell_M\ S}{\Gamma \vdash \texttt{write}\ s\ t : \texttt{DIO}_\tau\ \ell_L\ ()}$$

$$\text{T-ReadRef}\quad \frac{\ell_L \sqsubseteq \ell_H \qquad \Gamma \vdash s : \texttt{Ref}_\tau\ \ell_H\ S}{\Gamma \vdash \texttt{read}\ s : \texttt{DIO}_\tau\ \ell_L\ (\texttt{Labeled}_\tau\ \ell_H\ S)}$$

Fig. 10. Typing rules for TT_{sec}.

$$\text{C-Conv}\quad \frac{\Gamma \vdash S \simeq T}{\Gamma \vdash S \preceq T} \qquad\qquad \text{C-Forall}\quad \frac{\Gamma \vdash S_1 \simeq S_2 \qquad \Gamma; x : S_1 \vdash T_1 \preceq T_2}{\Gamma \vdash \forall x : S_1.T_1 \preceq \forall x : S_2.T_2}$$

Fig. 11. Cumulativity.

A.4 Example: Concatenating strings

This example illustrates the adequacy of the TT_{sec} calculus. The concrete example immitates the `readTwoFiles` function presented in Section 4. It takes two labeled strings as input and returns the concatenated result of the content of these, labeled with the join of the original labels. We assume having a well-typed string concatenation function `++` and a well-defined `join` function for which the following rules hold:

$$\frac{\ell : \texttt{Label} \qquad \ell' : \texttt{Label}}{\ell \sqsubseteq \texttt{join}\ \ell\ \ell'} \qquad\qquad \frac{}{\texttt{join} : \forall x, y : \texttt{Label}.\texttt{Label}}$$

The implementation of a concatenation function for labeled strings TT_{sec} is presented in Figure 12.

```
concat : ∀ℓ,ℓ': Label.∀x:Labeled_τ ℓ String.
                    ∀y:Labeled_τ ℓ' String.DIO_τ (join ℓ ℓ') String
concat = λℓ,ℓ':Label.λx:Labeled_τ ℓ String.λy:Labeled_τ ℓ' String.
         unlabel x >>= (λux:DIO_τ (join ℓ ℓ') String.
         unlabel y >>= (λuy:DIO_τ (join ℓ ℓ') String.
         pure (ux ++ uy)))
```

Fig. 12. Concatenation of secure strings in TT_{sec}.

`concat` is typed through multiple applications of T-Lam, which reduces the problem of

typing `concat` to showing

$$\Gamma; (x : \texttt{Labeled}_\tau\ \ell\ \texttt{String}); (y : \texttt{Labeled}_\tau\ \ell'\ \texttt{String}) \\ \vdash \texttt{unlabel}\ x \ldots \texttt{pure}\ (ux +\!\!+ uy) : \texttt{DIO}_\tau\ (\texttt{join}\ \ell\ \ell')\ \texttt{String}$$

and

$$\Gamma \vdash \forall x : \texttt{Labeled}_\tau\ \ell\ \texttt{String}\ . \\ \forall y : \texttt{Labeled}_\tau\ \ell'\ \texttt{String}\ .\ \texttt{DIO}_\tau\ (\texttt{join}\ \ell\ \ell')\ \texttt{String} : \texttt{Type}.$$

The typing of the actual expression with the return type, $\texttt{DIO}_\tau\ (\texttt{join}\ \ell\ \ell')\ \texttt{String}$, continues by T-Bind. This judgment requires that

$$\Gamma; (\ell, \ell' : \texttt{Label}); (x : \texttt{Labeled}_\tau\ \ell\ \texttt{String}) \vdash \texttt{unlabel}\ x : \texttt{DIO}_\tau\ (\texttt{join}\ \ell\ \ell')\ \texttt{String}$$

and that the rest of the expression in fact has a function type which takes such an input. As this goes by rules already presented we continue with the typing of `unlabel` x. This follows by T-Unlabel which can be used by the assumption on `join` and by T-Var.

Showing that the proposed type is a type follows by T-Forall, T-Var, T-App, and the assumption on `join`.

A.5 Erasure

Definition 5 (Erasure function on terms). *Let ℓ_A be the attacker's security level. The function*

$$\varepsilon_{\ell_A} : \mathit{Term} \to \mathit{Term}$$

denotes the erasure function on terms *where values and primitive types like* `True`*,* `Int`*, etc. is unaffected but otherwise defined by:*

$$\varepsilon_{\ell_A}(\bullet) \triangleq \bullet$$

$$\varepsilon_{\ell_A}(\lambda x.t) \triangleq \lambda x.\varepsilon_{\ell_A}(t)$$

$$\varepsilon_{\ell_A}(t_1\ t_2 : \tau) \triangleq \begin{cases} \bullet & \text{if } \tau = \mathit{DIO}_\tau\ \ell\ \tau' \wedge \ell \not\sqsubseteq \ell_A \\ \varepsilon_{\ell_A}(t_1)\ \varepsilon_{\ell_A}(t_2) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{if}\ t_1\ \texttt{then}\ t_2\ \texttt{else}\ t_3 : \tau) \triangleq \begin{cases} \bullet & \text{if } \tau = \mathit{DIO}_\tau\ \ell\ \tau' \wedge \ell \not\sqsubseteq \ell_A \\ \texttt{if}\ \varepsilon_{\ell_A}(t_1)\ \texttt{then}\ \varepsilon_{\ell_A}(t_2)\ \texttt{else}\ \varepsilon_{\ell_A}(t_3) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{pure}\ t : \mathit{DIO}_\tau\ \ell\ \tau) \triangleq \begin{cases} \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \texttt{pure}\ \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\mathit{DIO}_v\ t : \mathit{DIO}_\tau\ \ell\ \tau) \triangleq \begin{cases} \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \mathit{DIO}_v\ \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(t_1 \gg\!= t_2 : \mathit{DIO}_\tau\ \ell\ \tau) \triangleq \begin{cases} \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \varepsilon_{\ell_A}(t_1) \gg\!= \varepsilon_{\ell_A}(t_2) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{Labeled}_v\ t : \texttt{Labeled}_\tau\ \ell\ \tau) \triangleq \begin{cases} \texttt{Labeled}_v\ \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \texttt{Labeled}_v\ \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{label}\ t : \texttt{Labeled}_\tau\ \ell\ \tau) \triangleq \begin{cases} \texttt{label}\ \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \texttt{label}\ \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{unlabel}\ t : \texttt{DIO}_\tau\ \ell\ \tau) \triangleq \begin{cases} \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \texttt{unlabel}\ \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{Ref}^n_v\ : \texttt{Ref}_\tau\ \ell\ \tau) \triangleq \begin{cases} \texttt{Ref}^\bullet_v & \text{if } \ell \not\sqsubseteq \ell_A \\ \texttt{Ref}^n_v & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{new}^{\ell'}\ t : \texttt{DIO}_\tau\ \ell\ (\texttt{Ref}_\tau\ \ell'\ \tau)) \triangleq \begin{cases} \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \texttt{new}^{\ell'}\ \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{write}\ t_1\ t_2 : \texttt{DIO}_\tau\ \ell\ \tau) \triangleq \begin{cases} \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \texttt{write}\ \varepsilon_{\ell_A}(t_1)\ \varepsilon_{\ell_A}(t_2) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{read}\ t : \texttt{DIO}_\tau\ \ell\ (\texttt{Labeled}_\tau\ \ell'\ \tau)) \triangleq \begin{cases} \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \texttt{read}\ \bullet & \text{if } \ell' \not\sqsubseteq \ell_A \\ \texttt{read}\ \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{plug}\ t : \texttt{DIO}_\tau\ \ell\ (\texttt{Labeled}_\tau\ \ell'\ \tau)) \triangleq \begin{cases} \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \texttt{plug}_\bullet\ \varepsilon_{\ell_A}(t) & \text{if } \ell' \not\sqsubseteq \ell_A \\ \texttt{plug}\ \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{plug}_\bullet\ t : \texttt{DIO}_\tau\ \ell\ (\texttt{Labeled}_\tau\ \ell'\ \tau)) \triangleq \begin{cases} \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \texttt{plug}_\bullet\ \varepsilon_{\ell_A}(t) & \text{otherwise} \end{cases}$$

$$\varepsilon_{\ell_A}(\texttt{DIO}_\tau\ \ell\ \tau) \triangleq \texttt{DIO}_\tau\ \varepsilon_{\ell_A}(\ell)\ \varepsilon_{\ell_A}(\tau)$$

$$\varepsilon_{\ell_A}(\texttt{Labeled}_\tau\ \ell\ \tau) \triangleq \texttt{Labeled}_\tau\ \varepsilon_{\ell_A}(\ell)\ \varepsilon_{\ell_A}(\tau)$$

$$\varepsilon_{\ell_A}(\texttt{Ref}_\tau\ \ell\ \tau) \triangleq \texttt{Ref}_\tau\ \varepsilon_{\ell_A}(\ell)\ \varepsilon_{\ell_A}(\tau)$$

$$\varepsilon_{\ell_A}(\forall x : \tau.t) \triangleq \forall x : \varepsilon_{\ell_A}(\tau).\varepsilon_{\ell_A}(t)$$

In most cases the definition of the erasure function is straightforward as we simply collapse sensitive information and computations to $\bullet$ if they are above the security level of the attacker and otherwise apply the function homomorphically. In one particular case this idea fails, namely the erasure of the term `plug` t.

Consider $\texttt{plug}\ t : \texttt{DIO}_\tau\ \ell\ (\texttt{Labeled}_\tau\ \ell'\ \tau)$ for some ℓ, ℓ', and τ. If the adversary is not allowed to see ℓ, i.e. $\ell \not\sqsubseteq \ell_A$, the computation should not be visible to the adversary and therefore it should be completely collapsed into $\bullet$. Unfortunately, this approach of rewriting entire computations fails if $\ell \sqsubseteq \ell_A$ and $\ell' \not\sqsubseteq \ell_A$ as it would not be possible to show that $\langle \varepsilon_{\ell_A}(\Sigma), \texttt{plug}\ \bullet\rangle \longrightarrow \langle \varepsilon_{\ell_A}(\Sigma'), \texttt{pure}\ \texttt{Labeled}_v\ \bullet\rangle$ as $\langle \Sigma, \bullet\rangle \Downarrow \langle \Sigma', \bullet\rangle$ since $\bullet \rightsquigarrow \bullet$ and it does therefore not have a normal form. Hence, we need a context-sensitive erasure function as the idea about simply erasing computations above the level of an attacker is too simple. To handle this case soundly we make use of *two-steps erasure* that works by introducing an extra semantic step for $\texttt{plug}_\bullet$ introduced by the erasure

function. The extension is presented in Figure 13. Note that this, from an attackers point of view, still looks exactly like one would expect when erasing data and this is therefore purely a technicality of the proof.

$$\textsc{Hole}\ \frac{}{\bullet \rightsquigarrow \bullet} \qquad \textsc{Plug}_\bullet\ \frac{}{\texttt{plug}_\bullet\ t \rightsquigarrow \texttt{pure}\ (\texttt{Labeled}_v\ \bullet)}$$

Fig. 13. Operational semantics of $\mathsf{TT}_{sec\bullet}$: extensions to TT_{sec}.

Extra typing rules for both $\bullet$ and $\texttt{plug}_\bullet$ are also introduced as presented in Figure 14.

$$\text{T-Hole}\ \frac{\Gamma \vdash \tau : \texttt{Type}}{\Gamma \vdash \bullet : \tau} \qquad \text{T-Plug}_\bullet\ \frac{\ell_L \sqsubseteq \ell_H \qquad \Gamma \vdash s : \texttt{DIO}_\tau\ \ell_H\ S}{\Gamma \vdash \texttt{plug}_\bullet\ s : \texttt{DIO}_\tau\ \ell_L\ (\texttt{Labeled}_\tau\ \ell_H\ S)}$$

Fig. 14. Typing rules for $\mathsf{TT}_{sec\bullet}$: extensions to TT_{sec}.

The definition of ε_{ℓ_A} on stores is straightforward as we have a compartmentalized memory. If the a store is erased up to a security level ℓ_A then all levels above this should simply be collapsed entirely.

Definition 6 (Erasure function on configurations). *Let ℓ_A be the attacker's security level. The function*

$$\varepsilon_{\ell_A} : \mathit{Conf} \to \mathit{Conf}_\bullet$$

denotes the erasure function for configurations defined by

$$\varepsilon_{\ell_A}(\langle \Sigma, t : \texttt{DIO}_\tau\ \ell\ \tau \rangle) \triangleq \begin{cases} \langle \varepsilon_{\ell_A}(\Sigma), \bullet \rangle & \text{if } \ell \not\sqsubseteq \ell_A \\ \langle \varepsilon_{\ell_A}(\Sigma), \varepsilon_{\ell_A}(t) \rangle & \text{otherwise} \end{cases}$$

where the store Σ is erased pointwise at each security level and in every cell, i.e. $\varepsilon_{\ell_A}(\Sigma) = \lambda\ell.\varepsilon_{\ell_A}(\Sigma(\ell))$, where

$$\varepsilon_{\ell_A}(\Sigma(\ell)) \triangleq \begin{cases} \bullet & \text{if } \ell \not\sqsubseteq \ell_A \\ \mathit{map}\ \varepsilon_{\ell_A}\ \Sigma(\ell) & \text{otherwise} \end{cases}$$

Note that writing to an erased cell yields no update, i.e. $(\Sigma(\ell)[\bullet] ::= t) \triangleq \Sigma(\ell)$, and reading from an erased compartment yields $\bullet$, i.e. $\bullet[n] \triangleq \bullet$.

Definition 7 (ℓ_A-equivalence). *Let $c_1, c_2 \in \mathit{Conf}$. c_1 and c_2 are said to be* indistinguishable *from security level ℓ_A, written $c_1 \approx_{\ell_A} c_2$, if and only if $\varepsilon_{\ell_A}(c_1)$ and $\varepsilon_{\ell_A}(c_2)$ are structurally equivalent, written $\varepsilon_{\ell_A}(c_1) \equiv \varepsilon_{\ell_A}(c_2)$.*

B Results

Lemma 1 (Erasure of substitution). *Let* $t, v \in \mathsf{Term}$. *Then* $\varepsilon_{\ell_A}(t[v/x]) \equiv \varepsilon_{\ell_A}(t)[\varepsilon_{\ell_A}(v)/x]$.

Proof. The statement follows by case splitting on t and v and the definition of ε_{ℓ_A}.

Lemma 2 (Distributivity on pure term reduction). *Let* $t_1, t_2 \in \mathsf{Term}$. *If* $t_1 \rightsquigarrow t_2$ *then* $\varepsilon_{\ell_A}(t_1) \rightsquigarrow \varepsilon_{\ell_A}(t_2)$.

Proof. The proof goes by structural induction in the derivation of $t_1 \rightsquigarrow t_2$.

- App$_1$: Assume $t_1\ t_2 \rightsquigarrow t_1'\ t_2$. If $t_1\ t_2$ has type $\texttt{DIO}_\tau\ \ell\ \tau'$ and $\ell \not\sqsubseteq \ell_A$, the statement follows from the definition of ε_{ℓ_A} and Hole. Otherwise, $t_1 \rightsquigarrow t_1'$ holds by App$_1$ and by the induction hypothesis $\varepsilon_{\ell_A}(t_1) \rightsquigarrow \varepsilon_{\ell_A}(t_1')$. By App$_1$ and definition of ε_{ℓ_A} it holds that $\varepsilon_{\ell_A}(t_1\ t_2) \rightsquigarrow \varepsilon_{\ell_A}(t_1'\ t_2)$.
- App$_2$: The argument is identical to the App$_1$ case.
- Beta: Assume $(\lambda x.t)\ v \rightsquigarrow t[v/x]$. By Beta, $\lambda x.\varepsilon_{\ell_A}(t)\ \varepsilon_{\ell_A}(v) \rightsquigarrow \varepsilon_{\ell_A}(t)[\varepsilon_{\ell_A}(v)/x]$. By definition of ε_{ℓ_A} and Lemma 1 then $\varepsilon_{\ell_A}((\lambda x.t)\ v) \rightsquigarrow \varepsilon_{\ell_A}(t[v/x])$.
- If$_1$: The argument is identical to the App$_1$ case.
- If$_2$ and If$_3$: If t_1 and t_2 have type $\texttt{DIO}_\tau\ \ell\ \tau'$ and $\ell \not\sqsubseteq \ell_A$, the statement follows from the definition of ε_{ℓ_A} and Hole. Otherwise, the statement follows directly by the definition ε_{ℓ_A} and If$_i$.
- Bind$_1$: Assume $\texttt{pure}\ t_1 \gg\!= t_2 \rightsquigarrow t_2\ t_1$. As we assume well-typed terms, $\texttt{pure}\ t_1 \gg\!= t_2$ has type $\texttt{DIO}_\tau\ \ell\ \tau$ for some ℓ and τ. If $\ell \not\sqsubseteq \ell_A$ the statement follows from the definition of ε_{ℓ_A} and Hole. Otherwise, the statement follows from Bind-Pure and the definition of ε_{ℓ_A}.
- Pure$_1$: Assume $\texttt{pure}\ t \rightsquigarrow \texttt{pure}\ t'$. As we assume well-typed terms, $\texttt{pure}\ t$ and $\texttt{pure}\ t'$ have type $\texttt{DIO}_\tau\ \ell\ \tau$ for some ℓ and τ. If $\ell \sqsubseteq \ell_A$ then the statement follows from the induction hypothesis, the definition of ε_{ℓ_A} and Pure$_1$. If $\ell \not\sqsubseteq \ell_A$ then the statement follows from the definition ε_{ℓ_A} and Hole.
- Pure$_2$: Assume $\texttt{pure}\ v \rightsquigarrow \texttt{DIO}_v\ v$. As we assume well-typed terms, $\texttt{pure}\ v$ and $\texttt{DIO}_v\ v$ have type $\texttt{DIO}_\tau\ \ell\ \tau$ for some ℓ and τ. If $\ell \sqsubseteq \ell_A$ then the statement follows from the definition of ε_{ℓ_A} and Pure$_2$. If $\ell \not\sqsubseteq \ell_A$ then the statement follows from the definition ε_{ℓ_A} and Hole.
- Label$_1$: Assume $\texttt{label}\ t \rightsquigarrow \texttt{label}\ t'$. As we assume well-typed terms, $\texttt{label}\ t$ and $\texttt{label}\ t'$ have type $\texttt{Labeled}_\tau\ \ell\ \tau$ for some ℓ and τ. If $\ell \sqsubseteq \ell_A$ then the statement follows from the induction hypothesis, the definition of ε_{ℓ_A}, and Label$_1$. If $\ell \not\sqsubseteq \ell_A$ then the statement follows from the definition ε_{ℓ_A} Hole and Label$_1$.
- Label$_2$: Assume $\texttt{label}\ t \rightsquigarrow \texttt{Labeled}_v\ t$. As we assume well-typed terms, $\texttt{label}\ t$ and $\texttt{Labeled}_v\ t$ have type $\texttt{Labeled}_\tau\ \ell\ \tau$ for some ℓ, ℓ', and τ. If $\ell \not\sqsubseteq \ell_A$ the statement follows from the definition of ε_{ℓ_A} Hole and Label$_1$. Otherwise the statement follows by Label$_2$, the definition of ε_{ℓ_A} and the induction hypothesis.
- Unlabel$_1$: Assume $\texttt{unlabel}\ t \rightsquigarrow \texttt{unlabel}\ t'$. As we assume well-typed terms, $\texttt{unlabel}\ t$ and $\texttt{unlabel}\ t'$ have type $\texttt{DIO}_\tau\ \ell\ \tau$ for some ℓ and τ. If $\ell \not\sqsubseteq$

ℓ_A the statement follows from the definition of ε_{ℓ_A} and Hole. Otherwise, the statement follows from the induction hypothesis, Unlabel$_1$, and the definition of ε_{ℓ_A}.

Unlabel$_2$: Assume $\texttt{unlabel}\ (\texttt{Labeled}_v\ t) \rightsquigarrow \texttt{pure}\ t$. As we assume well-typed terms, $\texttt{unlabel}\ (\texttt{Labeled}_v\ t)$ and $\texttt{pure}\ t$ have type $\texttt{DIO}_\tau\ \ell\ \tau$ where $\texttt{Labeled}_v\ t$ have type $\texttt{Labeled}_\tau\ \ell'\ \tau$ such that $\ell' \sqsubseteq \ell$. If $\ell \not\sqsubseteq \ell_A$ the statement follows from the definition of ε_{ℓ_A} and Hole. If $\ell \sqsubseteq \ell_A$ then by transitivity of the partial ordering $\ell' \sqsubseteq \ell_A$ holds and the statement then follows by Unlabel$_2$ and the definition of ε_{ℓ_A}.

New$_1$: Assume $\texttt{new}^\ell\ t \rightsquigarrow \texttt{new}^\ell\ t'$. As we assume well-typed terms, $\texttt{new}^\ell\ t$ and $\texttt{new}^\ell\ t'$ have type $\texttt{DIO}_\tau\ \ell\ \tau$ for some ℓ and τ. If $\ell \not\sqsubseteq \ell_A$ then the statements follows from the definition of ε_{ℓ_A} and Hole. If $\ell \sqsubseteq \ell_A$ the statement follows from the induction hypothesis, the definition of ε_{ℓ_A}, and New$_1$.

Write$_1$: Assume $\texttt{write}\ t_1\ t_2 \rightsquigarrow \texttt{write}\ t_1'\ t_2$. As we assume well-typed terms, $\texttt{write}\ t_1\ t_2$ and $\texttt{write}\ t_1'\ t_2$ have type $\texttt{DIO}_\tau\ \ell\ ()$ and t_2 has type $\texttt{Labeled}_\tau\ \ell'\ \tau$. If $\ell \not\sqsubseteq \ell_A$ the statement follows from the definition of ε_{ℓ_A} and Hole. If $\ell \sqsubseteq \ell_A$ then the statement follows by the induction hypothesis, Write$_1$, and the definition of ε_{ℓ_A}.

Write$_2$: The argument is identical to the Write$_1$ case.

Read$_1$: Assume $\texttt{read}\ t \rightsquigarrow \texttt{read}\ t'$. As we assume well-typed terms $\texttt{read}\ t$ and $\texttt{read}\ t'$ have type $\texttt{DIO}_\tau\ \ell'\ (\texttt{Labeled}_\tau\ \ell\ \tau)$. If $\ell' \not\sqsubseteq \ell_A$ the statement follows from the definition of ε_{ℓ_A} and Hole. Otherwise, if $\ell' \sqsubseteq \ell_A$ then consider whether $\ell \sqsubseteq \ell_A$ holds. If $\ell \not\sqsubseteq \ell_A$ the statement follows by Read$_1$ and Hole. If $\ell \sqsubseteq \ell_A$ then the statement follows by the induction hypothesis, Read$_1$, and the definition of ε_{ℓ_A}.

DIO$_i$, Labeled$_i$, Ref$_i$, Forall$_i$: In all cases, the statement follows directly from the induction hypothesis, the definition of ε_{ℓ_A}, and the inference rules.

Plug$_\bullet$: Assume $\texttt{plug}_\bullet\ t \rightsquigarrow \texttt{pure}\ (\texttt{Labeled}_v\ \bullet)$. As we assume well-typed terms, $\texttt{plug}_\bullet\ t$ has type $\texttt{DIO}_\tau\ \ell\ (\texttt{Labeled}_\tau\ \ell'\ \tau)$. Consider whether $\ell' \sqsubseteq \ell_A$. In both cases, the statement follows by the definition ε_{ℓ_A} and Plug$_\bullet$.

Hole: The statement follows directly from the definition of ε_{ℓ_A} and Hole.

Lemma 3 (Erasure of a computation). *Let $t \in \mathit{Term}$. If t has type $\texttt{DIO}_\tau\ \ell\ \tau$ and $\ell \not\sqsubseteq \ell_A$ then $\varepsilon_{\ell_A}(t) \equiv \bullet$.*

Proof. The statement follows directly by case splitting on t and the definition of ε_{ℓ_A}.

Lemma 4 (Single step erased store equivalence). *Let $c_1 = \langle \Sigma_1, t_1 \rangle$ and $c_2 = \langle \Sigma_2, t_2 \rangle$. If t_1 and t_2 have type $\texttt{DIO}_\tau\ \ell\ \tau$, $\ell \not\sqsubseteq \ell_A$, and $c_1 \longrightarrow c_2$ then $\varepsilon_{\ell_A}(\Sigma_1) \equiv \varepsilon_{\ell_A}(\Sigma_2)$.*

Proof. The proof goes by case splitting in the derivation of $c_1 \longrightarrow c_2$.

Plug: Assume $\langle \Sigma, \texttt{plug}\ t \rangle \longrightarrow \langle \Sigma', \texttt{pure}\ (\texttt{Labeled}_v\ t') \rangle$. As we assume well-typed terms, t has type $\texttt{DIO}_\tau\ \ell'\ \tau'$ for some ℓ' and τ' where $\ell \sqsubseteq \ell'$. By transitivity of $\sqsubseteq$ it follows that $\ell' \not\sqsubseteq \ell_A$ and then the statement follows by Lemma 5 as t is structurally smaller than $\texttt{plug}\ t$.

New$_2$: Assume $\langle \Sigma, \texttt{new}^{\ell'}\ (\texttt{Labeled}_v\ t)\rangle \longrightarrow \langle \Sigma(\ell')[n] ::= t, \texttt{pure}\ (\texttt{Ref}_v^{\ell'}\ n)\rangle$. As we assume well-typed terms, $\texttt{new}^{\ell'}\ (\texttt{Labeled}_v\ t)$ and $\texttt{pure}\ (\texttt{Ref}_v^{\ell'}\ n)$ have type $\texttt{DIO}_\tau\ \ell\ (\texttt{Ref}_\tau\ \ell'\ \tau)$ where $\ell \sqsubseteq \ell'$. By transitivity of $\sqsubseteq$ it follows that $\ell' \not\sqsubseteq \ell_A$. Note that the only memory compartment changed is the one of ℓ' and as writing to an erased cell makes no update the statement follows.

Write$_3$: The argument is identical to the New case.

The remaining cases does not change the store and the statement follows immediately.

Lemma 5 (Multi-step erased store equivalence). *Let $c_1 = \langle \Sigma_1, t_1\rangle$ and $c_2 = \langle \Sigma_2, t_2\rangle$. If t_1 and t_2 have type $\texttt{DIO}_\tau\ \ell\ \tau$, $\ell \not\sqsubseteq \ell_A$, and $c_1 \longrightarrow^* c_2$ then $\varepsilon_{\ell_A}(\Sigma_1) \equiv \varepsilon_{\ell_A}(\Sigma_2)$.*

Proof. The statement follows from repeated applications of Lemma 4.

Lemma 6. *Let Σ be a store, $n \in \mathbb{N} \cup \{\bullet\}$, and $t \in \mathit{Term}$. Then $\varepsilon_{\ell_A}(\Sigma(\ell)[n] ::= t) \equiv \varepsilon_{\ell_A}(\Sigma)(\ell)[n] ::= \varepsilon_{\ell_A}(t)$.*

Proof. Note that the erasure of a store erases each compartment in both the old and the updated store. Only the $\Sigma(\ell)$ compartment is changed, hence the remaining compartments are preserved by definition. It remains to show that the updated compartments are equivalent. Let the updated store $\Sigma(\ell)[n] ::= t$ be denoted by Σ'. If $\ell \sqsubseteq \ell_A$ then $\varepsilon_{\ell_A}(\Sigma'(\ell)) = \mathit{map}\ \varepsilon_{\ell_A}\ \Sigma'(\ell)$ cf. the definition of ε_{ℓ_A} and hence the statement follows from properties of *map*. If $\ell \not\sqsubseteq \ell_A$ then $\varepsilon_{\ell_A}(\Sigma'(\ell)) \equiv \bullet$ and the statement follows as updating an erased cell yields no update.

Lemma 7. *Let Σ be a store and $n \in \mathbb{N}$. If $\ell \sqsubseteq \ell_A$ then $\varepsilon_{\ell_A}(\Sigma(\ell)[n]) \equiv \varepsilon_{\ell_A}(\Sigma)(\ell)[n]$.*

Proof. The statements follows from the definition of ε_{ℓ_A} and properties of *map*.

Proposition 2 (Distributivity of ε_{ℓ_A} over $\longrightarrow$). *Let $c_1, c_2 \in \mathit{Conf}$. If $c_1 \longrightarrow c_2$ then $\varepsilon_{\ell_A}(c_1) \longrightarrow \varepsilon_{\ell_A}(c_2)$.*

Proof. Let $c_1 = \langle \Sigma, t\rangle$ and $c_2 = \langle \Sigma', t'\rangle$. The proof goes by structural induction in the derivation of $c_1 \longrightarrow c_2$.

Lift: Note $\Sigma = \Sigma'$. Necessarily, $t \rightsquigarrow t'$ must hold. By Lemma 2 it holds $\varepsilon_{\ell_A}(t) \rightsquigarrow \varepsilon_{\ell_A}(t')$ and by Lift and the definition of ε_{ℓ_A} on configurations the statement follows.

Note that in the remaining cases t and t' necessarily have type $\texttt{DIO}_\tau\ \ell\ \tau$ as we assume well-typed terms. If $\ell \not\sqsubseteq \ell_A$ the statement in all cases follows by the definition of ε_{ℓ_A}, Lemma 3, Lemma 4, Lift, and Hole. Now, assume $\ell \sqsubseteq \ell_A$.

Bind$_2$: Assume $\langle \Sigma, t_1 \gg\!= t_2\rangle \longrightarrow \langle \Sigma', t_1' \gg\!= t_2\rangle$. The statement follows by Bind$_2$, the induction hypothesis and the definition of ε_{ℓ_A}.

PLUG: Assume $\langle \Sigma, \texttt{plug}\ t\rangle \longrightarrow \langle \Sigma', \texttt{pure}\ (\texttt{Labeled}_v\ t')\rangle$. As we assume well-typed terms, t has type $\texttt{DIO}_\tau\ \ell'\ \tau'$ for some ℓ' and τ'. From PLUG it holds that $\langle \Sigma, t\rangle \Downarrow \langle \Sigma', \texttt{DIO}_v\ t'\rangle$. If $\ell' \sqsubseteq \ell_A$, cf. Lemma 9 and that t is structurally smaller than $\texttt{plug}\ t$, it follows $\varepsilon_{\ell_A}(\langle \Sigma, t\rangle) \Downarrow \varepsilon_{\ell_A}(\langle \Sigma', \texttt{DIO}_v\ t'\rangle)$, and from the definition of ε_{ℓ_A} and PLUG the statement holds. If $\ell' \not\sqsubseteq \ell_A$, it follows from Lemma 5 that $\varepsilon_{\ell_A}(\Sigma) \equiv \varepsilon_{\ell_A}(\Sigma')$, and using LIFT, the definition of ε_{ℓ_A} and PLUG$_\bullet$ the statement follows.

NEW$_2$: Assume $\langle \Sigma, \texttt{new}^{\ell'''}\ (\texttt{Labeled}_v\ t)\rangle \longrightarrow \langle \Sigma(\ell'')[n] ::= t, \texttt{pure}\ (\texttt{Ref}_v^{\ell'''}\ n)\rangle$. As we assume well-typed terms, terms $\texttt{new}^{\ell'''}\ (\texttt{Labeled}_v\ t)$ and $\texttt{pure}\ \texttt{Ref}_v^{\ell'''}\ n$ have type $\texttt{DIO}_\tau\ \ell\ (\texttt{Ref}_\tau\ \ell''\ \tau)$, and the term $\texttt{Labeled}_v\ t$ has type $\texttt{Labeled}_\tau\ \ell'\ \tau$ where $\ell \sqsubseteq \ell' \sqsubseteq \ell''$. If $\ell'' \sqsubseteq \ell_A$ then by transitivity $\ell' \sqsubseteq \ell_A$. From the definition of ε_{ℓ_A} it follows $\varepsilon_{\ell_A}(\Sigma(\ell'')) = \mathit{map}\ \varepsilon_{\ell_A}\ \Sigma(\ell'')$ and hence $|\varepsilon_{\ell_A}(\Sigma(\ell''))| = |\Sigma(\ell'')| = n$. From Lemma 6, the definition of ε_{ℓ_A}, and NEW$_2$ the statement follows. If $\ell'' \not\sqsubseteq \ell_A$ then then from the definition of ε_{ℓ_A} it follows $\varepsilon_{\ell_A}(\Sigma(\ell'')) \equiv \bullet$ and the size of an erased label segment is also erased, hence $|\varepsilon_{\ell_A}(\Sigma(\ell''))| = |\bullet| = \bullet$. Consider whether $\ell' \sqsubseteq \ell_A$. In both cases the statement follows from Lemma 6, the definition of ε_{ℓ_A}, and NEW$_2$.

WRITE$_3$: Assume

$$\langle \Sigma, \texttt{write}\ (\texttt{Ref}_v^{\ell'}\ n)\ (\texttt{Labeled}_v\ t)\rangle \longrightarrow \langle \Sigma(\ell')[n] ::= t, \texttt{pure}\ ()\rangle.$$

As we assume well-typed terms, $\texttt{Ref}_v^{\ell'}\ n$ has type $\texttt{Ref}_\tau\ \ell'\ \tau$ and $\texttt{Labeled}_v\ t$ has type $\texttt{Labeled}_\tau\ \ell''\ \tau$ for some ℓ', ℓ'', and τ where $\ell'' \sqsubseteq \ell'$. If $\ell' \sqsubseteq \ell_A$ then by transitivity $\ell'' \sqsubseteq \ell_A$ and hence

$$\begin{aligned}&\texttt{write}\ (\texttt{Ref}_v^{\ell'}\ n)\ (\texttt{Labeled}_v\ \varepsilon_{\ell_A}(t))\\ \equiv\ &\texttt{write}\ \varepsilon_{\ell_A}(\texttt{Ref}_v^{\ell'}\ n)\ \varepsilon_{\ell_A}(\texttt{Labeled}_v\ t)\end{aligned}$$

by definition of ε_{ℓ_A}. The statement now follows from Lemma 6, WRITE$_3$, and the definition of ε_{ℓ_A}. If $\ell' \not\sqsubseteq \ell_A$ then consider whether $\ell'' \sqsubseteq \ell_A$. In either case, the statement follows from Lemma 6, WRITE$_3$, and the definition of ε_{ℓ_A}.

READ$_2$: Assume $\langle \Sigma, \texttt{read}\ (\texttt{Ref}_v^{\ell}\ n)\rangle \longrightarrow \langle \Sigma, \texttt{pure}\ (\texttt{Labeled}_v\ \Sigma(\ell)[n])\rangle$. As we assume well-typed terms, $\texttt{Ref}_v^{\ell}\ n$ has type $\texttt{Ref}_\tau\ \ell\ \tau$ and $\texttt{pure}\ (\texttt{Labeled}_v\ \Sigma(\ell)[n])$ has type $\texttt{DIO}_\tau\ \ell'\ (\texttt{Labeled}_\tau\ \ell\ \tau)$ for some ℓ, ℓ' and τ where $\ell' \sqsubseteq \ell$. If $\ell \sqsubseteq \ell_A$ then by transitivity $\ell' \sqsubseteq \ell_A$ and the statement follows from Lemma 7, READ$_2$, and the definition of ε_{ℓ_A}. If $\ell \not\sqsubseteq \ell_A$ the statement follows by the fact that reading from an erased compartment yields $\bullet$, the definition of ε_{ℓ_A}, and READ$_2$.

Lemma 8 (Distributivity of ε_{ℓ_A} over $\longrightarrow^*$). *Let $c_1, c_2 \in \mathit{Conf}$. If $c_1 \longrightarrow^* c_2$ then $\varepsilon_{\ell_A}(c_1) \longrightarrow^* \varepsilon_{\ell_A}(c_2)$.*

Proof. The statement follows from repeated applications of Proposition 2.

Lemma 9 (Distributivity of ε_{ℓ_A} over $\Downarrow$). *Let $c_1, c_2 \in$ Conf. If $c_1 \Downarrow c_2$ then $\varepsilon_{\ell_A}(c_1) \Downarrow \varepsilon_{\ell_A}(c_2)$.*

Proof. The statement follows directly from Lemma 8.

Lemma 10 (Determinacy of pure reductions). *Let $t_1, t_2, t_3 \in$ Term. If $t_1 \rightsquigarrow t_2$ and $t_1 \rightsquigarrow t_3$ then $t_2 \equiv t_3$.*

Proof. The proof goes by structural induction in the derivation of $t_1 \rightsquigarrow t_2$ and $t_1 \rightsquigarrow t_3$.

Proposition 3 (Single step determinacy). *Let $c_1, c_2, c_3 \in$ Conf. If $c_1 \longrightarrow c_2$ and $c_1 \longrightarrow c_3$ then $c_2 \equiv c_3$.*

Proof. The proof goes by structural induction in the derivation of $c_1 \longrightarrow c_2$ and $c_1 \longrightarrow c_3$. Note that both $c_1 \longrightarrow c_2$ and $c_1 \longrightarrow c_3$ have to have been derived from the same inference rule, syntactically decidable from c_1.

Lift: The statement follows by Lemma 10.
Plug: Assume $\langle \Sigma_1, \texttt{plug}\ t_1 \rangle \longrightarrow \langle \Sigma_2, \texttt{pure}\ (\texttt{Labeled}_v\ t_2) \rangle$ and $\langle \Sigma_1, \texttt{plug}\ t_1 \rangle \longrightarrow \langle \Sigma_3, \texttt{pure}\ (\texttt{Labeled}_v\ t_3) \rangle$. The statement follows from Lemma 11 as t_1 is structurally smaller than $\texttt{plug}\ t_1$.

In the remaining cases the statement follows by standard structural induction.

Lemma 11 (Big step determinacy). *Let $c_1, c_2, c_3 \in$ Conf. If $c_1 \Downarrow c_2$ and $c_1 \Downarrow c_3$ then $c_2 \equiv c_3$.*

Proof. The statement follows from repeated applications of Proposition 3.

Proposition 4 (Single step $\approx_{\ell_A}$ preservation). *Let $c_1, c_1', c_2, c_2' \in$ Conf. If $c_1 \approx_{\ell_A} c_2$, $c_1 \longrightarrow c_1'$, and $c_2 \longrightarrow c_2'$ then $c_1' \approx_{\ell_A} c_2'$.*

Proof. Proposition 2 states that $\varepsilon_{\ell_A}(c_1) \longrightarrow \varepsilon_{\ell_A}(c_1')$ and $\varepsilon_{\ell_A}(c_2) \longrightarrow \varepsilon_{\ell_A}(c_2')$. From Definition 7 it is known that $\varepsilon_{\ell_A}(c_1) \equiv \varepsilon_{\ell_A}(c_2)$ and from Proposition 3 it follows that $\varepsilon_{\ell_A}(c_1') \equiv \varepsilon_{\ell_A}(c_2')$. Hence $c_1' \approx_{\ell_A} c_2'$.

Theorem 2 (Progress-insensitive noninterference). *Let $c_1, c_1', c_2, c_2' \in$ Conf. If $c_1 \approx_{\ell_A} c_2$, $c_1 \Downarrow c_1'$, and $c_2 \Downarrow c_2'$ then $c_1' \approx_{\ell_A} c_2'$.*

Proof. The statement follows from repeated application of Proposition 4.

C DepSec

DIO.idr

```
module DepSec.DIO

% access public export

||| Security Monad
||| @ l security label of wrapped value
```

```
||| @ valueType type of wrapped value
data DIO : l
        -> (valueType : Type)
        -> Type where
  ||| TCB
  MkDIO : IO valueType -> DIO l valueType

||| Executes secure computation
||| TCB
||| @ dio secure computation
run : (dio : DIO l a) -> IO a
run (MkDIO m) = m

||| Lifts arbitrary IO monad into security monad
||| TCB
||| @ io computation
lift : (io : IO a) -> DIO l a
lift = MkDIO

Functor (DIO l) where
  map f (MkDIO io) = MkDIO (map f io)

Applicative (DIO l) where
  pure = MkDIO . pure
  (<*>) (MkDIO f) (MkDIO a) = MkDIO (f <*> a)

Monad (DIO l) where
  (>>=) (MkDIO a) f = MkDIO (a >>= run . f)
```

Labeled.idr

```
module DepSec.Labeled

import public DepSec.DIO
import public DepSec.Poset

% access public export

||| Labeled value
||| @ label label
||| @ valueType type of labeled value
data Labeled : (label : labelType)
            -> (valueType : Type)
            -> Type where
  ||| TCB
  MkLabeled : valueType -> Labeled label valueType

||| Label values
||| @ value value to label
label : Poset labelType
      => {l : labelType}
```

```
      -> (value : a)
      -> Labeled l a
label = MkLabeled

||| Unlabel values
||| @ flow evidence that l may flow to l'
||| @ labeled labeled value to unlabel
unlabel : Poset labelType
       => {l,l' : labelType}
       -> {auto flow : l `leq` l'}
       -> (labeled : Labeled l a)
       -> DIO l' a
unlabel (MkLabeled val) = pure val

||| Upgrade the security level of a labeled value
||| @ flow evidence that l may flow to l'
||| @ labeled labeled value to relabel
relabel : Poset labelType
       => {l, l' : labelType}
       -> {auto flow : l `leq` l'}
       -> (labeled : Labeled l a)
       -> Labeled l' a
relabel (MkLabeled x) = MkLabeled x

unlabel' : Poset labelType
        => {l,l' : labelType}
        -> {auto flow : l `leq` l'}
        -> (labeled : Labeled l a)
        -> DIO l' (c : a ** label c = labeled)
unlabel' (MkLabeled x) = pure (x ** Refl)

||| Plug a secure computation into a less secure computation
||| @ flow evidence that l may flow to l'
||| @ dio secure computation to plug into insecure computation
plug : Poset labelType
    => {l,l' : labelType}
    -> (dio : DIO l' a)
    -> {auto flow : l `leq` l'}
    -> DIO l (Labeled l' a)
plug dio = lift . run $ dio >>= pure . MkLabeled
```

Poset.idr

```
module DepSec.Poset

% access public export
% default total
% hide Prelude.Monad.join

||| Verified partial ordering
interface Poset a where
```

```
  leq : a -> a -> Type
  reflexive :  (x : a) -> x `leq` x
  antisymmetric : (x, y : a) -> x `leq` y -> y `leq` x -> x = y
  transitive : (x, y, z : a) -> x `leq` y -> y `leq` z -> x `leq` z
```

Lattice.idr

```
module DepSec.Lattice

import public DepSec.Poset

% access public export
% hide Prelude.Monad.join
% default total

||| Verified join semilattice
interface JoinSemilattice a where
  join : a -> a -> a
  joinAssociative : (x, y, z : a)
                 -> x `join` (y `join` z) = (x `join` y) `join` z
  joinCommutative : (x, y : a)    -> x `join` y = y `join` x
  joinIdempotent  : (x : a)       -> x `join` x = x

||| A well defined join induces a partial ordering.
implementation JoinSemilattice a => Poset a where
  leq x y = (x `join` y = y)
  reflexive = joinIdempotent
  antisymmetric x y lexy leyx =
    rewrite sym $ lexy in
    rewrite joinCommutative x y in
    rewrite sym $ leyx in Refl
  transitive x y z lexy leyx =
    rewrite sym $ leyx in
    rewrite joinAssociative x y z in
    rewrite sym $ lexy in Refl

||| A join-semilattice with an identity element (the lattices' bottom)
||| of the join operation.
interface JoinSemilattice a => BoundedJoinSemilattice a where
  Bottom : a
  bottomUnitaryElement : (e : a) -> e `join` Bottom = e
```

Ref.idr

```
module DepSec.Ref

import public DepSec.Labeled
import public DepSec.DIO
import Data.IORef

% access export
```

```
8
9   ||| Data type for secure references.
10  data SecRef : (l : labelType) -> (valueType : Type) -> Type where
11    |||TCB
12    MkSecRef : (ref : IORef a) -> SecRef l a
13
14  ||| Creating a reference to a labeled value.
15  ||| @ flow evidence that l may flow to l'
16  ||| @ flow' evidence that l' may flow to l''
17  ||| @ value The initial value for the reference.
18  newRef : Poset labelType
19        => {l, l', l'' : labelType}
20        -> {auto flow : l `leq` l'}
21        -> {auto flow' : l' `leq` l''}
22        -> (value : Labeled l' a)
23        -> DIO l (SecRef l'' a)
24  newRef (MkLabeled v)
25    = lift $ newIORef v >>= pure . MkSecRef
26
27  ||| Reading a secure reference.
28  ||| @ flow evidence that l may flow to l'
29  ||| @ ref The reference which we wish to read.
30  readRef : Poset labelType
31         => {l, l' : labelType}
32         -> {auto flow : l `leq` l'}
33         -> (ref  : SecRef l' a)
34         -> DIO l (Labeled l' a)
35  readRef (MkSecRef ioRef)
36    = lift $ map MkLabeled $ readIORef ioRef
37
38  ||| Wrting a labeled value to a secure reference.
39  ||| @ flow evidence that l may flow to l'
40  ||| @ flow' evidence that l' may flow to l''
41  ||| @ ref The reference which we wish to write to.
42  ||| @ content The content which we wish too read.
43  writeRef : Poset labelType
44          => {l, l', l'' : labelType}
45          -> {auto flow : l `leq` l'}
46          -> {auto flow' : l' `leq` l''}
47          -> (ref : SecRef l'' a)
48          -> (content : Labeled l' a)
49          -> DIO l ()
50  writeRef (MkSecRef ioRef) (MkLabeled content)
51    = lift $ writeIORef ioRef content
```

File.idr

```
1   module DepSec.File
2
3   import public DepSec.DIO
4   import public DepSec.Labeled
```

```
5
6   %access export
7
8   ||| Secure file
9   data SecFile : {label : Type} -> (l : label) -> Type where
10    ||| TCB
11    MkSecFile : (path : String) -> SecFile l
12
13  ||| Make a secure file from string
14  ||| TCB
15  ||| @ path path to file
16  makeFile : (path : String) -> SecFile l
17  makeFile = MkSecFile
18
19  ||| Read a secure file
20  ||| @ flow evidence that l may flow to l'
21  ||| @ file secure file to read
22  readFile : Poset labelType
23          => {l,l' : labelType}
24          -> {auto flow : l `leq` l'}
25          -> (file : SecFile l')
26          -> DIO l (Labeled l' (Either FileError String))
27  readFile (MkSecFile path) = lift $ map MkLabeled $ readFile path
28
29  ||| Write to a secure file
30  ||| @ file secure file to write to
31  ||| @ flow evidence that l may flow to l'
32  ||| @ flow' evidence that l' may flow to l''
33  ||| @ content labeled content to write
34  writeFile : Poset labelType
35           => {l,l',l'' : labelType}
36           -> {auto flow : l `leq` l'}
37           -> (file : SecFile l'')
38           -> {auto flow' : l' `leq` l''}
39           -> (content : Labeled l' String)
40           -> DIO l (Labeled l'' (Either FileError ()))
41  writeFile (MkSecFile path) (MkLabeled content)
42    = lift $ map MkLabeled $ writeFile path content
```