

# Encoding DCC in Haskell

Maximilian Alghed  
alghed@chalmers.se

Chalmers University of Technology, Sweden

Alejandro Russo  
russo@chalmers.se

Chalmers University of Technology, Sweden

## ABSTRACT

The seminal work on the Dependency Core Calculus (DCC) [1] shows how monads not only can be used for embedding effects in purely functional languages but also to statically track data dependencies. Such types of analysis have applications in research areas like security, partial evaluation, and slicing, where DCC plays the role of a unifying formalism. For a Haskell programmer, putting DCC into practice raises many interesting conceptual and implementation concerns. Specifically, DCC uses a non-standard bind operator, i.e., with a different type signature than that provided by monads. In fact, embedding such non-standard bind operator opens the door for many design decisions. Furthermore, it is unclear if DCC extends to traditional methods used by Haskell programmers to handle effects (such as monad transformers). In this work, we describe a novel encoding of DCC in Haskell, with a focus on its use for security—although our results also apply to the other domains. We address the concerns mentioned above and show how our implementation of DCC can be seen as a direct translation from its typing rules via the use of *closed type families* and *type classes*—two advanced type system features of Haskell. We also analyze what kind of effects DCC is compatible with and which ones it cannot secure. We also derive an alternative formulation of DCC based on *fmap* and a corresponding non-standard *join*.

## CCS CONCEPTS

• **Security and privacy** → **Information flow control; Formal security models**; • **Software and its engineering** → **Software libraries and repositories**;

## KEYWORDS

Dependency Core Calculus, Information-Flow Control, Haskell

## 1 INTRODUCTION

Building applications which preserve privacy of data is an open and difficult problem. It is tempting to think that applying some sort of mechanism to grant or deny access to information will be sufficient to preserve secrets; *this is not the case*. Often, code is required to have access to sensitive data in order to deliver its functionality, but what we would like to know is if the code treats private data adequately, i.e., it does not leak it! The following example illustrates this point. It

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

PLAS 2017, October 30 2017, Dallas, Texas

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-5099-0/17/10...\$15.00

<https://doi.org/https://doi.org/10.1145/3139337.3139338>

involves developers Alice and Bob, who are collaboratively writing a web application for *something very important* (we leave that up to the imagination of the reader).

*Example 1.1.* In an effort to make sure users' private data is kept safe, the web application is implemented using separated *public* and *private* databases. While implementing the user registration page, Alice has requested Bob's help in managing what data to commit to the database. Alice writes the following code.

```
Alice
webApp :: IO ()
webApp = do
  ...
  formData ← getUserRegisterForm
  let (pubDBEnts, privDBEnts) = mkDBEntries formData
  commitToPublicDB pubDBEnts
  commitToPrivateDB privDBEnts
  ...
```

The symbol `::` is used to describe the type of terms in Haskell. The function `getUserRegisterForm` retrieves the registration page data from the user and the `...` denotes some part of the code that is not relevant for the point being made. On Alice's request, Bob has written the function `mkDBEntries`<sup>1</sup>. We use the standard list syntax `[x1, x2, ... , xn]` for describing extensible lists and `[a]` for denoting the type of lists of elements of type `a`.

```
Bob
mkDBEntries :: FormData → ([DBEntry], [DBEntry])
mkDBEntries form =
  ([screenName form, userEmail form, userName form],
   [hash (password form)])
```

What Bob does not know is that the name `(userName form)` of the user is not meant to be public information. This fact was hidden deep in a design document somewhere, but Bob was too lazy to read that document.

Information-Flow Control (IFC) is a research area dedicated to providing guarantees that private data is kept confidential by applying programming languages techniques [12, 20]. Haskell is one of a few languages capable of guaranteeing IFC through libraries [15]. As long as applications adhere to the library's API, they are guaranteed to keep private data confined [6, 18, 19, 26]. These IFC libraries enforce security in the presence of many advanced programming languages features (e.g., concurrency) and often intertwine effect-free and effectful computations into a single monad.

<sup>1</sup> To simplify the presentation we take the type `DBEntry` to be `String` and consequently `hash :: String → String`. While a real database library will have a much richer representation of database entries, this simplification has no impact on the point being made.

$$\frac{\Gamma \vdash x : a}{\Gamma \vdash \text{return}_\ell x : T_\ell(a)}$$

$$\frac{\Gamma \vdash e : T_\ell(a) \quad \Gamma \vdash f : a \rightarrow s \quad \ell \leq s}{\Gamma \vdash e \gg>= f : s}$$

Figure 1: Core typing rules for DCC.

$$(R1) \frac{\ell \sqsubseteq \ell'}{\ell \leq T_{\ell'}(s)} \quad \frac{\ell \leq s \quad \ell \leq t}{\ell \leq (s \times t)} \quad \frac{\ell \leq t}{\ell \leq (s \rightarrow t)}$$

$$(R2) \frac{\ell \leq s}{\ell \leq T_{\ell'}(s)} \quad (R3) \frac{}{\ell \leq ()}$$

Figure 2: The  $\leq$ -relation

In this paper, we take a different approach by building an IFC library based on DCC [1]—a calculus designed to secure *pure computations*. We show how it is possible to decouple the core (pure) security concerns of applications from its (monadic) effects. As illustrated by the example above, getting IFC right is a tricky business and we are going to show how DCC can help us get it right!

## 2 BACKGROUND

*Attacker model.* We adopt the same attacker model as the one considered in the original DCC work [1]: *potentially malicious or careless programmers provide code that manipulates secret data; data which should not be leaked into public channels*. This choice is also aligned with the attacker model considered by other IFC libraries (e.g., LIO [26] and MAC [18]). The security policy that our implementation will enforce is *non-interference* [11]. Similarly to previous work [4, 30], we focus on the terminating fragment of DCC, thus ignoring leaks due to abnormal termination of programs—a security condition known as *termination-insensitive non-interference*. We also consider the treatment of other covert channels [14] like power consumption or memory footprint out of scope.

*Dependency Core Calculus.* DCC is a variant of the lambda calculus which features sums, products, and a family of identity monads, written  $T$ , which are indexed by elements of a lattice. For the purpose of this paper, elements of the lattice represent security levels denoting the sensitivity of data [8]. Importantly, the order relation of the lattice indicates how information is allowed to flow:  $\ell \sqsubseteq \ell'$  dictates that information with sensitivity  $\ell$  can flow into entities of sensitivity  $\ell'$ . For simplicity, we restrict ourselves to the classic two-point lattice with levels  $H$  (for high or private) and  $L$  (for low or public), where  $H \not\sqsubseteq L$  is the only disallowed flow.

The key property of DCC is that, once a value gets *protected* (wrapped) by  $T_\ell$ , any computation that depends on it is forced to reside in the same or a more sensitive family member of  $T$ . That is to say that, once data is considered private, it may only influence private computations. In contrast, public data is allowed

to influence both private and public computations. To capture this property, DCC provides a special typing rule for *bind*—after all, this is the operator which extracts values from computations of type  $T_\ell$ . Such a typing rule prevents, for instance, public computations, of type  $T_L(a)$ , from depending on sensitive ones, of type  $T_H(a)$ .

Figure 1 shows the typing rules for return and bind, respectively. We denote bind as  $(\gg>=)$ <sup>2</sup>—observe the extra  $>$  character compared to the bind operation  $(\gg=)$  from Haskell. The notation  $\Gamma \vdash e : \tau$  denotes the judgement “given the assumptions in  $\Gamma$ , the expression  $e$  has type  $\tau$ ”, where  $\Gamma$  is on the form  $a_0 : \tau_0, a_1 : \tau_1, \dots$  for some variables  $a_0, a_1, \dots$  and types  $\tau_0, \tau_1, \dots$ . The  $(\gg>=)$  operator takes a computation  $m$  and a function  $f$  and computes the result of  $m$  and applies  $f$  to that result, yielding a new computation. For a functional programmer familiar with monads, DCC’s bind has the *non-standard type*

$$\ell \leq s \Rightarrow T_\ell(a) \rightarrow (a \rightarrow s) \rightarrow s,$$

where the type  $s$  is subject to the restriction  $\ell \leq s$  (explained below), rather than the conventional type

$$T_\ell(a) \rightarrow (a \rightarrow T_\ell(b)) \rightarrow T_\ell(b).$$

In Section 3, we will show that  $(\gg>=)$  subsumes the standard bind operator, we therefore refer to it as *super-bind* from now on.

The typing rules for *return* $_\ell$  and  $(\gg>=)$  differ from the traditional typing rules for monadic operations in two ways. Because DCC is a label-monomorphic calculus, the primitive *return* $_\ell$  features a sub-index  $\ell$  denoting the family member which protects the given argument<sup>3</sup>. Other than that, this primitive is completely standard. Intuitively, the rule for super-bind forces the result to be protected at the same, or higher, level as the input. This last invariant is enforced by the side condition  $\ell \leq s$  (pronounce “ $s$  is protected at  $\ell$ ”).

Figure 2 formally introduces the  $\leq$ -relation. The judgement  $\ell \leq s$  is defined by structural induction on  $s$ :

- $T_{\ell'}(s)$  is protected at  $\ell$  when  $\ell \sqsubseteq \ell'$  or  $s$  is protected at  $\ell$ .
- $(s \times t)$  is protected at  $\ell$  when  $s$  and  $t$  are both protected at  $\ell$ .
- $s \rightarrow t$  is protected at  $\ell$  when  $t$  is protected at  $\ell$ .
- $()$  is protected at  $\ell$ .

This definition is very similar to the original definition by Abadi et al., it only differs in the rule (R3), which was first introduced by Tse and Zdancewic in [31].

We provide some examples to clarify how super-bind works.

*Example 2.1 (Respecting the security lattice).* The super-bind allows computations to be built as long as they follow the order-relation described by the security lattice. For instance, super-bind type-checks if the computation stays within the same level<sup>4</sup>.

$$x : T_L(\text{Bool}) \vdash x \gg>= \lambda y. \text{return}_L \neg y : T_L(\text{Bool})$$

Likewise, super-bind type-checks if the result is more sensitive than the input.

$$x : T_L(\text{Int}) \vdash x \gg>= \lambda n. \text{return}_H n : T_H(\text{Int})$$

<sup>2</sup>In the seminal work on DCC, the authors introduce bind as a let-like binding, **bind**  $x = e$  in  $e'$ . Instead, we denote it as  $(\gg>=)$  since it looks more familiar for a Haskell programmer. This choice is, however, only for improving the presentation.

<sup>3</sup>Observe, however, that sub-indexes are not strictly necessary since types can provide those indexes too.

<sup>4</sup>Technically, we take the type *Bool* to denote the DCC sum type  $() + ()$ , and type *Int* to denote finite size integers represented as products of *Bools*.

```

{-# LANGUAGE DataKinds #-}
{-# LANGUAGE TypeFamilies #-}
data Lattice = L | H
type family ℓ ⊆ ℓ' where
  H ⊆ L = False
  _ ⊆ _ = True

```

Figure 3: A lattice implemented with type families.

Observe that we started with a public value and end up with a sensitive one. However, downgrading the sensitivity of the data is disallowed. For example, the following program is ill-typed as  $H \not\sqsubseteq L$ .

```
returnH 42 >>>= λn. returnL n
```

*Example 2.2 (Statically predicted values).* DCC leverages types to identify certain values which are independent of secrets. This aspect provides a more flexible super-bind, i.e., a super-bind which does not demand protection on the resulting value when it is unitary. For instance, code below type-checks in DCC.

```
∅ ⊢ returnH 42 >>>= λn.() : ()
```

The symbol  $\emptyset$  denotes the empty typing environment. By examining the type of  $\lambda n.()$ , it is known that the produced value is unit and independent of the sensitive information—thus, it type-checks!

In a similar manner, protecting products is simply protecting the two components. After all, it is statically known that there are only two components in a product, so this fact is independent of secrets! The next piece of code is therefore well-typed.

```
∅ ⊢ returnH 41 >>>= λn.(returnH (n + 1), ()) : (TH(Int), ())
```

*Example 2.3 (Sums).* Unlike products, super-bind cannot be flexible when it comes to sums. To protect a sum, it is not enough to just protect the content of their options. This is mainly due to the fact that the type of sums does not reflect the option being taken—thus, it can be misused to reveal information. The following ill-typed program shows that.

```
returnH s >>>= λn. if n == 0 then inl(returnH 1)
                    else inr(returnH 0)
```

If this code were well-typed, it would be enough to observe if the returned value is a left (*inl*) or right (*inr*) injection to know if the secret is zero—despite the content of the options being protected.

### 3 EMBEDDING DCC IN HASKELL

In this section, we show how to implement DCC in Haskell<sup>5</sup> by leveraging on some of the recently added features of its type system. Importantly, the obtained code can be seen as a syntactic translation of the rules described in Figures 1 and 2, which gives us not only an elegant code but also confidence about its correctness.

<sup>5</sup>The code from this paper is available online at <https://github.com/MaximilianAlgehed/DCC>

```

-- definition of T as an abstract data type
data T (ℓ :: Lattice) a = TTCB { unTTCB :: a }
-- type signatures
return :: a → T ℓ a
(≫) :: (ℓ ≤ s) ~ True ⇒ T ℓ a → (a → s) → s
-- implementation
return = TTCB
t ≫ f = f (unTTCB t)

```

Figure 4: Core DCC primitives

```

type family ℓ ≤ t where
  ℓ ≤ (T ℓ' s) = ℓ ⊆ ℓ' ∨ ℓ ≤ s
  ℓ ≤ (s, t)   = ℓ ≤ s ∧ ℓ ≤ t
  ℓ ≤ (s → t) = ℓ ≤ t
  ℓ ≤ ()       = True
  ℓ ≤ s        = False

```

Figure 5: The  $\leq$ -relation as a closed type family.

*Security lattice.* Because we want to use Haskell’s type system to implement the security features of DCC, we need to represent our lattice at the type level. Figure 3 shows our implementation for the two-point lattice. It uses Haskell’s extensions for *closed type families* [10] and *data kinds* [35]. In a nutshell, closed type families are no more than type-level functions. The code defines such a function by introducing the type-level  $\sqsubseteq$ -relation on two arguments  $\ell$  and  $\ell'$ . The *DataKinds* extension, on the other hand, is not strictly necessary but allows booleans and security level constructors ( $L$  and  $H$ ) to be promoted to the type level—hence, it yields more elegant and self-descriptive code. In other words, constructors  $L$  and  $H$  belonging to the data type *Lattice* can be interpreted as types too! In fact, the equations in Figure 3 involve types (and not terms)  $L$ ,  $H$ , *True*, and *False*. The key part of the implementation is that the equations capture the expected security policy. Namely, as pointed out in Section 2,  $H \not\sqsubseteq L$  is the only disallowed flow. It is also possible to encode more complex (potentially infinite) security lattices using type families [6]—e.g., the lattice induced by *disjunction category labels* [25], where security levels express concerns of mutually distrusting principals.

*Dependency tracking.* The key component necessary to obtain DCC-style dependency tracking is the implementation of the *T*-monad family and its corresponding operations. Figure 4 shows our Haskell definitions for *T*, *return* (*return*), and super-bind ( $\gg$ ). At first glance, the code looks almost like a mere syntactic translation of the typing rules from Figure 1—except for the equality type constraint  $(\ell \leq s) \sim \text{True}$  (explained below). Just as in DCC, the type constructor *T* is a family of identity monads indexed by a security level  $\ell$  coming from the kind *Lattice*. The definition of *T* involves “constructor”  $T^{\text{TCB}} :: a \rightarrow T \ell a$  and the “destructor”  $\text{unT}^{\text{TCB}} :: T \ell a \rightarrow$

a. These component are part of the *trusted computing base* or TCB (observe the supra-index). Consequently, it is not accessible to the users of our DCC implementation but is part of its internals. Otherwise, for instance, access to the constructor  $T^{\text{TCB}}$  would allow attacker’s code to pattern match on a value  $v :: T^{\text{TCB}} H a$  to simply extract the secret value (e.g., `case v of  $T^{\text{TCB}} secret \rightarrow secret$` ).

Because Haskell is polymorphic, the type of `return` can capture the typing rules for the entire family of `return $_{\ell}$` , for any given  $\ell$ . Similarly, the implementation of `( $\gg$ )` has almost the same type ( $\ell \leq s$ ) $\sim True \Rightarrow T \ell a \rightarrow (a \rightarrow s) \rightarrow s$  as the super-bind `( $\gg\gg$ )` :  $T l a \rightarrow (a \rightarrow s) \rightarrow s$  in DCC. The difference being that, as  $s$  is restricted by  $\ell \leq s$  in DCC, our implementation requires discharging the type equality constraint  $(\ell \leq s) \sim True$  to enforce the same characteristic over  $s$ . More specifically, the constraint  $(\ell \leq s) \sim True$  means that the application of the type-level function  $\ell \leq s$  evaluates to (the type) `True`. In the whole implementation, `( $\gg$ )` is the only operation which utilizes `un $T^{\text{TCB}}$` , and therefore is the *only primitive that we must trust to work as expected*—clearly, we must trust our type system and type declarations too! As the Haskell language permits certain unsafe operations<sup>6</sup> which might break type-safety and data encapsulation, we also rely on `SafeHaskell` [28]—an extension of the Haskell compiler GHC which restricts the use of unsafe features to trusted code. We now turn our attention to the type family implementing the  $\leq$ -relation.

*Protecting data.* Figure 5 encodes the  $\leq$ -relation of DCC as a closed type family with the same name. The code is almost an entirely syntactic translation of the  $\leq$ -relation from Figure 2 except for the first two equations which use type-level versions of operators  $\wedge$  and  $\vee$ . However, these cases follow directly from the structure of the  $\leq$ -relation. For instance, the first equation collapses rules (R1) and (R2) into one disjunction. The rest of the cases are self-explanatory and therefore we omit describing them<sup>7</sup>.

*Why closed type families?* Haskell supports type families, i.e., type-level functions, both in an open and closed variant. An open type family allows equations defining it being scattered throughout the code base. The type family definitions for the  $\sqsubseteq$ - and  $\leq$ -relations are instead closed. This means that (i) equations may not be added outside the initial ones, and (ii) equations are always tried in order of appearance. Condition (i) is essential to guarantee security. Otherwise, for instance, if these type families were not closed, an attacker (or careless programmer) could inadvertently add the equation  $H \sqsubseteq L = True$ , or  $\ell \sqsubseteq s = True$ —thereby nullifying all security properties given by `( $\gg$ )`. In contrast, condition (ii) is essential to avoid being overly restrictive. For instance, if equation  $\ell \leq s = False$  were to trigger first, many programs would be unnecessarily classified as insecure.

### 3.1 Revisited example

Let us see how the solution to Alice and Bob’s problem can be translated from the theoretical calculus to Haskell.

*Example 3.1.* Alice starts by defining the security levels of her data as follows.

<sup>6</sup>like `unsafePerformIO :: IO a  $\rightarrow$  a` and `unsafeCoerce :: a  $\rightarrow$  b`

<sup>7</sup>We note that, while the implementation requires the extension `UndecidableInstances`, the definition of the type family is always terminating.

```
instance Monad (T  $\ell$ ) where
  return :: a  $\rightarrow$  T  $\ell$  a
  return = return
  ( $\gg$ ) :: T  $\ell$  a  $\rightarrow$  (a  $\rightarrow$  T  $\ell$  b)  $\rightarrow$  T  $\ell$  b
  ( $\gg$ ) = ( $\gg\gg$ )
```

Figure 6: Monad instance for  $T \ell$

```
Alice
data FormData = FD { password :: T H String,
  userName     :: T H String,
  userEmail    :: T L String,
  screenName   :: T L String }
```

With this implementation of `FormData`, Bob’s old code will no longer be type correct. This is because `userName` now returns a `String` protected at  $H$ , but the return type of the `mkDBEntries` expects a `T L String`. As noted previously, Bob has no way to go from a `T H String` to a `T L String`. The type system guarantees that the real name of the user is kept safe. He is forced to rewrite his implementation of `mkDBEntries` as follows.

```
Bob
mkDBEntries :: FormData  $\rightarrow$ 
  (T L [DBEntry], T H [DBEntry])
mkDBEntries form =
  (sequence [screenName form, userEmail form]
   , sequence [hash (password form), userName form])
```

The `sequence :: [T  $\ell$  a]  $\rightarrow$  T  $\ell$  [a]` function collects the results produced by a list of protected computations and it is implemented as a derived operation, i.e., it does not break DCC abstractions. Observe that the name of the user is now in the correct list of database entries.

*Monads in Haskell and super-bind.* To convince Haskell’s compiler that  $T \ell$  is a monad, it is necessary to provide a *Monad* type-class instance for  $T \ell$ , i.e., we need to implement the `return` and `( $\gg$ )` operators over that type. As claimed above, super-bind is general enough to implement `( $\gg$ )`. Figure 6 shows that, by specializing the type of `( $\gg$ )`, each family member  $T \ell$ , for a certain  $\ell$ , can indeed be made an instance of the *Monad* type class.

*Type errors.* The reader may be concerned that the use of advanced features of the Haskell type system, like closed type families and data kinds give rise to complicated and difficult to interpret type errors. Indeed, this appears to be the case, in the worst cases type errors are on the form “Could not unify ‘False with ‘True, arising from the use of `>>>=`”. However, recent work by Serrano and Hage [22] provides methods for improving type errors in Haskell programs by defining domain specific type errors. We note the incorporation of such techniques into our encoding of DCC as future work.

$$\begin{aligned}
m \gg f & \equiv \text{join } (\text{fmap } f \ m) \\
m \gg \text{id} & \equiv \text{join } m \\
m \gg (\text{return } \circ f) & \equiv \text{fmap } f \ m
\end{aligned}$$

**Figure 7: Laws relating ( $\gg$ ) and return with join and fmap.**

*Discussion.* The experienced functional programmer might propose the following combinator-based interface as an alternative to our formulation of DCC.

$$\begin{aligned}
\text{upgrade} & :: (\ell \leq h) \sim \text{True} \Rightarrow T \ell \ a \rightarrow T \ h \ a \\
\text{split} & :: T \ell \ (a, b) \rightarrow (T \ell \ a, T \ell \ b)
\end{aligned}$$

The advantage of such an approach would be that it makes the security-related properties of the  $T$  monad family more explicit. However, it is not obvious what a sound and complete set of such combinators would look like. Furthermore, basing the translation on such combinators rather than our direct syntactic translation nullifies our argument that the translation is correct by construction. One observation is that any correct combinator-based translation should be implementable using our translation. For example, the *upgrade* and *split* combinators are trivially implementable using super-bind, thereby providing an interface which would be familiar to functional programmers while retaining the confidence provided by the syntactic translation.

Another alternative is the use of generalized monads [5] to keep track of the dependencies. Possibly employing a type like:

$$(\gg) :: T \ell \ a \rightarrow (a \rightarrow T \ell' \ b) \rightarrow T (\ell \sqcup \ell') \ b$$

However, this would represent a significant departure from the original DCC formalism and we therefore consider an analysis of this approach outside the scope of this paper. The interested reader may consult [7] for a detailed discussion of this alternative approach.

## 4 AN ALTERNATIVE FORMULATION

It is known that every computations involving a monad  $m$  may be built by (at least) two different interfaces, one using

$$\begin{aligned}
\text{return} & :: a \rightarrow m \ a \\
\text{fmap} & :: (a \rightarrow b) \rightarrow m \ a \rightarrow m \ b \\
\text{join} & :: m \ (m \ a) \rightarrow m \ a
\end{aligned}$$

and the other using

$$\begin{aligned}
\text{return} & :: a \rightarrow m \ a \\
(\gg) & :: m \ a \rightarrow (a \rightarrow m \ b) \rightarrow m \ b
\end{aligned}$$

Every interface has its pros and cons and programmers are likely to use both—we refer the interested reader to [34] for further details. In this light, given the non-standard notion of ( $\gg$ ), a natural question is what such an isomorphism would look like for DCC? Before we start answering that question, Figure 7 recaps the laws establishing the correspondence between these two formulations.

We start by finding the equivalent notion of *join* and *fmap* in DCC. However, beare in mind that we expect to end up with a

$$\begin{aligned}
\text{join}^{\text{DCC}} & :: (\ell \leq s) \sim \text{True} \Rightarrow T \ell \ s \rightarrow s \\
\text{join}^{\text{DCC}} \ t & = t \gg \text{id} \\
\text{fmap}^{\text{DCC}} & :: (\ell \sqsubseteq \ell') \sim \text{True} \Rightarrow (s \rightarrow t) \rightarrow T \ell \ s \rightarrow T \ell' \ t \\
\text{fmap}^{\text{DCC}} \ t \ f & = t \gg (\text{return } \circ f)
\end{aligned}$$

**Figure 8: Types for  $\text{join}^{\text{DCC}}$  and  $\text{fmap}^{\text{DCC}}$ .**

“super-join” and “super-fmap” so we will refer to them as  $\text{join}^{\text{DCC}}$  and  $\text{fmap}^{\text{DCC}}$ , respectively. The code for these functions are given in Figure 8. These functions are simply computed by replacing all occurrences of ( $\gg$ ) with ( $\gg$ ), *join* with  $\text{join}^{\text{DCC}}$ , and *fmap* with  $\text{fmap}^{\text{DCC}}$  in Figure 7. The derived  $\text{join}^{\text{DCC}}$  inherits the constraint from ( $\gg$ ) that checks if a certain type  $s$  is protected by  $\ell$ —i.e.,  $(\ell \leq s) \sim \text{True}$ . This function basically removes the  $(T \ell)$ -monad when the value that it protects is already protected by  $\ell$ . The derived  $\text{fmap}^{\text{DCC}}$ , on the other hand, allows us to map a function over the content of the  $(T \ell)$ -monad but potentially returning the result in a more sensitive family member  $T \ell'$ , where  $\ell \sqsubseteq \ell'$ .

Unfortunately, we need to go through some contortions when using  $\text{join}^{\text{DCC}}$  and  $\text{fmap}^{\text{DCC}}$  in a way similar to how we would use *join* and *fmap*, respectively. To illustrate this point, we consider the case where we have a computation  $t :: T \ H \ (T \ L \ a)$  that we want to “join” (collapse) in order to return something of type  $T \ H \ a$ . We then proceed to use super-join as follow.

$$\begin{aligned}
t' & :: T \ H \ a \\
t' & = \text{join}^{\text{DCC}} (\text{join}^{\text{DCC}} (\text{fmap}^{\text{DCC}} (\text{fmap}^{\text{DCC}} \ \text{return}_H) \ t))
\end{aligned}$$

Observe that the problem with this example is that  $T \ L \ a$  is not protected by  $H$ , therefore, we cannot simply apply  $\text{join}^{\text{DCC}}$ —the type constraint  $H \leq (T \ L \ a)$  does not hold! Instead, we need to transform  $t$  into something of type  $T \ H \ (T \ L \ (T \ H \ a))$  by calling  $\text{fmap}^{\text{DCC}} (\text{fmap}^{\text{DCC}} \ \text{return}_H) \ t$ . After that, we take the two inner  $T$ 's, i.e.,  $T \ L \ (T \ H \ a)$ , and collapse them to  $T \ H \ a$ —this is done by the inner  $\text{join}^{\text{DCC}}$ . To conclude, the outer  $\text{join}^{\text{DCC}}$  takes something of type  $T \ H \ (T \ H \ a)$  and collapses it to  $T \ H \ a$  as wanted.

As types become more complex, for example when we have a term  $t$  of type  $T \ H \ (T \ L \ a, T \ H \ b)$ , we would expect to be able to use  $\text{join}^{\text{DCC}}$  to get a (collapsed) term  $t'$  of type  $(T \ H \ a, T \ H \ b)$ , but once again we arrive at a verbose implementation.

$$\begin{aligned}
t' & :: (T \ H \ a, T \ H \ a) \\
t' & = \text{join}^{\text{DCC}} (\text{fmap}^{\text{DCC}} \\
& \quad (\lambda(t_l, t_h) \rightarrow (\text{join}^{\text{DCC}} (\text{fmap}^{\text{DCC}} \ \text{return}_H \ t_l), t_h)) \ t)
\end{aligned}$$

As in the example above, we need to convert the first component of the pair from something of type  $T \ L \ a$  to something of type  $T \ L \ (T \ H \ a)$  to be able to collapse it to something of type  $T \ H \ a$ . Once that is done, the outer  $\text{join}^{\text{DCC}}$  collapses  $T \ H \ (T \ H \ a, T \ H \ a)$  to  $(T \ H \ a, T \ H \ a)$  as expected.

All these contortions arise from having public values wrapped in a sensitive  $T$ -family member, which makes them effectively non-public values. To convince the type system that this is the case,

```

class Upgrade a b where
  up :: a → b
instance (ℓ ⊆ ℓ') ~ True ⇒ Upgrade (T ℓ a) (T ℓ' a) where
  up t = t ≫≡ return

```

Figure 9: A type class for automatic upgrading.

```

joinDCC' :: (Upgrade s s', (ℓ ≤ s) ~ True) ⇒ T ℓ s → s'
joinDCC' t = t ≫≡ up

```

Figure 10:  $join^{DCC'}$  with automatic upgrading.

we need to take those public values and *upgrade* them to sensitive ones, i.e., transforming them from something of type  $T L a$  into something of type  $T L (T H a)$  with  $fmap^{DCC}$ . After that, the  $(T L)$ -monad can be removed using  $join^{DCC}$ . To simplify the life of the programmer, we implement the type class *Upgrade a b* in Figure 9. The type class provides the (overloaded) operators  $up :: a \rightarrow b$  which precisely implements the procedure for upgrading values of type  $a$  into values of type  $b$  to have potentially more sensitive labels. The most interesting cases for  $up$  is the instance involving the  $T$ -monad family present in Figure 9—see Appendix A.1 for more details. In that case, to upgrade  $T \ell a$  into  $T \ell' a$ , it must be the case that  $\ell'$  is at least as sensitive as  $\ell$  ( $\ell \sqsubseteq \ell'$ ). Observe that the instance demands the type constraint  $(\ell \sqsubseteq \ell') \sim True$ , which arises from the use of  $(\gg\equiv)$ .

With the type class *Upgrade* in place, we proceed to change the definition of  $join^{DCC}$  to automatically adjust labels when needed—see Figure 10. Different from  $join^{DCC}$ ,  $join^{DCC'}$  allows for a more elegant code. We revisit the examples shown before. When we have a term  $t :: T L (T H a)$ , we can transform it to a term of type  $t :: T H a$  as follows.

```

t' :: T H a
t' = joinDCC' t

```

Similarly, when we have a term of type  $T H (T L a, T H b)$ , we can make an expression of type  $(T H a, T H b)$  as follows.

```

t' :: (T H a, T H b)
t' = joinDCC' t

```

The equations in Figure 7 hold for  $join^{DCC'}$  provided that the type constraint *Upgrade s s'* can be satisfied. For instance, if the  $up$  method does not change the type, that is to say *Upgrade s s* holds, then we have exactly the same equations that we had previously.

It is also possible to derive  $(\gg\equiv)$  from  $join^{DCC'}$  and  $fmap^{DCC}$ . This encoding requires  $join^{DCC'}$  and  $fmap^{DCC}$  to use  $unT^{TCB}$  in their definitions. A natural question is then if such occurrences of  $unT^{TCB}$  are sound—i.e., it does not provide opportunities to leak information. We argue that, as long as the equations in Figure 7 hold for the

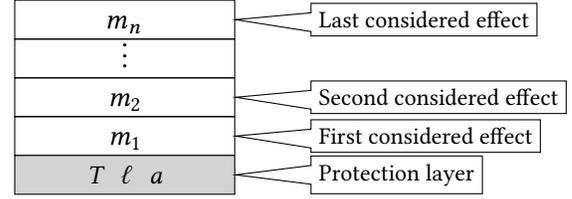


Figure 11:  $(T \ell)$ -monad at the base of monadic stacks

```

type Outputs o
type WT o ℓ a = WriterT (Outputs o) (T ℓ) a

```

Figure 12: Adding outputs.

DCC variants, it is possible to translate every occurrence of  $join^{DCC'}$  and  $fmap^{DCC}$  back into  $(\gg\equiv)$ , which in turn uses  $unT^{TCB}$  safely due to its type constraint. We leave the primitive definitions of  $join^{DCC'}$  and  $fmap^{DCC}$  as well as the derivation of  $(\gg\equiv)$  as an exercise for the reader.

## 5 SIDE-EFFECTS

In this section, we show how to extend DCC to handle effects. Being a family of identity monads,  $T$  does not provide any interesting side-effect on its own. At a first glance, it is not clear that effects can be naturally handled by DCC without having to trust the implementation of each and every effect. Approaches for information-flow control in Haskell have typically *intertwined security concerns with monadic effects*, i.e., side-effects are performed by primitive operations of their respective security monads (see for example [18, 19, 26]). Instead, we show that, as long as side-effects can be mimicked or denoted by pure computations [33], they can be added "on top" of our implementation of DCC. In other words, we show that the implementation of such effects need not be inherently trusted, but may be added to a monad of type  $T \ell$  without using  $unT^{TCB}$ —the only primitive that must be handled with care. The main idea that makes this possible is to place the  $(T \ell)$ -monad at the bottom of a stack of monads (where each layer implements an effect). By doing so, data handled by upper layers is wrapped, i.e., protected, by  $T \ell$ . Figure 11 depicts this idea.

The rest of the section shows the derivation process of securely adding the most common effects to DCC, namely outputs, state, and error handling. For simplicity, we focus initially on effects provided by the *Writer* monad, a monad which equips computations with outputs.

### 5.1 Outputs

To start adding outputs, we resort to one of the most frequent tools that Haskell programmers rely on: *monad transformers* [16]. In a nutshell, a monad transformer provides effects similarly to regular

```

class Project (ℓ :: Lattice) where
  type SegmentsAbove ℓ o
  project :: Segments o → (T ℓ o, SegmentsAbove ℓ o)
  merge :: (T ℓ o, SegmentsAbove ℓ o) → Segments o
         → Segments o

```

Figure 13: The *Project* type class.

monads, however a monad transformer cannot stand on its own; instead, it modifies the behavior of an underlying monad by means of pure functions to incorporate the effects associated with the transformer. Figure 12 shows the interface for *WT* (the *Writer-T* monad). The  $WT\ o\ \ell$  monad is built by applying the monad transformer *WriterT* to  $T\ \ell$ , creating the type *WriterT* (*Outputs*  $o$ ) ( $T\ \ell$ )  $a$ , which extends the  $(T\ \ell)$ -monad with outputs of type *Outputs*  $o$ . In IFC systems, end-points get associated with a security level denoting the sensitivity of the data entering or leaving the system. Having classified the sensitivity of outputs, it is possible to detect when information might get leaked, e.g., when a program tries to produce a sensitive output in the public partition or part of the system. Hence, we consider outputs as partitioned per security level—this is exactly what the type *Outputs*  $o$  denotes. For the two-point lattice, outputs are partitioned into two segments: one for sensitive and one for public outputs—observe that every partition is protected by  $T$ .

```

type Segments o = (T L o, T H o)
type Outputs o = Segments o

```

We use the name *Segments* instead of *Outputs* as the idea of partitioning data by security level is more general than outputs and we will show how it extends when adding state as side-effect later.

In the writer monad, the operation producing outputs is the function  $tell :: MonadWriter\ o\ m \Rightarrow o \rightarrow m\ ()$ , where  $m$  is the type of the writer monad and  $o$  is the type of the output being produced. What would be a secure version of *tell*? Intuitively, it should be the case that taking something at confidentiality level  $\ell_L$ , it should be only possible to write in the output partition label with  $\ell_H$ , where  $\ell_H$  is the same as or higher than  $\ell_L$ , that is  $\ell_L \sqsubseteq \ell_H$ . This is known as the no-write-down policy [3]. A computation is secure as long as it does not write data to less sensitive sinks. We will show that DCC guides us to a secure implementation of *tell*, called  $tell^{WT}$ , which implements precisely the no-write-down policy. To help readers, we indicate the relationship between type variables in their subindexes, we use type variables  $\ell_L$  and  $\ell_H$  to attest that  $\ell_L \sqsubseteq \ell_H$ . Before diving into the details of  $tell^{WT}$ , we need to implement some auxiliary functions.

For the rest of the section, we need to consider that outputs have a monoid structure, i.e., they have both a neutral element and an (associative) operations to concatenate outputs—see details in Appendix A.2. Figure 13 shows an auxiliary type class needed to provide secure outputs. The type class provide operations to *manipulate the space of outputs which are, at least, as sensitive as a given security level ℓ*. The associated type *SegmentsAbove ℓ* defines the

```

instance Project H where
  type SegmentsAbove H o = ()
  project (_, h) = (h, ())
  merge (h, ()) (ℓ, _) = (ℓ, h)

instance Project L where
  type SegmentsAbove L o = T H o
  project (ℓ, h) = (ℓ, h)
  merge (ℓ, h) _ = (ℓ, h)

```

Figure 14: Instances for *Project H* and *Project L*

```

tellWT :: (Monoid o, Monoid (SegmentsAbove ℓ o)
           , Project ℓ o) ⇒ T ℓ o → WT o ℓ' ()
tellWT t = tell (merge t mempty)

```

Figure 15: Secure version of primitive *tell*.

structure of all the levels strictly above  $\ell$ . The *project* function takes the outputs and returns the partitions at level  $\ell$  and higher. The *merge* function, on the other hand, takes all segments (*Segments*  $o$ ) and replaces the values at level  $\ell$  and higher by the first argument ( $T\ \ell\ o, SegmentsAbove\ \ell\ o$ ). To illustrate how this type class works, Figure 14 shows instances useful for the two-point lattice, where outputs have type  $o$ . The instance *Project H* indicates that there are no outputs strictly more sensitive than  $H$  by defining type *SegmentsAbove H*  $o$  as unit, i.e.,  $()$ . Similarly, the instance *Project L* indicates that the memory strictly above  $L$  is of type  $T\ H\ o$  by defining *SegmentsAbove L*  $o$  as  $T\ H\ o$ . The implementation for *project* and *merge* of both instances are self-explanatory. Note that, so far, we do not need to trust any other aspects of our code than what we trusted before—we have not used  $unT^{CB}$  at all!

Figure 15 shows the secure, *derived*, implementation of *tell*. The type constraints in the type signature of  $tell^{WT}$  reflect that outputs and partitions have a monoid structure—see constraints *Monoid o* and *Monoid (SegmentsAbove ℓ o)*, respectively. The constraint *Project ℓ o* arising from the use of *merge* in the definition of  $tell^{WT}$ . This function is simply defined as merging the output  $t$  into an empty output (*mempty*). Ostensibly, the type signature looks insecure, i.e., the argument to  $tell^{WT}$  is protected at  $\ell$  ( $T\ \ell\ o$ ) but the resulting computation handles (protects) data at arbitrary label  $\ell'$  ( $WT\ o\ \ell'\ ()$ )—recall the definition of *WT* in Figure 12. However, precisely because the argument is of type  $T\ \ell\ o$ , it means that it will be *always* protected at  $\ell$ , no matter in what context we attempt to output it. As a result, calls to  $tell^{WT}$  with public arguments in sensitive contexts are allowed but no effects would be ever performed. To illustrate this point, we proceed to describe some examples.

*Example 5.1.* Consider the code below.

$$\begin{aligned} \text{cast}^{\text{WT}} &:: (\text{Monoid } o, \text{Project } \ell_{\mathbf{H}} o, \\ &\quad (\ell_{\mathbf{H}} \leq \text{OutputsAbove } \ell_{\mathbf{H}} o) \sim \text{True}) \\ &\Rightarrow \text{WT } o \ell_{\mathbf{H}} a \rightarrow \text{WT } o \ell_{\mathbf{L}} (T \ell_{\mathbf{H}} a) \end{aligned}$$

**Figure 16: Primitive  $\text{cast}^{\text{WT}}$**

```
secure :: Monoid o => T L o -> T H Bool -> WT o H ()
secure output secret = do bool <- lift secret
  when bool (tellWT output)
```

Where the  $\text{lift} :: T \ell a \rightarrow \text{WT } o \ell a$  function unlabels a  $T \ell$  value at the correct security level in the  $\text{WT}$  monad. The example seems to be insecure: depending on the value of the secret ( $\text{bool}$ ), the program produces a public output ( $\text{tell}^{\text{WT}} \text{ output}$ ). However, since the resulting type is  $\text{WT } o \text{H } ()$ , it means that all the outputs are indeed protected by  $\text{H}$ —thus, the code does not leak information! In other words, once in the monad  $\text{WT } o \text{H } a$ , any public output has no effect since the whole output domain is protected by  $\text{H}$ .

The reason why the example above is secure relies on the fact that monad transformers are implemented using pure functions and that the  $T \ell$ -monad is placed at the bottom of the monadic stack. By definition, the monad transformer  $\text{WriterT}$  (as most monad transformers do) manipulates the underlying monad to introduce the output effect. More concretely, the type  $\text{WT } o \text{H } a$ , i.e.,  $\text{WriterT } (\text{Outputs } o) (T \text{H } a)$ , is isomorphic to the type  $T \text{H } (a, \text{Outputs } o)$ . Observe how the monad transformer leverages on the  $T \text{H}$ -monad to carry around outputs. This fact implies that all the outputs are considered sensitive by  $\text{WT } o \text{H } a$ . Then, it is not surprising that  $\text{tell}^{\text{WT}}$  can be called with a public argument in sensitive contexts; after all, outputs are effectively secret!

*Example 5.2.* Consider instead the following code.

```
secure :: Monoid o => T H o -> T L Bool -> WT o L ()
secure output public = do bool <- lift public
  when bool (tellWT output)
```

This function is very similar to that in the example above, except that the value to output is now sensitive ( $T \text{H } o$ ), while the boolean ( $T \text{L } \text{Bool}$ ) and the writer monad are not ( $\text{WT } o \text{L } ()$ ). As before, the code is safe to execute, i.e., a sensitive output gets performed based on the value of a public boolean.

One could still argue that the type of  $\text{tell}^{\text{WT}}$  should restrict the programmer and give a type error in the case where  $\ell \not\sqsubseteq \ell'$  in order to hint about possible mistakes in the code. However, this would have no impact on security. The generality of  $\text{tell}^{\text{WT}}$  is comfortable; it allows for *securely* and indiscriminately writing to different outputs at different levels.

## 5.2 Label creep

Although the interface provided by  $\text{tell}^{\text{WT}}$  is secure, it leaves some things to be desired. As illustrated by Example 5.1, computations depending on sensitive data force the entire computation to be protected at  $\text{H}$ . In other words, *after inspecting a secret, public outputs*

are tainted with the label of the secret. This phenomenon is known as *label creep*: the label associated with the computation reaches a point where it is not possible to perform any useful side-effects—even if the effects are benign and do not attempt to leak information. We illustrate the label creep problem in the following example.

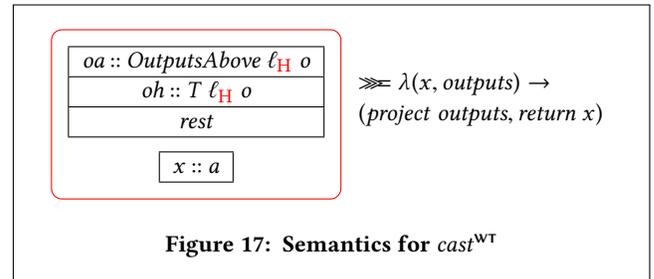
*Example 5.3.* Even though the following example is well typed, and therefore secure, the intended public output on the final line ( $\text{tell}^{\text{WT}} \text{ somethingPublic}$ ) occurs inside the  $(\text{WT } o \text{H})$ -monad. Consequently, it is considered secret.

```
creep :: Monoid o => T H Bool -> WT o H ()
creep th = do
  secret <- lift th
  -- From now on, no more public outputs
  when secret (tellWT somethingSecret) -- secret output
  tellWT somethingPublic -- intended public output
```

To mitigate the label creep problem, we introduce another *derived* function called  $\text{cast}^{\text{WT}}$ . Figure 16 shows the type signature for  $\text{cast}^{\text{WT}}$ . Starting with the types, the  $\text{Monoid } o$  constraint arises from the use of the  $\text{Writer}$  effect, while the other two deserve some explanation. The constraint  $\text{Project } \ell_{\mathbf{H}} o$  tells us that some manipulation on the sensitive part of the outputs takes place in  $\text{cast}^{\text{WT}}$  (explained below). Finally, the constraint  $(\ell_{\mathbf{H}} \leq \text{OutputsAbove } \ell_{\mathbf{H}} o) \sim \text{True}$  is redundant but appears in the type signature due to the inability of Haskell to figure out that  $\ell_{\mathbf{H}}$  protects all the outputs about itself.

The function  $\text{cast}^{\text{WT}}$  runs a computation in the  $(\text{WT } o \ell_{\mathbf{H}})$ -monad (given as argument) and injects the produced sensitive outputs as part of the outputs handled by the resulting  $(\text{WT } o \ell_{\mathbf{L}})$ -monad. It also protects the result of type  $a$  with label  $\ell_{\mathbf{H}}$ , that is  $T \ell_{\mathbf{H}} a$ . By applying  $\text{cast}^{\text{WT}}$ , we do not need to remain in the  $(\text{WT } o \text{H})$ -monad every time we need to inspect a secret. The role of  $\text{cast}^{\text{WT}}$  is to convert (cast) sensitive computations (at level  $\ell_{\mathbf{H}}$ ) to public ones (at level  $\ell_{\mathbf{L}}$ ) without leaking any information, while still retaining both the sensitive result and side-effects. Note that any public side-effects that may have been performed in the sensitive context are discarded by  $\text{cast}^{\text{WT}}$ , this is possible because we are in a pure setting. In the jargon of more traditional IFC systems,  $\text{cast}^{\text{WT}}$  is an operation which allows to temporarily raised the *program counter* [32].

Figure 17 shows a graphic rendering of the core operational semantics of  $\text{cast}^{\text{WT}}$ —we refer readers to Appendix A.3 for the concrete implementation. As a first step,  $\text{cast}^{\text{WT}}$  runs its argument of type  $\text{WT } \ell_{\mathbf{H}} a$ , which results in a result  $x :: a$  and outputs of type  $\text{Outputs } o$ —see left part of the graphic. From the produced outputs, we identify three segments of outputs: the outputs at level



**Figure 17: Semantics for  $\text{cast}^{\text{WT}}$**

$\ell_H$  and higher, respectively denoted by variable  $oh :: T \ell_H o$  and  $oa :: OutputsAbove \ell_H o$ , and outputs which are strictly less sensitive or incomparable to  $\ell_H$ , denoted by  $rest$  in the graphic. Function  $cast^{WT}$  then proceeds to deconstruct, by using super-bind, the returned value and the outputs ( $\lambda(x, outputs) \rightarrow \dots$ ). The outputs at level  $\ell_H$  and above are extracted (*project outputs*), an action that produces the tuple  $(oh, oa) :: (T \ell_H o, OutputsAbove \ell_H o)$ . Hence, it disregards any other outputs (i.e.,  $rest$ ) and considers *only* those that can be safely affected by the computation of type  $WT \ell_H a$ , i.e., those outputs which do not write down from the security level  $\ell_H$ . The key realisation behind the implementation of  $cast^{WT}$  is that, while sensitive computations may depend on low data, they may not influence low outputs; a property ensured by the type system!

The idea of  $cast^{WT}$  is common in other Haskell IFC libraries [18, 19, 26]. However, different from them,  $cast^{WT}$  is a *derived* operation rather than a trusted primitive. We should remark that the core idea of  $cast^{WT}$  is present in the appendix of the DCC paper [1] when showing an embedding of the imperative (stateful) SV-calculus in DCC. In the next section, we show that the core idea behind  $cast^{WT}$  can be extended to cover other effects mimicked by pure expressions, namely, state and error handling.

**5.2.1 Revisited running example.** With  $cast^{WT}$  in place, Alice and Bob can rewrite their application to make use of the power of monadic effects.

*Example 5.4.* As Alice decides to add more and more functionality to her website, she realises that a *Writer* effect is necessary to handle the many functions for generating database entries. Furthermore, Alice has decided to add a function  $isSecure :: String \rightarrow Bool$  which checks if a password is deemed secure.

#### Bob

```
mkDBEntries :: FormData → WT [DBEntry] L (T H Bool)
mkDBEntries form = do
  -- public output
  tellWT (sequence [screenName form, userEmail form])
  -- including a sensitive output
  castWT (mkSecretEntries form)
```

#### Bob

```
mkSecretEntries :: FormData → WT [DBEntry] H Bool
mkSecretEntries form = do
  pwd ← lift (password form)
  when (isSecure pwd)
    (tellWT (sequence
      [fmap hash (password form), userName form]))
  return (isSecure pwd)
```

With this implementation, Alice can decide when she, in the trusted code, obtains the result of  $mkDBEntries$  and if she is to commit the result to the database based on the result in the  $T H Bool$  coming from the action  $return sec$ . Note that  $cast^{WT}$  allows Bob to write the sensitive part of the code as a stand-alone module

```
-- Definition
type WET o ℓ a =
  WriterT (Outputs o) (EitherT String (T ℓ)) a

-- Outputs
tellWET :: (Monoid o, Project ℓ o) ⇒ T ℓ o → WET o ℓ' ()

-- API for error handling
throwError :: Monoid o ⇒ String → WET o ℓ ()
catchError :: Monoid o ⇒
  WET o ℓ a → (String → WET o ℓ a) → WET o ℓ a

-- Label creep
castWET :: (Monoid o, Project ℓH o,
  (ℓH ≤ OutputsAbove ℓH o) ~ True)
  ⇒ WET o ℓH a → WET o ℓL (T ℓH (Either Error a))
```

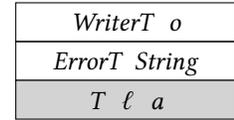
**Figure 18: Combination of exception and writer effects.**

## 6 COMBINING EFFECTS

In this section, we show that it is possible to derive secure operations capable of handling multiple effects via monad transformers.

### 6.1 Error handling

We begin by extending the  $WT$ -monad with exceptions to create the  $WET$ -monad (where the  $E$  stands for errors)—see Figure 18. The type  $WET o \ell a$  is a monad stack which consists on the  $(T \ell)$ -monad at the bottom, the monad transformer  $ErrorT String$  in the middle (which extend the  $(T \ell)$ -monad with error messages of type  $String$ ), and  $WriterT (Outputs o)$  at the top—Figure 19 depicts the monadic stack.



**Figure 19: The monadic stack**

The interface for  $WET$  is similar to that of  $WT$ , i.e., it provides the primitive  $tell^{WT}$ . It also provides the functionality for throwing ( $throwError$ ) and catching ( $catchError$ ) errors. The key difference between the interface of  $WT$  and  $WET$  lies in the implementations of  $cast^{WT}$  and  $cast^{WET}$ . The derived operation  $cast^{WET}$  requires checking if the sensitive computation threw an error—an error that must be kept secret ( $T \ell_H (Either Error a)$ ). This occurs because the monad family  $T$  keeps us from communicating errors from a sensitive ( $\ell_H$ ) to a less sensitive ( $\ell_L$ ) context. Since the error might occur within the  $(WET o \ell_H)$ -monad, the error must then remain protected by  $\ell_H$ . This restriction is in agreement with how other IFC libraries deal with exceptions [18, 26]. Differently from such libraries, we derive how to securely deal with errors by simply using  $T \ell$  and the DCC typing rules.

```

-- Definition
type ST s ℓ a = StateT (State s) (T ℓ) a
-- State
get :: ST s ℓ (State s)
put :: State s → ST s ℓ a
-- Avoiding label creep
castST :: (Project ℓH s,
           (ℓH ≤ StateAbove ℓH s) ~ True)
         ⇒ ST s ℓH a → ST s ℓL (T ℓH a)

```

**Figure 20: Secure stateful computations.**

```

-- Definition
type SET s ℓ a = StateT s (EitherT String (T ℓ)) a
-- State
get :: SET s ℓ (State s)
put :: State s → SET s ℓ a
-- Label creep
castSET :: (Project ℓH,
            (ℓH ≤ StateAbove ℓH s) ~ True)
          ⇒ SET s ℓH a → SET s ℓL (T ℓH (Either String a))

```

**Figure 21: Combining state and error handling.**

## 6.2 State

Another effect which can be implemented on top of DCC is state. In order to indicate that the focus is now on state, we change the type variable  $o$  denoting “outputs” for type variable  $s$  denoting “state”. We also introduce type synonyms for *Segments*  $s$  and *SegmentsAbove*  $ℓ s$  as *State*  $s$  and *StateAbove*  $ℓ s$ , respectively.

Figure 20 shows the definition and interface of stateful computations as the  $(ST\ s\ ℓ)$ -monad. It is simply the application of the stateful  $(StateT)$  monad transformer over *State*  $s$  and  $T\ ℓ$ . The interface presents specific primitives to fetch (*get*) and store state (*put*). The type signature for  $cast^{ST}$  is exactly the same as  $cast^{WT}$  except for mentioning  $ST$  and *StateAbove* instead of  $WT$  and *OutputsAbove*, respectively—Appendix A.4 shows the complete, *derived*, implementation of  $cast^{ST}$ .

As the  $WT$  monad can be extended to handle errors, so it can be the  $ST$  monad. Figure 21 shows the interface for the  $(SET\ s\ ℓ)$ -monad, a monad with state and error handling. The idea behind  $cast^{SET}$  is the same as  $cast^{WET}$ , catch the error and put it in the result, adjusting the effect accordingly (depending on *EitherT* being on the outside or inside of *StateT*).

## 6.3 I/O

We explore if DCC is capable of securing effects provided by one of the most popular monads in Haskell, i.e., the  $IO$  monad. This monad

provides the API to perform effects via the underlying I/O system, e.g., the file system, network communication, running executables, etc. The  $IO$  monad, different from the ones we have seen so far, is opaque and lacks a manner to “get out of it” without possibly breaking type-safety [29]—i.e., there is no safe function of type  $IO\ a \rightarrow a$ . Once a computation has type  $IO$ , it will remain there.

We argue that DCC can only secure effects which can be represented in a pure way (like state, exceptions, etc.), but not the ones provided by the (opaque)  $IO$  monad. To illustrate this point, let us consider  $IO$ -effects conceived inside the  $(T\ H)$ -monad, i.e., a computation of type  $T\ H\ (IO\ a)$ . Are such  $IO$ -effects safe to perform? The answer is no, we have *no guarantee* how such effects were constructed or what they do! If those effects were internal to the program, like state or exceptions, they could potentially be managed as described before, i.e., building a monadic stack where the  $T\ ℓ$ -monad is at the bottom. However,  $IO$ -effects could also involve external communication or any other type of effects. For instance, they could observe sensitive information and send it over the network.

One could be tempted to conceive a primitive of the form

```
secureIO :: T H (IO a) → IO (T H a)
```

where the  $IO$ -actions are being taken out of the  $T\ H$ -monad—after all, they cannot be observed (the  $IO$  monad is opaque). Unfortunately, *secureIO* does not honor to its name. Consider the following leaking program.

```

attack :: T H Bool → T H (IO Bool)
attack tb =
  tb >>= λb → do if b then putStrLn "True"
                else putStrLn "False"
                returnH (return b)

leak :: T H Bool → IO (T H Bool)
leak = secureIO ∘ attack

```

This code always returns the same boolean wrapped in the  $IO$ -monad ( $return\ b$ ), while it performs two different effects depending on the value of the secret boolean  $b$ —i.e., either showing “True” or “False” on the screen. Consequently, when applying *secureIO* and taking the  $IO$ -effects “out” of  $T\ H$ , the attacker can execute them and infer the value of  $b$  from reading the console—which is a public channel. This argument has been previously brandished to justify the introduction of another security monad capable to handle  $IO$ -effects [18, 19, 26]. In this work, however, we restrict ourselves to only having the  $T$ -monad family.

Alternatively, it could be argue that there could be a pure representation of the  $IO$ -effects rather than the  $IO$ -monad itself [27]. Having that, it is just a matter of placing  $(T\ ℓ)$ -monads in the pure representation of actions according to the sensitivity of inputs and outputs. We declare this intriguing idea as future work.

## 7 RELATED WORK

DCC. The seminal work on DCC [1] has inspired a large amount of work. Tse and Zdancewic translated DCC into System F [30].

They conceived the rules for  $()$  and use parametricity to prove non-interference. Authors also provided an implementation of DCC in Haskell [30]. However, this implementation differs from ours in two main aspects. First, the  $T \ell a$  monad is isomorphic to the type  $\ell \rightarrow a$  rather than an identity monad as originally conceived in DCC [1]—this representation is inspired by the translation of DCC into System F. Secondly, the super-bind operator is implemented as an overloaded primitive (type-class). As a result, it is necessary to give an implementation of super-bind *for each possible type constructor* present in DCC. Instead, we provide a single implementation of super-bind and provide a type-level function to enforce the corresponding security checks—which aligns with having a single typing rule for super-bind in DCC. Shikuma and Igarashi gave a counterexample to show that Tse and Zdancewic translation was not sound [23, 24]. Instead, Shikuma and Igarashi propose a calculus very much inspired by DCC, called *sealing calculus*. To make their translation proof work, they weaken the source language, i.e., admit some additional well-typed terms compared with DCC, at the same time strengthening the target language to be a restricted version of System F. Recently, Bowman and Ahmed show a sound translation from DCC into System F by the use of existential types [4]—thus being the first ones to prove non-interference by encoding it in parametricity. As in [4, 23, 24], our work focuses on the terminating fragment of DCC. We carefully craft an implementation which looks as a syntactic translation of DCC rules. By doing so, we can rely on the results of [4] to give us confidence about the correctness of our code, where parametricity lies behind polymorphic languages like Haskell.

*Security libraries.* Many monadic IFC security libraries, e.g., LIO [26], HLIO [6], and MAC [18], provide security monads which *intertwine pure and effectful computations*—this is one of the main differences with our work. Instead, we describe a methodology for using DCC as a foundational approach to secure effects mimicked by pure computations via the use of monad transformers. As a result, and contrary to existing IFC libraries, we are able to *derive* (by complying with the type system) a secure API to safely support exceptions, references, and any other effects mimicked by pure computations. On the other hand, IFC libraries (e.g., LIO) often come accompanied with proof showing that programs satisfy the non-interference security policy. We argue that providing soundness proofs about our encoding would end up in essentially the same semantics/proofs as in the original DCC paper due to the nature of our implementation, i.e., a syntactic translation. (We note that stacks of monad transformers can be modeled in DCC by leveraging sum- and product-types.) This work also identifies DCC’s limits: handling arbitrary I/O effects (see Section 6.3). In that situation, it is required to introduce security-specific monads much along the lines of what SecLib, LIO, and MAC do. Rajani et al. [17] study the trade-offs and equivalences between statically tracking dependencies using monads with more fine-grained traditional IFC type systems. Devriese and Piessens provide a monad transformer to extend imperative-like APIs with support for IFC in Haskell [9]. Similar to our work, they use type families to encode the security lattice; however, they provide *open* definitions for type-level functions. In principle, it would be possible to rephrase some of their static analyzes into our encoding of DCC. Jaskelioff

and Russo implements a library which dynamically enforces IFC using secure multi-execution (SME) [13]—a technique that runs programs multiple times (once per security level) and varies the semantics of inputs and outputs to protect confidentiality. Rather than running multiple copies of a program, Schmitz et al. provide a library with *faceted values* [21], where values present different behavior according to the privilege of the observer. These last two libraries apply different ideas than those behind DCC: it enforces IFC dynamically by repeating computations rather than tracking dependencies.

## 8 FINAL REMARKS

By leveraging closed type families and type classes, we describe a novel implementation of DCC in Haskell which looks like a direct translation of DCC’s typing rules—an elegant outcome of this work. We clarify some aspects of DCC’s non-standard bind operation from the perspective of a Haskell programmer. Furthermore, we also provide a type directed approach (the *cast* functions) to avoid the label creep problem present in DCC. Once embedded in Haskell, DCC can be used as a solid base to secure programs with effects like outputs, state, exceptions, or combinations thereof. For that, we utilize monad transformers, well known tools for combining effects. Most importantly, to secure such effects, we *only* need to trust our type system (via the use of SafeHaskell [28]) and the implementation of super-bind. It becomes now possible to start envisioning the construction of larger case studies and applications in Haskell, where DCC is the underlying security mechanism—an interesting direction for future work.

In a recent paper, Austin *et al.* [2] study how to secure imperative languages constructs by translating them into a pure functional language. While the paper focuses on dynamic techniques, the articles states “... *one can translate various static information-flow systems for imperative languages to the Dependency Core Calculus (DCC) [1] ...*” However, the only translation in [1] is an IFC system for a imperative language with state. This work completes that statement by showing how to handle a richer set of effects—by simply embedding DCC into Haskell and letting the type system guiding us.

*Acknowledgments.* We would like to thank Deepak Garg, Vineet Rajani, Marco Vassena, and the anonymous reviewers for their comments of this work. This work was funded in part by the Swedish research agency VR and in part by the GRACeFUL (grant agreement No 640954) project.

## REFERENCES

- [1] M. Abadi, A. Banerjee, N. Heintze, and J. Riecke. 1999. A Core Calculus of Dependency. In *Proc. ACM Symp. on Principles of Programming Languages*. 147–160.
- [2] Thomas H. Austin, Cormac Flanagan, and Martın Abadi. 2012. A Functional View of Imperative Information Flow. *Lecture Notes in Computer Science* (2012), 344–349.
- [3] David E. Bell and L. La Padula. 1976. *Secure Computer System: Unified Exposition and Multics Interpretation*. Technical Report MTR-2997, Rev. 1. MITRE Corporation, Bedford, MA.
- [4] William J. Bowman and Amal Ahmed. 2015. Noninterference for free. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015, Vancouver, BC, Canada, September 1-3, 2015*.
- [5] Jan Bracker and Henrik Nilsson. 2016. Supermonads: one notion to bind them all. In *Proceedings of the 9th International Symposium on Haskell*. ACM, 158–169.

[6] P. Buiras, D. Vytiniotis, and A. Russo. 2015. HLIO: Mixing Static and Dynamic Typing for Information-Flow Control in Haskell. In *Proc. of the ACM SIGPLAN International Conference on Functional Programming (ICFP '15)*. ACM.

[7] Pablo Buiras, Dimitrios Vytiniotis, and Alejandro Russo. 2015. HLIO: Mixing static and dynamic typing for information-flow control in Haskell. In *ACM SIGPLAN Notices*, Vol. 50. ACM, 289–301.

[8] D. E. Denning and P. J. Denning. 1977. Certification of Programs for Secure Information Flow. *Commun. ACM* 20, 7 (July 1977), 504–513.

[9] D. Devriese and F. Piessens. 2011. Information flow enforcement in monadic libraries. In *Proc. of the ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI '11)*. ACM.

[10] R. A. Eisenberg, D. Vytiniotis, S. Peyton Jones, and S. Weirich. 2014. Closed Type Families with Overlapping Equations. In *Proc. of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '14)*. ACM.

[11] J.A. Goguen and J. Meseguer. 1982. Security policies and security models. In *Proc of IEEE Symposium on Security and Privacy*. IEEE Computer Society.

[12] D. Hedin and A. Sabelfeld. 2011. A Perspective on Information-Flow Control. In *Proc. of the 2011 Marktoberdorf Summer School*. IOS Press.

[13] M. Jaskelioff and A. Russo. 2011. Secure multi-execution in Haskell. In *Proc. Andrei Ershov International Conference on Perspectives of System Informatics (LNCS)*. Springer-Verlag.

[14] B. W. Lampson. 1973. A Note on the Confinement Problem. *Commun. ACM* 16, 10 (Oct. 1973).

[15] P. Li and S. Zdancewic. 2006. Encoding Information Flow in Haskell. In *Proc. of the IEEE Workshop on Computer Security Foundations*. IEEE Computer Society.

[16] Sheng Liang, Paul Hudak, and Mark Jones. 1995. Monad Transformers and Modular Interpreters. In *In Proc. of the 22nd ACM Symposium on Principles of Programming Languages*. ACM Press.

[17] Vineet Rajani, Iulia Bastys, Willard Rafnsson, and Deepak Garg. 2017. Type systems for information flow control: the question of granularity. *SIGLOG News* 4, 1 (2017).

[18] A. Russo. 2015. Functional Pearl: Two Can Keep a Secret, if One of Them Uses Haskell. In *Proc. of the ACM SIGPLAN International Conference on Functional Programming (ICFP 2015)*. ACM.

[19] A. Russo, K. Claessen, and J. Hughes. 2008. A library for light-weight information-flow security in Haskell. In *Proc. ACM SIGPLAN symposium on Haskell (HASKELL '08)*. ACM.

[20] A. Sabelfeld and A. C. Myers. 2003. Language-Based Information-Flow Security. *IEEE J. Selected Areas in Communications* 21, 1 (Jan. 2003), 5–19.

[21] Thomas Schmitz, Dustin Rhodes, Thomas H. Austin, Kenneth Knowles, and Cormac Flanagan. 2016. Faceted Dynamic Information Flow via Control and Data Monads. In *Principles of Security and Trust - 5th International Conference, POST 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings*.

[22] Alejandro Serrano and Jurriaan Hage. 2016. *Type Error Diagnosis for Embedded DSLs by Two-Stage Specialized Type Rules*. Springer Berlin Heidelberg, Berlin, Heidelberg, 672–698. [https://doi.org/10.1007/978-3-662-49498-1\\_26](https://doi.org/10.1007/978-3-662-49498-1_26)

[23] Naokata Shikuma and Atsushi Igarashi. 2006. Proving Noninterference by a Fully Complete Translation to the Simply Typed lambda-Calculus. In *Advances in Computer Science - ASIAN 2006. Secure Software and Related Issues, 11th Asian Computing Science Conference, Tokyo, Japan, December 6-8, 2006, Revised Selected Papers*.

[24] Naokata Shikuma and Atsushi Igarashi. 2008. Proving Noninterference by a Fully Complete Translation to the Simply Typed Lambda-Calculus. *Logical Methods in Computer Science* 4, 3 (2008).

[25] D. Stefan, A. Russo, D. Mazières, and J. C. Mitchell. 2011. Disjunction Category Labels. In *Proc. of the Nordic Conference on Information Security Technology for Applications (NORDSEC '11)*. Springer-Verlag.

[26] D. Stefan, A. Russo, J. C. Mitchell, and D. Mazières. 2011. Flexible Dynamic Information Flow Control in Haskell. In *Proc. of the ACM SIGPLAN Haskell symposium (HASKELL '11)*.

[27] Wouter Swierstra and Thorsten Altenkirch. 2007. Beauty in the Beast: A Functional Semantics of the Awkward Squad. In *Haskell '07: Proceedings of the ACM SIGPLAN Workshop on Haskell*. 25–36.

[28] David Terei, Simon Marlow, Simon Peyton Jones, and David Mazières. 2012. Safe Haskell. In *ACM SIGPLAN Notices*, Vol. 47. ACM, 137–148.

[29] D. Terei, S. Marlow, S. Peyton Jones, and D. Mazières. 2012. Safe Haskell. In *Proc. of the ACM SIGPLAN Haskell symposium (HASKELL '11)*. ACM.

[30] S. Tse and S. Zdancewic. 2004. Translating dependency into parametricity. In *Proc. of the Ninth ACM SIGPLAN International Conference on Functional Programming*. ACM.

[31] Stephen Tse and Steve Zdancewic. 2004. Translating Dependency into Parametricity. *SIGPLAN Not.* 39, 9 (Sept. 2004), 115–125.

[32] D. Volpano, G. Smith, and C. Irvine. 1996. A Sound Type System for Secure Flow Analysis. *J. Computer Security* 4, 3 (1996), 167–187.

[33] Philip Wadler. 1990. Comprehending Monads. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming (LFP '90)*. ACM, 61–78.

[34] Philip Wadler. 1995. Monads for functional programming. In *International School on Advanced Functional Programming*. Springer, 24–52.

[35] Brent A. Yorgey, Stephanie Weirich, Julien Cretin, Simon Peyton Jones, Dimitrios Vytiniotis, and José Pedro Magalhães. 2012. Giving Haskell a Promotion. In *Proc. of the 8th ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI '12)*. ACM.

## A APPENDIX

### A.1 Automatic relabeling

In this Appendix, we provide the rest of the instances for the type class *Upgrade*: pairs and functions. As the following code shows, these instances simply work by inductively traversing the type being upgraded.

```
instance (Upgrade s s', Upgrade t t')
  => Upgrade (s, t) (s', t') where
  up (s, t) = (up s, up t)
```

```
instance Upgrade s s'
  => Upgrade (t -> s) (t -> s') where
  up f = up o f
```

Upgrading a pair (*up (s, t)*) is the same thing as upgrading both elements (i.e., *up s* and *up t*). Upgrading a function (*up f*) is the same as upgrading the result of the function (*up o f*).

### A.2 Structure for outputs

As explain in the main text of this work, we consider outputs of type *o* as having a monoid structure. Consequently, we consider every output partition *T ℓ o* as a monoid too since we need both neutral elements, denoted by *mempty*, and a way to combine *protected* outputs from different computation—denoted by the primitive *mappend*.

```
instance Monoid o
  => Monoid (T ℓ o) where
  mempty      = return mempty
  mappend m0 m1 = do
    o0 <- m0
    o1 <- m1
    return (mappend o0 o1)
```

The neutral element is simply the *protected* computations which returns the neutral element of type *o*. Combining, i.e., “*mappend*ing”, two *protected* computations which produce *os* is done by running the first computation *m0* and then the second *m1* in order to return their *protected* combined results (*return (mappend o0 o1)*).

### A.3 Label creep

In this section, we provide the implementation of *cast<sup>WT</sup>*. As described before, the key observation for this primitive is that it requires getting rid of the low output created by running the high computation being casted while retaining the high part of the output. Specifically, the code for *cast<sup>WT</sup>* is as follows.

```
castWT m = do
  let (highOutputs, tlx) =
      runWriterT m ≫ λ(x, outputs) → (project outputs, return x)
  tell (merge highOutputs mempty)
  return tlx
```

First, it runs the sensitive computation  $m$  ( $\text{runWriterT } m$ ). By using super-bind, it extracts the result ( $x$ ) and produced outputs ( $\text{outputs}$ ) by  $m$ —observe the function  $\lambda(x, \text{outputs}) \rightarrow \dots$ ). Then,  $\text{cast}^{\text{WT}}$  only keeps the sensitive part of the outputs ( $\text{project outputs}$ ) and the result produced by  $m$  ( $\text{return } x$ ). Consequently, the pair ( $\text{highOutputs}, \text{tl}_x$ ) respectively holds the sensitive outputs and the sensitive result produced by  $m$ . The next step by  $\text{cast}^{\text{WT}}$  fills the tuple of empty segments of outputs ( $\text{mempty}$ ) with the sensitive part produced by  $m$  by running  $\text{tell } (\text{merge highOutputs mempty})$ . At this point, the produced ( $\text{WT } o \ell_{\text{L}}$ )-monad computation contains as sensitive output  $\text{highOutputs}$  while its public segment is empty—nevertheless, other ( $\text{WT } o \ell_{\text{L}}$ )-monadic computations could populate such segment as shown by Example 5.4.

#### A.4 State

When considering state, the implementation of  $\text{cast}^{\text{ST}}$  requires keeping track of what the secret computation has done to the state. While it is different from  $\text{cast}^{\text{WT}}$ , the typing rules of the  $T$  monad family keep us from leaking secrets! Below is the implementation of  $\text{cast}^{\text{ST}}$ :

```

castST m = do
  s ← get
  let (sH, tlx) =
      runStT m s ≫≧ λ(x, st) → (project st, return x)
  put (merge sH s)
  return tlx

```

The implementation of works by the following steps:

- Obtain the current state ( $s \leftarrow \text{get}$ ).
- Run the sensitive computation with the current state as starting state ( $\text{runStT } m \ s$ ) and obtain the segments of the resulting state which are protected at  $\ell_{\text{H}}$ .
- Return the result ( $x$ ) of the computation wrapped in  $T \ell_{\text{H}}$ .
- Merge the new, sensitive, state with the old public state—observe the instruction  $\text{put } (\text{merge } s_{\text{H}} \ s)$ .
- Return the result protected by  $T \ell_{\text{H}}$  ( $\text{return } \text{tl}_x$ ).

Similar as  $\text{cast}^{\text{WT}}$ , note the use of super-bind to obtain the value produced from the sensitive computation  $m$  into two parts—observe function ( $\lambda(x, st) \rightarrow \dots$ ). In this case, super-bind respectively produces the result of the sensitive computation ( $\text{return } x$ ) and the sensitive part of the state ( $\text{project } st$ ) which will be subsequently merged back into the state ( $\text{put } (\text{merge } s_{\text{H}} \ s)$ ).