

# THE UNIVERSITY OF NEW SOUTH WALES



SYDNEY • AUSTRALIA

School of Computer Science and Engineering  
Faculty of Engineering

## FLOGENT

AN INFORMATION FLOW SECURITY FEATURE FOR COGENT

A Thesis (Part C) Report By

**Vivian Ye-Ting Dang**

Supervised By

**Christine Rizkallah**

June 27, 2020

# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
<b>2</b>	<b>Background</b>	<b>7</b>
2.1	Compromises in MAC . . . . .	9
<b>3</b>	<b>Flogent: Implementation</b>	<b>11</b>
3.1	Case Study: Password Confidentiality . . . . .	11
3.1.1	Implementation using MAC . . . . .	12
3.2	Design . . . . .	13
3.2.1	Using a "Key" . . . . .	14
3.2.2	World Labelling . . . . .	16
3.3	Implementation . . . . .	18
3.3.1	Typing Rules . . . . .	20
3.3.2	Constraint Generation and Solving . . . . .	21
<b>4</b>	<b>Testing</b>	<b>22</b>
4.1	Confidentiality . . . . .	22
4.2	Integrity . . . . .	23
4.3	Test result . . . . .	24
<b>5</b>	<b>Formalisation</b>	<b>25</b>
5.1	Noninterference . . . . .	25
5.1.1	Definition of the Erasure Function . . . . .	26
5.1.2	Small-step Semantics . . . . .	28
5.2	Applying to Flogent . . . . .	32
5.3	Summary . . . . .	32

# List of Figures

1.1	Identifying memory-safety-related CVEs (from Miller (2019)) . . . . .	5
2.1	Core API for MAC (Vassena et al., 2018) . . . . .	9
3.1	Using "World" instead of Monads . . . . .	14
3.2	Proposed type signatures using "Key" . . . . .	14
3.3	Implementation using World-Labeling . . . . .	17
3.4	Summary of syntactic definitions . . . . .	19
3.5	Typing rules . . . . .	20
3.6	Constraint Generation . . . . .	21
5.1	Single-step simulation . . . . .	26
5.2	Summary of value semantics . . . . .	27
5.3	Erasure Function for expressions and values . . . . .	28
5.4	Small step semantics of the COGENT language (part 1) . . . . .	29
5.5	Small step semantics of the COGENT language (part 2) . . . . .	30
5.6	Small step semantics of new operations . . . . .	31
5.7	Single-step simulation for <b>unlock</b> . . . . .	32

# Acknowledgements

I would like to thank Christine Rizkallah, my supervisor, and Amos Robinson for their support and guidance while I am undertaking this honours thesis, I deeply appreciate the mentorship and advice from both of them.

I am extremely grateful to Liam O'Connor, for being an exceptional lecturer, tutor and friend, without whom I would not have been inspired to undertake my thesis in this field, nor would I have been able to get through all the hurdles that I have come across during my years in university.

I would like to also thank Kai Engelhardt for being a great educator who opened the door of Formal Methods and the more theoretical side of computer science to me.

Of course, I want to also thank my mum Monica and my brother Justin for their unlimited support throughout the years. Thank you to all my friends, who have been there for me at different times. I would not have been able to get through my degree without them: Alexander Rowell, Anna Azzam, Annie Liu, David Lei, Edward Thomson, Emmet Murray, Hussein Debel, Kristian Mansfield, Mahika Suri, Nethan Tran, Subramanya Vajiraya, Victor Phan and many others to whom I am very much in-debted.

# Chapter 1

## Introduction

In a world where the amount of sensitive information being put online is increasing, the risk of this information falling into the hands of malicious people also increases; this calls for more emphasis on information security within systems.

This thesis introduces FLOGENT, an information flow control system for the programming language COGENT (O'Connor et al., 2016). COGENT is a higher-order, purely functional language with *uniqueness types*. The semantics of COGENT are designed to be easy to reason about, in accordance with its aim of reducing the cost of *verification*.

Software verification aims to formally establish by proof that a program fully satisfies its expected behaviour and outcomes (i.e. its specification). Proving program correctness often involves an interactive theorem prover like *Isabelle/HOL* (Nipkow, Paulson, and Wenzel, 2002).

Given a COGENT program, the compiler generates a high-level *shallow embedding* of the program's semantics in Isabelle/HOL, a C program, and a proof that the generated C program is a refinement of the generated Isabelle/HOL embedding. This means that any functional correctness property proved about the shallow embedding also holds for the generated C program. As a result, the need to verify manually written C programs is significantly reduced, which in turn brings down the cost of verification, as C is significantly harder to reason about than COGENT.

COGENT is designed for the implementation of operating systems components such as file-systems. Vulnerabilities in operating systems can compromise the confidentiality and integrity of the entire system. As such, ensuring COGENT software is free of vulnerabilities is particularly important.

A recent presentation by Miller (2019) indicates that approximately 70% of *Com-*



**Figure 1.1:** Identifying memory-safety-related CVEs (from Miller (2019))

*mon Vulnerabilities and Exposures* (CVEs) (MITRE Corporation, 2019) that have been reported by Microsoft in the recent decade are the result of vulnerabilities related to *memory-safety* (see Figure 1.1).

COGENT is already equipped to deal with memory safety related problems. Memory safety can be divided into two main categories: temporal safety and spatial safety (Szekeres et al., 2013). Temporal safety issues involve accessing memory at the wrong time, such as trying to make use of uninitialised memory, or using memory after it has been freed. COGENT ensures temporal safety by having uniqueness types, a kind of *linear types*. Uniqueness types ensure that, at any given time, there is only one writable reference to a mutable object; this is done by requiring that variables of a linear type are used exactly once. Such a type system allows memory allocation to be tracked statically, enabling the COGENT compiler to ensure that memory is only accessed when it is safe to do so.

Spatial safety issues involve accessing memory at the wrong place. Common problems include buffer overflows and accessing arrays out-of-bounds, which creates opportunities for stack smashing. Spatial safety is ensured as a consequence of Cogent’s type safety theorem, where unsafe operations are not exposed to the users.

Problems that are unrelated to memory safety can be split into two subcategories: those that are not observable on the programming language level, including hard-

ware problems such as Spectre (Kocher et al., 2019) and Meltdown (Lipp et al., 2018), as well as side channel attacks involving timing or power signatures; and problems that are observable on the programming language level, specifically, secret data being directly accessed or modified by unauthorised processes which violates *confidentiality* and *integrity*. The latter kind of vulnerability can be overlooked when trusted programs or applications are handling sensitive information, leaving potential risks to data breaches. FLOGENT aims to reduce these types of mistakes by extending Cogent with an information flow control system.

## Chapter 2

# Background

One of the very first attempts to express *information flow control* (IFC) was made by D. E. Denning and P.J. Denning (1977), where an information flow *policy* is defined as a set of *security classes* or *labels* for information and a *flow relation* defining the permissible flows between these classes in a *lattice* structure. Each runtime object is *bound* to a security class.<sup>1</sup> Operations that use an object, say  $x$ , to derive another, say  $y$ , are only allowed if the security class of  $x$  flows into that of  $y$  according to the flow relation. For example, the classes might be 'high privilege' and 'low privilege' where 'high' privileged data cannot leak to low. Our flow relations would be high to high and low to high, thus disallowing high to low. This control over data dependency is intended to ensure “the absence of information flows from private data to public observers” (Nipkow and Klein, 2014). Information flow control is an easy and potentially automatable method to check if programs protect private data such as passwords.<sup>2</sup>

This initial concept has been applied in a variety of domains. One such domain is that of *program logics*, proof calculi for software verification. Based on separation logic (Reynolds, 2002), Costanzo and Shao (2014) present a logic that allows for manual reasoning about IFC, either by pen-and-paper proof or in an interactive theorem prover. Murray et al. (2016) combine this program logic approach with dependent types to produce a compositional verification framework. However, because COGENT generates an embedding in a purely functional subset of Isabelle/HOL, putting this embedding into a different logic would require reworking all of the existing infrastructure as well as losing the advantage of purely functional reasoning, such approaches are not well-suited to a COGENT context.

Another style of IFC is *dynamic* (i.e. run-time) security enforcement. Stefan et

---

<sup>1</sup>More recent literature calls this operation *labelling*.

<sup>2</sup>Note that this is under the assumption that the *trusted computing base* (TCB) is well-behaved.



al. (2011) describe a dynamic, language-based IFC mechanism implemented as a Haskell library called *Labeled IO* (LIO). LIO keeps track of labels by inserting them directly into run-time values and enforcing dynamic checks to ensure IFC policies are adhered to. The *Hybrid LIO* (HLIO) system of Buiras, Vytiniotis, and Russo (2015) extends this by enforcing some checks statically; however, any amount of dynamism would not be suitable in the context of COGENT. This is because COGENT is entirely statically typed and is used for low-level operating system components, where the performance overheads introduced by dynamic checking are not acceptable.

Volpano, Irvine, and Smith (1996) re-formulate this idea of IFC into a *type system* that enforces security properties, along with a proof of soundness. A simplified version of their security type system is approachably introduced by Nipkow and Klein (2014). Nipkow and Klein (2014) define security levels using natural numbers, whereas Volpano, Irvine, and Smith (1996) make use of a more general lattice structure. The soundness result ensures *noninterference* from high security level to low security level, i.e. high does not interfere with low, information only flows from low to high. Here, a judgement  $\ell \vdash c$  denotes that the program  $c$  contains no information flows to variables lower than level  $\ell$  and the only flows would be from variables greater than or equal to  $\ell$ .

Abadi et al. (1999) generalise security type systems to accommodate a functional language with higher order functions. They propose a general calculus, specifically a  $\lambda$ -calculus for type-based dependency analyses called *Dependency Core Calculus* (DCC). They give DCC a denotational semantics that formalises the idea of noninterference. DCC incorporates a family of monads (Moggi, 1991), indexed by a security level, to express the different levels under which a program executes.

This calculus for static IFC was introduced by Russo (2015) and is adapted into a HASKELL (Haskell, 2010) library called MAC by Vassena et al. (2018). MAC places labels on the type level, and defines a HASKELL monad indexed by these type-level labels. In MAC, *labelled values*—values that are associated with a security level—are an explicit type, written *Labeled*  $\ell$   $a$ . This denotes a value of type  $a$  labelled at security class  $\ell$ .

Figure 2.1 outlines the type signatures for the core operations on labelled values in MAC. The library makes use of the *MAC* monad, which associates the computation environment with a label. To create a labelled value, the join operation is used, which transforms a computation of level  $\ell_H$  to a computation of lower level  $\ell_L$ , but the result is labelled with the higher level  $\ell_H$ . This means that the result of the computation cannot be accessed by those who do not have authorisation of level  $\ell_H$ . They also define a simpler constructor for labelled values, *label*, which can be defined in terms of join as follows:

$$\begin{aligned} \text{join} &:: \ell_L \leq \ell_H \Rightarrow \text{MAC } \ell_H \ a \rightarrow \text{MAC } \ell_L \ (\text{Labeled } \ell_H \ a) \\ \text{unlabel} &:: \ell_L \leq \ell_H \Rightarrow \text{Labeled } \ell_L \ a \rightarrow \text{MAC } \ell_H \ a \end{aligned}$$

**Figure 2.1:** Core API for MAC (Vassena et al., 2018)

$$\begin{aligned} \text{label} &:: (\ell_L \leq \ell_H) \Rightarrow a \rightarrow \text{MAC } \ell_L \ (\text{Labeled } \ell_H \ a) \\ \text{label} &= \text{join} \circ \text{return} \end{aligned}$$

The second operation defined by MAC is `unlabel`. To unlabel a value of level  $\ell_L$ , the computation level  $\ell_H$  must be at least as high as  $\ell_L$ .

As `HASKELL` and `COGENT` are both functional languages with static types, the MAC approach is the most directly applicable to `COGENT`. However, the use of monads is not appropriate in a `COGENT` context, because monadic programming requires closures, which `COGENT` does not support (O’Connor et al., 2016). Since the use of monads is vital to the MAC approach, this presents a challenge for us, to implement the same functionality in `COGENT`.

There are other statically typed functional languages, such as `ML` (Milner, Tofte, and MacQueen, 1997) and `Mercury` (Somogyi, Henderson, and Conway, 1995), where monadic programming is also not common. Pottier and Simonet (2003) present an implementation of Abadi’s framework for `ML`, suggesting that Abadi’s framework can be applied without the use of monads.

## 2.1 Compromises in MAC

Any practical implementation of IFC will come with some caveats not present in the calculus of Abadi et al. (1999). For example, in MAC, a number of compromises had to be made in order to implement IFC in Haskell (Russo, 2015):

**unsafePerformIO** Haskell includes a built-in unsafe function, `unsafePerformIO`, which allows the user to perform arbitrary side-effects but hide them from Haskell’s type system. In MAC, they must assume that any use of `unsafePerformIO` does not violate security properties. For `COGENT`, a similar assumption must be made about abstract functions imported from C.

**Exceptions** Haskell includes exceptions that can be thrown from pure functions. This means that it is possible for an attacker to learn secret information by observing which inputs cause an exception to be thrown. MAC includes some support for exceptions, but there are still scenarios where exceptions can lead to vulnerabilities (Hritcu et al., 2013). In `COGENT`, however, all

functions are total and exceptions are not permitted. Therefore, this is not a problem faced by FLOGENT.

**Non-termination** Similarly to exceptions, Haskell functions may fail to terminate. This also opens a potential side-channel through which secret information may be leaked. COGENT ensures termination of all functions which limits the problem of non-termination to only the use of abstract functions imported from C.

Because of COGENT's restrictive design, fewer compromises have to be made in the implementation of FLOGENT. This means that our type system can be more informative and better at detecting vulnerabilities, however care must be taken to ensure that the system remains expressive enough for practical use-cases.

## Chapter 3

# Flogent: Implementation

As we have established in previous chapters, FLOGENT aims to extend COGENT with an information flow control static type system. This includes design, compiler implementation, and formalisation. To ensure that the security properties hold (i.e. confidentiality and integrity) for all well-typed FLOGENT programs, these properties must also be characterised on top of our formalisation.

The main contribution so far in Thesis B is the design and implementation of two primitive operations within MINIGENT, a miniature version of COGENT. We call the operations `join` and `unlock`, derived from the `join` and `unlabel` operations of Abadi et al. (1999). These are sufficient to ensure information flow control in programming language libraries such as MAC (Vassena et al., 2018). Details of these operations are explained later in this chapter.

The calculus of Abadi et al. (1999) is widely implemented in functional languages (Vassena et al., 2018; Pottier and Simonet, 2003). It is foundational work in the field of information flow security and, according to experts in the field (Murray, 2019), this is a good basis for our work in COGENT. As mentioned in Chapter 2, MAC is a direct HASKELL implementation of this calculus. This adaptation is particularly useful because both HASKELL and COGENT are purely functional languages with a static type system.

### 3.1 Case Study: Password Confidentiality

To illustrate the design process and the usage of the two operations, we will examine a case study, first presented by Russo (2015). The model has been slightly modified for the example to be more relevant to this project.

We have Alice, a programmer who wants to write a password manager that asks the user for passwords. She wants to prevent users from using common passwords, so she wishes to include a feature that will tell the user whether or not the provided password is found in a list of common passwords. A colleague of hers, Bob, has already implemented such functionality in a different project and Alice wants to make use of that function in her code. Observe that Bob will need access to the password in order to provide this functionality—this breaks confidentiality as sensitive data is leaked. Therefore, we want to make sure that the password is somehow passed to Bob for the check without Bob being able to act maliciously, i.e. using this information for any other (potentially nefarious) purpose.

### 3.1.1 Implementation using MAC

Alice
<pre> password :: IO String password =   do putStr "Please, select your password:"      pass ← getLine      lbool ← runMAC \$ do        lpass ← label pass :: MAC L (Labeled H String)        Bob.commonPass lpass      let b = unRes lbool      if then putStrLn "Your password is too common!" &gt;&gt; password      else return pass </pre>

Recall in Figure 2.1 that the operations require the two arbitrary security levels  $\ell_L$  and  $\ell_H$  to satisfy the condition of  $\ell_L \leq \ell_H$ . For this example, we need only a two-element security lattice, so we instantiate  $\ell_L$  to a concrete low level  $L$  and  $\ell_H$  to a concrete high level  $H$ . First, Alice asks the user for a password. Once the password is obtained, it is then labelled at  $H$  using the label function, before passing to Bob, who only has access to operations of level  $L$ . The function *unRes*, which unsafely removes a label, is only permitted in functions that are part of the trusted computing base.

The code for Bob is presented below:

```
Bob  
  
commonPass :: Labeled H String → MAC L (Labeled H Bool)  
commonPass lpass =  
  do str ← wgetMAC "http://common-password.list"  
    let passwords = filter (not ∘ null) (lines str)  
    join $ do pass ← unlabel lpass  
        return $ isJust $ find (== pass) passwords
```


Here, Bob first fetches the list of common passwords via an HTTP request (the *wgetMAC* function). The password given to Bob is labelled at level *H*, so Bob can only perform the password comparison in a computation of level *H*. Bob uses *join* to transform this comparison computation to one of level *L*, however the resultant boolean value is now labelled at level *H*. This prevents Bob from leaking information, because during the password comparison computation, he would be unable to make outside requests of level *L*; and outside the password comparison computation, he would be unable to access the inputs or outputs of this secure computation.

## 3.2 Design

We want to design FLOGENT to have the same information flow control capabilities that MAC demonstrates in the example above. Despite the similarities between COGENT and HASKELL, there are two main differences: The first is the lack of support for monadic programming in COGENT. This necessitates a modification of MAC's approach to implement security-labelling in FLOGENT.

The second main difference is the uniqueness type system in COGENT, which is not a feature that is present in HASKELL or any widely-used security type system. As we will see, our final design for FLOGENT depends crucially on this feature of COGENT.

HASKELL makes use of monads primarily as a way to model side effects while maintaining a purely functional semantics. In COGENT, however, the uniqueness type system is used for this purpose. Specifically, effectful functions are given a type that takes a (model of the) world as input and produces a world as output (O'Connor et al., 2016). Figure 3.1 describes this translation, where the type of the world is denoted as  $\mathbb{W}$ .



**Figure 3.1:** Using "World" instead of Monads

### 3.2.1 Using a "Key"

An intuitive way to represent authorisation for operations of a certain security level is to have a special type that acts like a *key*. This can be thought of as a witness of authorisation—analogue to possessing a key to a room to signify authorisation to access that room. Thus, the unlabel operation of Abadi et al. (1999) becomes *unlock* (*mutatis mutandis* for *Labeled*  $\ell$  and *Locked* $_{\ell}$ ).

Proposed type signatures based on this idea for the two primitive operations *join* and *unlock* are shown in Figure 3.2. Similarly to MAC, this design also has a *lock* function (cf. *label*) as follows:


$$\begin{aligned} \text{lock} &:: a \rightarrow (\text{Locked}_{\ell} a) \\ \text{lock } a &= \text{join } (\lambda x \rightarrow a) \end{aligned}$$

Note that the high level computation in MAC is now represented as a function that accepts a *Key* type of the given level.

To return to our case study, with this design Alice's code would be as follows<sup>1</sup>:

---

<sup>1</sup>For the purposes of illustration, we use a hypothetical uniqueness-typed language with syntax superficially resembling Haskell.



**Figure 3.2:** Proposed type signatures using "Key"

## Alice

```
password :: World → (World, Password)
password w0 = let w1 = putString "pick a password" w0
                (w2, pass) = getLine w1
                lockedPass = lock pass
                (w3, b) = Bob.commonPassword (w2, lockedPass)
            in if b then
                let w4 = putString "your password is weak" w3
                in password w4
            else
                (w3, pass)
```

Note that, unlike the MAC example, no monads are used, and instead a world value (World) is passed to every effectful function. Bob's code would be as follows:

## Bob

```
commonPassword :: (World, LockedH String) → (World, LockedH Bool)
commonPassword (w0, lockedS) =
    let (w1, str) = wget "http://common-password.list" w0
        passwords = filter (not ∘ null) (lines str)
    in (w1, join (λ k → let s = unlock k lockedS
                        in elem s passwords))
```

Here, Bob uses the join operation to gain a key to the high level locked string, which he can then use to check if the password is in the list.

## Confidentiality Loophole

This design ensures that high level Locked values are not accessed without authorisation, but it does not ensure that Bob does not communicate over insecure channels while authorised to access secure values.

For example, if Bob wished to leak passwords to his own server, he could make a malicious call to *wget* inside the join computation as follows:



**Bob (Evil)**

```
commonPassword (w0, lockedS) =  
  let (w1, str) = wget "http://common-password.list" w0  
      passwords = filter (not ∘ null) (lines str)  
  in join (λ k → let s = unlock k lockedS  
              (w2, _) = wget ("http://bob.evil/" ++ s) w1  
          in (w2, elem s passwords))
```

This means that this type system does not prevent low-level computations from being executed in a high-level context. There is one caveat, however: The type signature assigned to this computation will make the world inaccessible from a low-level context:

$$\text{commonPassword} :: (\clubsuit, \text{Locked}_H \text{ String}) \rightarrow \text{Locked}_H (\clubsuit, \text{Bool})$$

Such a type signature would be useless in most contexts, as such a locked up world would make it impossible to implement a top-level main function, which would be of type  $\clubsuit \rightarrow \clubsuit$ . This means that the security of any part of the program, however, depends on the type of the main function, which is non-compositional and potentially difficult to verify. Another drawback of this design is that the type errors presented to the user in these insecure scenarios are not located at the point of vulnerability. Rather, they are located at some point further up the call chain. In addition to making debugging more difficult, it can also mask vulnerabilities inside trusted functions. For example, suppose a trusted function  $f$  calls a function  $g$ , and the function  $g$  is insecure. This could lead to the function  $f$  having an insecure type, without explicitly pointing to  $g$  as the source of the vulnerability. Thus, because the function  $f$  is trusted, vulnerabilities in  $g$  could remain hidden from the programmer.

### 3.2.2 World Labelling

Rather than have a separate type to denote authorisation, we could instead tag the world itself with a security level. Here, the type signatures, shown in Figure 3.3, require each operation to take a world and return a world. This makes use of the uniqueness typing applied to  $\clubsuit$  values: now, when the user calls  $\text{join}$ , they must provide their low-level  $\clubsuit$  value, which, having been used, cannot be used again. This means that the only  $\clubsuit$  value accessible to the user inside the  $\text{join}$  is the high-level one provided to the computation. This prevents the situation seen earlier, where Bob could invoke low-level computations from inside  $\text{join}$ . Furthermore,

$\begin{aligned} \text{join} &:: \ell_L \leq \ell_H \Rightarrow \bullet_{\ell_L} \rightarrow (\bullet_{\ell_H} \rightarrow (\bullet_{\ell_H}, \tau)) \rightarrow (\bullet_{\ell_L}, \text{Locked } \ell_H \tau) \\ \text{unlock} &:: \ell_L \leq \ell_H \Rightarrow \text{Locked } \ell_L \tau \rightarrow \bullet_{\ell_H} \rightarrow (\bullet_{\ell_H}, \tau) \end{aligned}$
---

**Figure 3.3:** Implementation using World-Labeling

because the computation must return this high level  $\bullet$  value, the high-level world cannot exist in the same scope as the low-level world.

As with all previous design iterations, we can implement the lock operation in terms of join:

$$\begin{aligned} \text{lock } a &:: (\ell_L \leq \ell_H) \Rightarrow a \rightarrow \bullet_{\ell_L} \rightarrow (\bullet_{\ell_L}, \text{Locked}_{\ell_H} a) \\ \text{lock } a &= \text{join } (\lambda w \rightarrow (w, a)) \end{aligned}$$

Returning to our case study again, we can now express the Alice code in much the same way as before, save that the lock operation now requires a world to be passed in:

**Alice**

```

password ::  $\bullet_L \rightarrow (\bullet_L, \text{Password})$ 
password w0 =
  let w1           = putString "pick a password" w0
      (w2, pass)   = getLine w1
      (w3, lockedPass) = lock pass w2
      (w4, b)       = Bob.commonPassword (w3, lockedPass)
  in if b then
      let w5 = putString "your password is weak" w4
          in password w5
      else
          (w4, pass)

```

For Bob's code, he must use the low-level world  $w_1$  before gaining access to the high-level world  $w_2$ . While he has access to  $w_2$  he is therefore unable to call

functions such as *wget*, which expect a low-level world:

Bob
<pre> <i>commonPassword</i> :: (<math>\bullet_L</math>, Locked<sub>H</sub> String) → (<math>\bullet_L</math>, Locked<sub>H</sub> Bool) <i>commonPassword</i> (<i>w</i><sub>0</sub>, <i>lockedS</i>) =   <b>let</b> (<i>w</i><sub>1</sub>, <i>str</i>) = <i>wget</i> "http://common-password.list" <i>w</i><sub>0</sub>       <i>passwords</i> = <i>filter</i> (<i>not</i> ∘ <i>null</i>) (<i>lines</i> <i>str</i>)   <b>in</b> join <i>w</i><sub>1</sub> (λ <i>w</i><sub>2</sub> → <b>let</b> (<i>w</i><sub>3</sub>, <i>s</i>) = <i>unlock</i> <i>lockedS</i> <i>w</i><sub>2</sub>               <b>in</b> (<i>w</i><sub>3</sub>, <i>elem</i> <i>s</i> <i>passwords</i>))  <i>wget</i> :: URL → <math>\bullet_L</math> → (<math>\bullet_L</math>, String) </pre>

As we can see, this design addresses the main issues seen in the first iteration, and is also sufficiently expressive for the case study example. Furthermore, we can see a correspondence between this design and the monadic design of Abadi et al. (1999). Specifically, the monadic type that would be expressed in MAC as  $MAC\ \ell\ \tau$  can be expressed in our design as  $\bullet_\ell \rightarrow (\bullet_\ell, \tau)$ .

### 3.3 Implementation

The COGENT project includes a simpler compiler for a miniature dialect of COGENT called MINIGENT. While the language of MINIGENT does not lack any essential features from COGENT, the syntax is drastically reduced and additional compiler phases not relevant to type system experimentation are removed. This makes it a perfect candidate for a prototype implementation of FLOGENT.

In our earlier design sketches, we made extensive use of higher order functions to model the primitive operations for IFC. In COGENT (and MINIGENT), however, higher order functions can be quite cumbersome to use, due to the absence of support for closures. To sidestep this problem, we have implemented the fundamental IFC operations as built-in language constructs.

Implementation of these constructs in FLOGENT involves a number of milestones:

- Implementation of the syntactic changes required in the lexer and parser of the compiler.
- Formulation of the typing rules for these constructs.
- Implementation of the typing constraint generation phases for these constructs in the compiler's type checker.
- Solving of security level inequalities in the type constraint solver.

expressions	$e$	$::=$	$x \mid v \mid \ell$ $  e_1 \ \wr \ e_2$ $  e_1 \ e_2 \mid f[\overline{\tau_i}]$ $  \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2$ $  \mathbf{if} \ e_1 \ \mathbf{then} \ e_2 \ \mathbf{else} \ e_3$ $  e :: \tau$ $  \mathbf{unlock} \ e_1 \ e_2$ $  \mathbf{join} \ x \leftarrow e_1 \ \mathbf{in} \ e_2$ $  \mathbf{join} \bullet$ $  \mathbf{f} \bullet$	<i>(O'Connor, 2019)</i> <i>(primops)</i> <i>(applications)</i> <i>(type signature)</i>
types	$\tau$	$::=$	$\alpha$ $\alpha$ $(\tau_1, \tau_2)$ $\bullet_\ell$ $\mathbf{Locked}_\ell \ \tau$	<i>(O'Connor, 2019)</i> <i>(type variables)</i> <i>(unification variables)</i> <i>(tuples)</i> <i>(worlds)</i> <i>(secure data)</i>
literals		$::=$	$\mathbf{True} \mid \mathbf{False} \mid \mathbb{N}$	
constraints	$C, A$	$::=$	$\dots$ $\top$ $\mathbf{Drop} \ \tau$ $C_1 \wedge C_2$ $\tau_1 \sqsubseteq \tau_2$ $\ell_1 \leq \ell_2$	<i>(O'Connor, 2019)</i> <i>(tautology)</i> <i>(discardability)</i> <i>(conjunction)</i> <i>(subtyping)</i> <i>(level order)</i>
levels	$\ell$	$::=$	$\underline{L} \mid \underline{H}$	
contexts	$\Gamma$	$::=$	$\overline{x : \tau}$	
algorithmic contexts	$G$	$::=$	$\overline{x :_n \tau}$	
variables	$x, y$			
unification variables	$\alpha, \beta$			
numbers	$n$	$\in$	$\mathbb{N}$	

overlines indicate repetition

<b>Summary of Judgements</b>	
$A \vdash \Gamma \rightsquigarrow \Gamma_1 \boxplus \Gamma_2$	context splitting
$A, \Gamma \vdash e : \tau$	(non-algorithmic) typing rules
$A \vdash C$	constraint entailment
$G_1 \vdash e : \tau \rightsquigarrow C \mid G_2$	constraint generation

**Figure 3.4:** Summary of syntactic definitions

$$\begin{array}{c}
\frac{A \vdash \ell_2 \leq \ell_1 \quad A, \Gamma_1 \vdash e_1 : \bullet_{\ell_1} \quad A, \Gamma_2 \vdash e_2 : \text{Locked}_{\ell_2} \tau}{A, \Gamma \vdash \mathbf{unlock} \ e_1 \ e_2 : (\bullet_{\ell_1}, \tau)} \\
\frac{A \vdash \Gamma \rightsquigarrow \Gamma_1 \boxplus \Gamma_2 \quad A \vdash \ell_1 \leq \ell_2 \quad A, \Gamma_1 \vdash e_1 : \bullet_{\ell_1} \quad A, w : \bullet_{\ell_2}, \Gamma_2 \vdash e_2 : (\bullet_{\ell_2}, \tau)}{A, \Gamma \vdash \mathbf{join} \ w \leftarrow e_1 \ \mathbf{in} \ e_2 : (\bullet_{\ell_1}, \text{Locked}_{\ell_2} \tau)}
\end{array}$$

**Figure 3.5:** Typing rules

- Preliminary tests and examples to exercise the compiler with these new primitives.

Figure 3.4 outlines the main syntactic changes made to MINIGENT. Broadly, this formalisation is derived from the formalisation by O’Connor (2019), where specific additions are highlighted in cyan. The previously higher-order function invocation  $\text{join } e_1 (\lambda x \rightarrow e_2)$  is now modelled as the built in operation  $\mathbf{join} \ x \leftarrow e_1 \ \mathbf{in} \ e_2$ . The expressions  $\mathbf{join}\bullet$  and  $\mathbf{f}\bullet$  are new additional expressions that are not yet implemented in MINIGENT, these are expressions that we will be using for *term erasure* in the formalisation technique described in Chapter 5.

### 3.3.1 Typing Rules

Figure 3.5 describes the typing rules for the two new constructs in MINIGENT. The typing judgement is parameterised by an additional constraint  $A$  which is the set of all assumptions about polymorphic type variables. This can include level inequalities (e.g.  $\ell_L \leq \ell_H$ ) for functions such as  $\text{lock}$ , defined in subsection 3.2.2.

Because COGENT (and MINIGENT) has a uniqueness type system, care must be taken to avoid allowing unique variable bindings to be used more than once. Thus, we use the judgement  $A \vdash \Gamma \rightsquigarrow \Gamma_1 \boxplus \Gamma_2$  to explicitly split a context  $\Gamma$  into two parts  $\Gamma_1$  and  $\Gamma_2$ , where a binding can only be distributed to both  $\Gamma_1$  and  $\Gamma_2$  if it is *not* of a unique type. This ensures the uniqueness property, but makes the typing rules *non-algorithmic*—they do not straightforwardly correspond to a type checking algorithm. We describe how algorithmic type checking is achieved in the next subsection.

Other than these changes, the typing rules for these operations resemble the hypothetical type signatures described in subsection 3.2.2.

$$\begin{array}{c}
\frac{\ell_1, \ell_2, \alpha \text{ fresh} \quad G_1 \vdash e_1 : \bullet_{\ell_1} \rightsquigarrow C_1 \mid G_2 \quad G_2 \vdash e_2 : \text{Locked}_{\ell_2} \alpha \rightsquigarrow C_2 \mid G_3 \quad C_3 = \ell_2 \leq \ell_1 \quad C_4 = (\bullet_{\ell_1}, \alpha) \sqsubseteq \tau}{G_1 \vdash \text{unlock } e_1 e_2 : \tau \rightsquigarrow \bigwedge_{i=1}^4 C_i \mid G_3} \\
\frac{\ell_1, \ell_2, \alpha \text{ fresh} \quad G_1 \vdash e_1 : \bullet_{\ell_1} \rightsquigarrow C_1 \mid G_2 \quad x : \bullet_{\ell_2}, G_2 \vdash e_2 : (\bullet_{\ell_2}, \alpha) \rightsquigarrow C_2 \mid x : \bullet_{\ell_2}, G_3 \quad C_4 = \ell_1 \leq \ell_2 \quad C_5 = (\bullet_{\ell_1}, \text{Locked}_{\ell_2} \alpha) \sqsubseteq \tau \quad \text{if } n = 0 \text{ then } C_3 = \mathbf{Drop} \bullet_{\ell_2} \text{ else } C_3 = \top}{G_1 \vdash \text{join } x \leftarrow e_1 \text{ in } e_2 : \tau \rightsquigarrow \bigwedge_{i=1}^5 C_i \mid G_3}
\end{array}$$

**Figure 3.6:** Constraint Generation

### 3.3.2 Constraint Generation and Solving

To efficiently check well-typedness of an expression, COGENT (and MINIGENT) divide the type checking process into two phases:

1. Given an expression and a desired type (which may involve some unknown *unification variables*), generate a constraint such that the given expression has the desired type if and only if a satisfying assignment can be found to all unification variables in the constraint.
2. A constraint solving process to find such a satisfying assignment, if one exists.

Figure 3.6 describes the rules of the constraint generation process for **join** and **unlock**. Unlike the typing rules above, these rules describe a computation that has an input and output *algorithmic context*. The difference between an algorithmic context and a non-algorithmic context is that an algorithmic context is part of the type checker state, keeping track of the number of uses of each variable. The constraint generation judgement is written  $G_1 \vdash e : \tau \rightsquigarrow C \mid G_2$ . This is read as: given an input context  $G_1$ , the expression  $e$ , and the type  $\tau$ , we generate the constraint  $C$  and an output context  $G_2$ .

The only type of constraint we now generate, that previously was not already handled by MINIGENT is the security level inequalities  $\ell_1 \leq \ell_2$ . Seeing as we have a simple two element lattice (with just  $L$  and  $H$ ), solving these inequalities is trivial.

# Chapter 4

## Testing

To check that FLOGENT has the information flow control that we have aimed to achieve, we want to ensure that well-typed FLOGENT programs will preserve confidentiality: when secrets can only be read by processes with the right security clearance, and integrity: where data can only flow up the security level lattice. Any program that violates confidentiality or integrity should not be successfully type-checked by the compiler. Since these are implemented and tested in the MINIGENT compiler, the MINIGENT concrete syntax are used for the examples presented in this chapter.

### 4.1 Confidentiality

One example of a well-typed program `ReadSecret`, implemented in MINIGENT below, takes in the user's world at security level H and a file that is locked at level H; `take` is used to access the fields of records in a way that respects the linearity of the record fields. This function returns the world at the same level and the file that is now unlocked:

#### ReadSecret

```
readSecret: { world: World H , input : Locked H U8 }#  
            -> { world : World H , output:U8}#;  
readSecret r = take r2 { world = w } = r in  
    take r3 { input = i } = r2 in  
        unlock w i  
    end  
end;
```

On the other hand, if we try to implement a function `ReadSecretDodgy` that does the same thing but returns a world that has a security level lower than the input world, this program will not be type-checked. This is because once this function returns the unlocked file along with the lower level world, the file is now accessible to the world with lower security clearance when it should not be:

<b>ReadSecretDodgy</b>
<pre> readSecretDodgy: { world: World H , input : Locked H U8 }#     -&gt; { world : World L , output:U8}#; readSecretDodgy r = take r2 { world = w } = r in     take r3 { input = i } = r2 in     unlock w i     end end;</pre>

Similarly, the type-checker will not allow an implementation of a function that unlock files which is locked with a higher security level than the level of the input world:

<b>ReadSecretBad</b>
<pre> readSecretBad : { world: World L , input : Locked H U8 }#     -&gt; { world : World L , output:U8}#; readSecretBad r = take r2 { world = w } = r in     take r3 { input = i } = r2 in     unlock w i     end end;</pre>

These three MINIGENT examples show the restrictions that FLOGENT placed upon the accessibility of secret data in order to preserve confidentiality.

## 4.2 Integrity

To preserve integrity, we want to make sure that information flows from low to high and not the other way around. The example below shows a function `WriteSecret` that takes in a World of level L and a file, then outputs a world of the same level and a Locked file of level H. This program is type-checked by the compiler as the function is writing data from low to high.



### WriteSecret

```
writeSecret : [U8]. {world: World L, input: f}#  
              -> {world: World L, output: Locked H f}#;  
writeSecret r = take r1 {world=w} = r  
  in take r2 {input=i} = r1  
    in join w1 <- w  
      in {world = w1, output = i}  
    end  
  end  
end;
```

If we reverse the direction and try to leak data from high to low, we will get a function that will not type-check such as WriteSecretBad below.

### WriteSecretBad

```
writeSecretBad : [U8]. {world: World H, input: f}#  
                 -> {world: World H, output: Locked L f}#;  
writeSecretBad r = take r1 {world=w} = r  
  in take r2 {input=i} = r1  
    in join w1 <- w  
      in {world = w1, output = i}  
    end  
  end  
end;
```

## 4.3 Test result

These examples are corner cases that illustrate how the MINIGENT type checker is capable of statically enforce both confidentiality and integrity at compile time, with no runtime checks required.

# Chapter 5

## Formalisation

So far in the previous chapters, we have seen the design, implementation and testing of FLOGENT. The next logical step is to formalise and prove the security guarantees that we want Flogent to have. This chapter presents the formalisation and proof techniques that we can use towards achieving that. The proof is not present as it is out of the scope of this thesis.


### 5.1 Noninterference

Noninterference (Goguen and Meseguer, 1982) is a property that restricts the information flow through a system that we want to preserve in FLOGENT; the formal definition of this property will be discussed later in this chapter.

A common proof technique used to prove noninterference in functional programs is known as *term erasure* (Li and Zdancewic, 2010). This technique uses an erasure function  $\varepsilon_\ell$  on terms, which rewrites data above the reader's security level  $\ell$ , into a special symbol  $\bullet$ , this represents data that are supposed to be secret to the reader will be hidden. I describe the erasure function in more detail later in subsection 5.1.1.

Once we have established the erasure function, the proof will involve the relationship between the erasure function and the reduction step, also known as the single-step simulation between the erased terms and the original ones.

We can intuitively understand this relationship by looking at the commutative diagram in Figure 5.1. In this figure, we can see that a program  $p$  can take a step to  $p'$  and then erase  $p'$  to reach  $\varepsilon_\ell(p')$ , or the same program  $p$  can first be erased to  $\varepsilon_\ell(p)$  then take a step and arrive at the same result  $\varepsilon_\ell(p')$ . This is important



**Figure 5.1:** Single-step simulation

for proving noninterference as it shows that any erased secret will not affect the behaviour of the program.

Note that in order to reason with the reduction step, we will define the small step semantic for FLOGENT in subsection 5.1.2.

We can formally state the noninterference property using the following propositions and definition:

**Proposition 1** (Single-Step Simulation).

If  $p_1 \mapsto p_2$  then  $\varepsilon_\ell(p_1) \mapsto \varepsilon_\ell(p_2)$ .

**Proposition 2** (Determinancy).

If  $p_1 \mapsto p_2$  and If  $p_1 \mapsto p_3$  then  $p_2 \equiv p_3$ .

**Proposition 3** ( $\approx_\ell$  Preservation).

If  $c_1 \approx_\ell c_2$ ,  $c_1 \mapsto c'_1$ ,  $c_2 \mapsto c'_2$ , then  $c'_1 \approx_\ell c'_2$ .

**Definition** ( $\ell$ -equivalence).

Two program  $p_1$  and  $p_2$  are indistinguishable for someone at security level  $\ell$ , written as  $c_1 \approx_\ell c_2$ , if and only if  $\varepsilon_\ell(c_1) \equiv \varepsilon_\ell(c_2)$ .

**Definition** (Big-step evaluation  $\Downarrow$ ).

We will define the big-step evaluation ( $\Downarrow$ ) as the reflexive transitive closure over small-step ( $\mapsto$ ) that results in a value.

**Theorem** (Noninterference Theorem).

If  $c_1 \approx_\ell c_2$ ,  $c_1 \Downarrow c'_1$  and  $c_2 \Downarrow c'_2$ , then  $c'_1 \approx_\ell c'_2$

Next, we will define the erasure function as well as small-step semantic rules.

### 5.1.1 Definition of the Erasure Function

This section presents the definition of the erasure function  $\varepsilon_\ell$  for every expression and value. Figure 5.2 and Figure 3.4 show the summaries of the value semantics and the syntactic expressions.

value semantic values	$v ::= \ell$	<i>(literals)</i>
	$\langle\langle \lambda x. e \rangle\rangle$	<i>(function values)</i>
	$\langle\langle \mathbf{abs}. f \mid \vec{\tau} \rangle\rangle$	<i>(abstract functions)</i>
	$K v$	<i>(variant values)</i>
	$\{\overline{f \rightarrow v}\}$	<i>(records)</i>
	$a_v$	<i>(abstract values)</i>
	$V_w$	<i>(World value)</i>
	<b>Locked</b> $v$	<i>(Locked value)</i>
	$\bullet$	<i>(Erased value)</i>

**Figure 5.2:** Summary of value semantics

Figure 5.3 shows the definitions for the erasure function when given different values and expressions. For most expressions, the erasure function is homomorphically applied to the sub-expressions (e.g.  $\varepsilon_\ell(e_1 \wr e_2) = (\varepsilon_\ell(e_1) \wr \varepsilon_\ell(e_2))$ ). Primitive values (e.g. `True`) that are not sensitive will not be affected by the erasure function and the function will just return what it was given.

The more interesting cases for the erasure functions are cases where sensitive content is rewritten into  $\bullet$  when it has a security level that is higher than the reader's access, ensuring that sensitive information is not visible to the reader (e.g. **Locked**  $v$  being erased to **Locked**  $\bullet$ ). In the case of  $f\bullet$ , erasure is delayed until the function body is available and can be erased.

$\varepsilon_\ell(x) = x$	<i>(variable)</i>
$\varepsilon_\ell(\ell) = \ell$	<i>(literal)</i>
$\varepsilon_\ell(v) = \begin{cases} \text{defined below} & \text{if } v = \mathbf{Locked } v \\ v & \text{otherwise} \end{cases}$	<i>(value)</i>
$\varepsilon_\ell(V_w) = V_w$	<i>(world value)</i>
$\varepsilon_\ell(e_1 \wr e_2) = (\varepsilon_\ell(e_1) \wr \varepsilon_\ell(e_2))$	<i>(primop)</i>
$\varepsilon_\ell(e_1 e_2) = \varepsilon_\ell(e_1) \varepsilon_\ell(e_2)$	<i>(application)</i>
$\varepsilon_\ell(\mathbf{let } x = e_1 \mathbf{ in } e_2) = \mathbf{let } x = \varepsilon_\ell(e_1) \mathbf{ in } \varepsilon_\ell(e_2)$	<i>(let)</i>
$\varepsilon_\ell(\mathbf{if } e_1 \mathbf{ then } e_2 \mathbf{ else } e_3) = \mathbf{if } \varepsilon_\ell(e_1) \mathbf{ then } \varepsilon_\ell(e_2) \mathbf{ else } \varepsilon_\ell(e_3)$	<i>(if)</i>
$\varepsilon_\ell(f[\vec{\tau}_i]) = f \bullet [\vec{\tau}_i]$	<i>(fun. app.)</i>
$\varepsilon_\ell(\langle\langle \lambda x. e[\alpha_i := \vec{\tau}_i] \rangle\rangle) = \langle\langle \lambda x. \varepsilon(e[\alpha_i := \vec{\tau}_i]) \rangle\rangle$	<i>(abs. function)</i>
$\varepsilon_\ell(\mathbf{join } e_1 e_2 :: \langle w, \mathbf{Locked } \ell' \tau \rangle) = \begin{cases} \mathbf{join} \bullet \varepsilon_\ell(e_1) \varepsilon_\ell(e_2) & \text{if } \ell' > \ell \\ \mathbf{join} \varepsilon_\ell(e_1) \varepsilon_\ell(e_2) & \text{otherwise} \end{cases}$	
Erasure for join.	
$\varepsilon_\ell(\mathbf{Locked } v :: \mathbf{Locked } \ell' \tau) = \begin{cases} \mathbf{Locked } \bullet & \text{if } \ell' > \ell \\ \mathbf{Locked } \varepsilon_\ell(v) & \text{otherwise} \end{cases}$	
Erasure for Locked value.	

**Figure 5.3:** Erasure Function for expressions and values

### 5.1.2 Small-step Semantics

COGENT has a big-step operational semantics. To establish the reduction step that we need to investigate the single-step relationship, we present a small-step semantics for COGENT in Figure 5.4 and Figure 5.5 as well as for the new operations that we introduced to FLOGENT: Figure 5.6. The big-step evaluation is defined as the reflexive transitive closure over small-step that results in a value.

$$\begin{array}{c}
\frac{c \mapsto \text{True} \quad e_1 \mapsto v}{\text{if } c \ e_1 \ e_2 \mapsto v} \text{IF-TRUE} \\
\frac{c \mapsto \text{False} \quad e_2 \mapsto v}{\text{if } c \ e_1 \ e_2 \mapsto v} \text{IF-FALSE} \\
\frac{c \mapsto c'}{\text{if } c \ e_1 \ e_2 \mapsto \text{if } c' \ e_1 \ e_2} \text{IF-COND} \\
\frac{}{\text{if True } e_1 \ e_2 \mapsto e_1} \text{IF-TRUE} \\
\frac{}{\text{if False } e_1 \ e_2 \mapsto e_2} \text{IF-FALSE} \\
\frac{e_1 \mapsto e'_1}{e_1 \ \wr \ e_2 \mapsto e'_1 \ \wr \ e_2} \text{OP}_1 \\
\frac{e_2 \mapsto e'_2}{v_1 \ \wr \ e_2 \mapsto v_1 \ \wr \ e'_2} \text{OP}_2 \\
\frac{}{v_1 \ \wr \ v_2 \mapsto v_1 \llbracket ? \rrbracket v_2} \text{OP}_3 \\
\frac{e_1 \mapsto e'_1}{\text{let } x = e_1 \ \text{in } e_2 \mapsto \text{let } x = e'_1 \ \text{in } e_2} \text{LET}_1 \\
\frac{}{\text{let } x = v \ \text{in } e_2 \mapsto e_2[x := v]} \text{LET}_2 \\
\frac{e_1 \mapsto e'_1}{(e_1)(e_2) \mapsto (e'_1)(e_2)} \text{APP}_1 \\
\frac{e_2 \mapsto e'_2}{(\langle\langle \lambda x. e \rangle\rangle)(e_2) \mapsto (\langle\langle \lambda x. e \rangle\rangle)(e'_2)} \text{APP}_2 \\
\frac{}{\langle\langle \lambda x. e \rangle\rangle v \mapsto e[x := v]} \text{APP}_3 \\
\frac{e_2 \mapsto e'_2}{(\langle\langle \text{abs. } f | \bar{\tau} \rangle\rangle)(e_2) \mapsto (\langle\langle \text{abs. } f | \bar{\tau} \rangle\rangle)(e'_2)} \text{ABS.} \\
\frac{}{\langle\langle \text{abs. } f | \bar{\tau} \rangle\rangle v \mapsto \llbracket f \rrbracket v} \text{ABS. APP}
\end{array}$$

**Figure 5.4:** Small step semantics of the COGENT language (part 1)

$$\begin{array}{c}
\frac{e_1 \mapsto e'_1}{\mathbf{case } e_1 \mathbf{ of } C \ x \rightarrow e_2 \ \mathbf{else } y \rightarrow e_3 \mapsto \mathbf{case } e'_1 \mathbf{ of } C \ x \rightarrow e_2 \ \mathbf{else } y \rightarrow e_3} \text{CASE}_1 \\
\\
\frac{}{\mathbf{case } C \ v \ \mathbf{of } C \ x \rightarrow e_2 \ \mathbf{else } y \rightarrow e_3 \mapsto e_2[x := v]} \text{CASE}_2 \\
\\
\frac{B \neq C}{\mathbf{case } B \ v \ \mathbf{of } C \ x \rightarrow e_2 \ \mathbf{else } y \rightarrow e_3 \mapsto e_3[y := B \ v]} \text{CASE}_3 \\
\\
\frac{e \mapsto e'}{\mathbf{esac } e \mapsto \mathbf{esac } e'} \text{ESAC}_1 \\
\\
\frac{}{\mathbf{esac } C \ v \mapsto v} \text{ESAC}_2 \\
\\
\frac{e_1 \mapsto e'_1}{\mathbf{take } x \{f_K = y\} = e_1 \ \mathbf{in } e_2 \mapsto \mathbf{take } x \{f_K = y\} = e'_1 \ \mathbf{in } e_2} \text{TAKE}_1 \\
\\
\frac{e_1 \mapsto e'_1}{\mathbf{take } x \{f_K = y\} = \overline{\{f_i = v_i\}} \ \mathbf{in } e_2 \mapsto e_2 [x := \overline{\{f_i = v_i\}}][y := v_K]} \text{TAKE}_2 \\
\\
\frac{e_1 \mapsto e'_1}{\mathbf{put } e_1.f_k := e_2 \mapsto \mathbf{put } e'_1.f_k := e_2} \text{PUT}_1 \\
\\
\frac{e_2 \mapsto e'_2}{\mathbf{put } \{\overline{f_i = v_i}\}.f_k := e_2 \mapsto \mathbf{put } \{\overline{f_i = v_i}\}.f_k := e'_2} \text{PUT}_2 \\
\\
\frac{\mathbf{for } i \neq k : v'_i = v_i}{\mathbf{put } \{\overline{f_i = v_i}\}.f_k := v'_k \mapsto \{\overline{f_i = v_i}\}} \text{PUT}_3 \\
\\
\frac{\mathbf{funDef}(f) = \langle f :: \forall(\overline{\alpha_i} ::_K \overline{K_i}).\tau \mapsto \tau_1 f x = e \rangle}{f[\overline{\tau_i}] \mapsto \langle \lambda x. e[\overline{\tau_i}/\alpha_i] \rangle} \text{VFUNC} \\
\\
\frac{e \mapsto e'}{C \ e \mapsto C \ e'} \text{VCONS}
\end{array}$$

Figure 5.5: Small step semantics of the COGENT language (part 2)

$$\begin{array}{c}
\frac{e_1 \mapsto e'_1}{\mathbf{unlock} \ e_1 \ e_2 \mapsto \mathbf{unlock} \ e'_1 \ e_2} \text{UNLOCK}_1 \\
\frac{e_2 \mapsto e'_2}{\mathbf{unlock} \ (\text{Locked } v_1) \ e_2 \mapsto \mathbf{unlock} \ (\text{Locked } v_1) \ e'_2} \text{UNLOCK}_2 \\
\frac{}{\mathbf{unlock} \ (\text{Locked } v_1) \ v_2 \mapsto v_1 \ v_2} \text{UNLOCK}_3 \\
\frac{e_1 \mapsto e'_1}{\mathbf{join} \ w \leftarrow e_1 \ \mathbf{in} \ e_2 \mapsto \mathbf{join} \ w \leftarrow e'_1 \ \mathbf{in} \ e_2} \text{JOIN}_1 \\
\frac{e_2[x := v_w] \Downarrow (v'_w, v)}{\mathbf{join} \ x \leftarrow v_w \ \mathbf{in} \ e_2 \mapsto (v'_w, \text{Locked } v)} \text{JOIN}_2 \\
\frac{e_1 \mapsto e'_1}{\mathbf{join} \bullet \ w \leftarrow e_1 \ \mathbf{in} \ e_2 \mapsto \mathbf{join} \ w \leftarrow e'_1 \ \mathbf{in} \ e_2} \text{JOIN}_{\bullet 1} \\
\frac{e_2[x := v_w] \Downarrow (v'_w, v)}{\mathbf{join} \bullet \ x \leftarrow v_w \ \mathbf{in} \ e_2 \mapsto (v'_w, \text{Locked } .)} \text{JOIN}_{\bullet 2} \\
\frac{\text{funDef}(f) = \langle f :: \forall (\alpha_i ::_K K_i). \tau \mapsto \tau'_1 f x = e \rangle}{f \bullet [\vec{\tau}_1] \mapsto \varepsilon(\langle \langle \lambda x. e[\vec{\tau}_1/\alpha_i] \rangle \rangle)} \text{VFUN}_{\bullet}
\end{array}$$

**Figure 5.6:** Small step semantics of new operations



## 5.2 Applying to Flogent

So far, we have designed the type system for FLOGENT that uses linear types which enforces information flow control. We have a case study that conceptualising the security features we want to achieve and also tested the implementation on several edge cases which yield results that we expected. We formalised the proof technique for proving noninterference which involve proving the single-step simulation relationship between program reduction step and term erasure. We defined a small step operational semantics for FLOGENT which allow us to reason with the reduction steps in FLOGENT. We intuitively believe that this system maintains confidentiality and integrity, however, further work is required to look into other potential proof techniques or improving on the current single step proof technique for it to work on our system. Here is a concrete example of why the current approach does not work: Figure 5.7; this example shows the blue path does not hold in the commutative diagram. This is because  $\epsilon_\ell$  would not erase  $\langle v_1, v_2 \rangle$  into  $\langle \bullet, v_2 \rangle$  as the erasure function would not see  $v_1$  as sensitive by itself. We will need to use a new technique to perform the security proofs since the existing technique does not apply to the current design in FLOGENT. Coming up with a novel proof technique to prove security in this context is beyond the scope of this thesis and remains open for future work.




Figure 5.7: Single-step simulation for **unlock**

## 5.3 Summary

To summarise my contribution in this thesis, I extended MINIGENT with an information flow aware language FLOGENT with a type system that utilise linear types. I implemented the core primitive operations **unlock** and **join** in MINIGENT, I conducted testing on some examples to check that confidentiality and integrity are preserved in different cases. I formalised the noninterference property that we want in FLOGENT which involves creating small-step semantics for the core language and new operations, defining the erasure function as well as investigating the single-step simulation relationship between the erasure function and the reduction step of a program.

# Bibliography

- Martín Abadi, Anindya Banerjee, Nevin Heintze, and Jon G. Riecke (1999). “A Core Calculus of Dependency”. In: *Principles of Programming Languages*. POPL ’99. San Antonio, Texas, USA: ACM, pp. 147–160. ISBN: 1-58113-095-3. DOI: 10.1145/292540.292555.
- Pablo Buiras, Dimitrios Vytiniotis, and Alejandro Russo (2015). “HLIO: Mixing Static and Dynamic Typing for Information-flow Control in Haskell”. In: *International Conference on Functional Programming*. ICFP 2015. Vancouver, BC, Canada: ACM, pp. 289–301. ISBN: 978-1-4503-3669-7. DOI: 10.1145/2784731.2784758.
- David Costanzo and Zhong Shao (2014). “A Separation Logic for Enforcing Declarative Information Flow Control Policies”. In: *Principles of Security and Trust*. Vol. 8414, pp. 179–198. DOI: 10.1007/978-3-642-54792-8\_10.
- Dorothy E. Denning and Peter J. Denning (1977). “Certification of Programs for Secure Information Flow”. In: *Communications of the ACM* 20.7, pp. 504–513. ISSN: 0001-0782. DOI: 10.1145/359636.359712.
- J. A. Goguen and J. Meseguer (1982). “Security Policies and Security Models”. In: *1982 IEEE Symposium on Security and Privacy*, pp. 11–11.
- Haskell (2010). *Haskell 2010 Language Report*. Ed. by Simon Marlow. URL: <https://www.haskell.org/onlinereport/haskell2010/>.
- C. Hritcu, M. Greenberg, B. Karel, B. C. Pierce, and G. Morrisett (2013). “All Your IFCEException Are Belong to Us”. In: *IEEE Symposium on Security and Privacy*, pp. 3–17. DOI: 10.1109/SP.2013.10.
- Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom (2019). “Spectre Attacks: Exploiting Speculative Execution”. In: *Symposium on Security and Privacy*.
- Peng Li and S. Zdancewic (2010). “Arrows for Secure Information Flow”. In: *Theor. Comput. Sci.* Pp. 1974–1994. DOI: 10.1016/j.tcs.2010.01.025.
- Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval

- Yarom, and Mike Hamburg (2018). “Meltdown: Reading Kernel Memory from User Space”. In: *USENIX Security Symposium*.
- Matt Miller (2019). *Trends, challenges, and strategic shifts in the software vulnerability mitigation landscape*. Accessed: 28-04-2019. URL: [https://github.com/Microsoft/MSRC-Security-Research/tree/master/presentations/2019\\_02\\_BlueHatIL](https://github.com/Microsoft/MSRC-Security-Research/tree/master/presentations/2019_02_BlueHatIL).
- Robin Milner, Mads Tofte, and David MacQueen (1997). *The Definition of Standard ML*. Cambridge, MA, USA: MIT Press. ISBN: 0262631814.
- MITRE Corporation (2019). *Common Vulnerabilities and Exposures*. Accessed: 26-04-2019. URL: <https://cve.mitre.org/>.
- Eugenio Moggi (1991). “Notions of Computation and Monads”. In: *Information and Computation* 93.1, pp. 55–92. ISSN: 0890-5401. DOI: 10.1016/0890-5401(91)90052-4.
- Toby Murray (2019). *Private Correspondance*.
- Toby Murray, Robert Sison, Edward Pierzchalski, and Christine Rizkallah (2016). “Compositional Verification and Refinement of Concurrent Value-Dependent Noninterference”. In: pp. 417–431. DOI: 10.1109/CSF.2016.36.
- Tobias Nipkow and Gerwin Klein (2014). *Concrete Semantics: With Isabelle/HOL*. Springer Publishing Company, Incorporated. ISBN: 978-3-319-10541-3.
- Tobias Nipkow, Lawrence Paulson, and Markus Wenzel (2002). *Isabelle/HOL: A Proof Assistant for Higher-Order Logic*. Vol. 2283. Lecture Notes in Computer Science.
- Liam O’Connor (2019). *Type Systems for Systems Types*. UNSW Sydney.
- Liam O’Connor, Zilin Chen, Christine Rizkallah, Sidney Amani, Japheth Lim, Toby Murray, Yutaka Nagashima, Thomas Sewell, and Gerwin Klein (2016). “Refinement Through Restraint: Bringing Down the Cost of Verification”. In: *International Conference on Functional Programming*. ICFP 2016. Nara, Japan: ACM, pp. 89–102. ISBN: 978-1-4503-4219-3. DOI: 10.1145/2951913.2951940.
- François Pottier and Vincent Simonet (2003). “Information Flow Inference for ML”. In: *Transactions in Programming Languages and Systems* 25.1, pp. 117–158. ISSN: 0164-0925. DOI: 10.1145/596980.596983.
- John C. Reynolds (2002). “Separation Logic: A Logic for Shared Mutable Data Structures”. In: *Symposium on Logic in Computer Science*. LICS ’02. IEEE Computer Society, pp. 55–74. ISBN: 0-7695-1483-9.
- Alejandro Russo (2015). “Functional Pearl: Two Can Keep a Secret, if One of Them Uses Haskell”. In: *International Conference on Functional Programming*. ICFP 2015. Vancouver, BC, Canada: ACM, pp. 280–288. ISBN: 978-1-4503-3669-7. DOI: 10.1145/2784731.2784756.
- Z. Somogyi, F. J. Henderson, and T. C. Conway (1995). “Mercury, an Efficient Purely Declarative Logic Programming Language”. In: *Australian Computer Science Conference*, pp. 499–512.

- Deian Stefan, Alejandro Russo, John C. Mitchell, and David Mazières (2011). “Flexible Dynamic Information Flow Control in Haskell”. In: *Haskell Symposium*. Haskell ’11. Tokyo, Japan: ACM, pp. 95–106. ISBN: 978-1-4503-0860-1. DOI: 10.1145/2034675.2034688.
- Laszlo Szekeres, Mathias Payer, Tao Wei, and Dawn Song (2013). “SoK: Eternal War in Memory”. In: *Symposium on Security and Privacy*. SP ’13. IEEE Computer Society, pp. 48–62. ISBN: 978-0-7695-4977-4. DOI: 10.1109/SP.2013.13.
- Marco Vassena, Alejandro Russo, Pablo Buiras, and Lucas Waye (2018). “MAC: A verified static information-flow control library”. In: *Journal of Logical and Algebraic Methods in Programming* 95, pp. 148–180. ISSN: 2352-2208. DOI: 10.1016/j.jlamp.2017.12.003.
- Dennis Volpano, Cynthia Irvine, and Geoffrey Smith (1996). “A Sound Type System for Secure Flow Analysis”. In: *Journal of Computer Security* 4.2-3, pp. 167–187. ISSN: 0926-227X.