

A Resourceful Monad for IO

Luke Lau

A DISSERTATION

Presented to the University of Dublin, Trinity College
In Partial Fulfilment of the Requirements for the Degree of

Masters in Computer Science

Supervisor of Dissertation
Glenn Strong

Declaration

I hereby declare that this project is entirely my own work and that it has not been submitted as an exercise for a degree at this or any other university.



Luke Lau
April 2020

Contents

1 Introduction	4
1.1 Overview	7
2 Background	8
2.1 Separation logic	8
2.2 Monads categorically	9
2.3 Type-level programming within GHC	11
2.4 Hindley-Damas-Milner	12
2.4.1 Syntax	13
2.4.2 Static semantics	13
2.4.3 Dynamic Semantics	16
3 The Resourceful System	17
3.1 Syntax	17
3.1.1 Substitution	19
3.1.2 Barendregt's variable convention	20
3.2 Static semantics	21
3.2.1 Examples	25
3.3 Dynamic semantics	26
3.3.1 Values	27
3.3.2 Small-step semantics	27
3.3.3 The reflexive and transitive closure	29
4 Properties	31
4.1 Type soundness	36
5 Mechanisation in Agda	44
5.1 Definitions	44
5.2 Type schemes and type variables	47
5.3 Properties	49
5.4 Postulations	50

6 Evaluation and further work	51
6.1 Separation logic	51
6.2 Well formed heaps	52
6.3 Let polymorphism	53
6.4 Almost syntax directed	54
6.5 Modelling state within the monad	54
6.6 Extending the inference algorithm	55
6.7 Heap polymorphism	56
6.8 Dependently typed resources	56
6.9 Casting heaps	58
7 Conclusion	59
A Denotational semantics of Hindley-Damas-Milner	62
B Proofs	64
B.1 Transitivity of the subheap relation	64
B.2 Helper lemmas	64

Chapter 1

Introduction

The immutability and purity of some functional languages make them seem like a perfect fit for parallelism and concurrency. The lack of side effects mean we are free to compute expressions in whatever order we please, and there is no shared mutable state to plague us with race conditions. However in the real world, code is never completely pure. Useful programs need to interact with the real world at some point, whether that be by reading from a keyboard or sending packets over a network. That is, they need to carry out side effects: And the order in which they are carried out is important.

So when it comes to side effects in a language with referential transparency, the language must model the sequence in which effects are carried out. One of the most successful approaches has been to capture the state of the outside world, and have functions with side effects that write to that state. This way, if writing to `stdout` causes some particles to be perturbed on a display monitor, then we can track that by saying the state of the world has changed, and that it matters what order it changed. We don't actually track what particles are there, where they might be or what their current charge is. All we need to know is that the world is slightly different than what it was before.

Concurrent Clean models this by threading the world in and out of functions. Uniqueness types guarantee that the same world is only used once, so that the programmer doesn't create an alternative timeline by duplicating it.

```
readFile :: !String !*World -> (!MaybeError FileError String, !*World)
```

Haskell also treats the world as a state, but without the uniqueness guarantee. Any function that interacts with the world returns a function, which returns a new version of the World alongside the function result.¹

```
type IO a = World -> (World, a)
```

How Haskell ensures that an old world isn't erroneously used is by hiding the updating of the state from the user, tucking it away into a *monad*. This idea of

¹The actual definition in GHC is
`newtype IO a = IO (State## RealWorld -> (## State## RealWorld, a ##))`

using monads to sequence together stateful computations was first introduced by Peyton Jones and Wadler [15][18]. The programmer no longer needs to keep track of the state of the world. They can keep their imperative code imperative, and their pure code pure. This marriage of monads and IO is one of the crown jewels of functional language research to come out of Haskell.

```
instance Monad IO where
  return x = \w -> (w, x)
  x >>= f = \w -> let (w', y) = x w in f y w'
```

Now IO actions can be easily chained together in a type-safe way that ensures their ordering. Unfortunately, this ordering imposes limitations. One of the main benefits of pure functional languages is that since expressions do not have side effects, there is no restriction on what order they need to be evaluated in. Take for example the following program:

```
f, g, h :: Int -> Int
f x = g x + h x
```

`g` could be evaluated before `h`, or `h` could be evaluated before `g`. It wouldn't make a difference because there are no side effects. One might be tempted then to evaluate the two expressions concurrently, and indeed that would be safe to do so. The same cannot be said for impure IO actions however, and Haskell's type system is well aware of that.

```
f, g, h :: Int -> IO Int
f x = g x >>= \y -> y + h x
```

We need to explicitly bind the actions and sequence evaluation. Does this mean that concurrency is impossible for IO actions? Not at all, many languages provide primitives to run these actions concurrently in a type safe way. Haskell's base library has `forkIO`, but for simplicity we are going to assume the existence of a higher level function that runs two IO actions simultaneously and collects the results.

```
concurrently :: IO a -> IO b -> IO (a, b)
```

Now we can use it to run our two IO actions side by side safely.

```
f, g, h :: Int -> IO Int
f x = g x `concurrently` h x >>= \(a, b) -> return (a + b)
```

But what if `g` and `h` actually looked like this?

```
g x = do
  txt <- readFile "foo.txt"
  return (x + (read txt))
h x = do
  writeFile "foo.txt" "hello"
  return (42 - x)
```

```
readFile :: FilePath -> IO String
writeFile :: FilePath -> String -> IO ()
```

Running these two functions concurrently could be disastrous, as the order in which they execute could affect the outcome of the program. These innocuous looking IO actions then end up introducing non-determinism and race conditions. We know statically that a program such as

```
writeFile "foo.txt" "a" `concurrently` writeFile "foo.txt" "b"
```

should probably not be allowed, because it is blatantly accessing the same resource simultaneously. But then why did the type system allow it? Has it failed us? The goal of the type system is to disallow as many incorrect programs as possible whilst allowing all correct programs. It is a fine line as to what programs are deemed correct and what are deemed incorrect. A type system too lenient and buggy programs will creep through: but a type system too strict and the programmer will waste time fighting the type checker.

In this dissertation we aim to find a point in the design space that rejects such programs as the one above. We do **not** want to allow programs that have glaring race conditions, where we can see that there is a contentious access of resources. Our work is based around the idea of keeping track of what resources are in use, at the type level. We first imagine what this might look like in Haskell by adding another type parameter to our IO type, representing what resource an IO action uses:

```
type IO r a = World -> (World, a)
```

r is a phantom type parameter, which only exists at the type level. Now our type signatures could look like this, annotating the API with what resources it might use.

```
data Resource = FileSystem | Net | Database | OpenGL | ...
readFile :: FilePath -> IO FileSystem ()
writeFile :: FilePath -> IO FileSystem String
readSocket :: Socket -> IO Net ()
runQuery :: Query a -> IO Database a
swapBuffers :: IO OpenGL ()
```

Keep in mind we are painting in broad strokes when we use the word “resource”. In the running example the resource has been a file, `foo.txt`, but the notion of a resource can be as broad or as specific as a function needs it to be. It could represent a specific database instance, or a network socket. For simplicity in our example we will consider the entire file system as a single resource, the entire network as a single resource, and so on.

Now that we know what resources each IO action is using, we would like to change the type of our concurrent function to take advantage of this new information. Perhaps we would like to reject any two functions that use the same resource, i.e. it only accepts IO actions with distinct resources.

```
concurrently :: r ~ s => IO r a -> IO s b -> IO ? (a, b)
```

You can read $r \not\sim s$ as “ r is distinct from s ”, or the opposite of a $r \sim s$ equality constraint that one might see in a type signature. This, however, does not exist in Haskell. And what does it exactly mean for two resources to “be distinct”? And what resources would the returned IO use?

These questions are answered with a formal definition of a type system that tracks resource usage. We explore a specific point in the design space, where the type system rejects programs like

$$\text{readFile} \Upsilon \text{readFile}$$

but accepts and assigns types to programs such as

$$\text{readFile} \Upsilon \text{readNet} : \text{IO}_{\text{File} \cup \text{Net}} \square \times \square$$

In short, we aim to create a type system that keeps tracks of the resources being used, so that the programmer doesn’t have to. It **does not** aim to solve concurrency — there will still be programs that have concurrency errors that the type system will still allow. We just aim to narrow down the scope of valid programs, by ruling out those with blatant, concurrent resource access errors.

1.1 Overview

This dissertation has been written in a style that doesn’t assume prior knowledge of type theory, and so it should double as a tutorial along our journey to create our resourceful type system.

In Chapter 2 we will talk about the inspirations of the type system, namely separation logic, and how it parallels with our monadic language. We will also briefly go over what we mean by a monad formally, and then look at the original design of Hindley-Damas-Milner which we will build upon. Chapter 3 introduces the language, its type system and its semantics. It gives a complete definition of all the parts needed to prove properties. We prove these properties in Chapter 4, in which we eventually build up and present a proof of its soundness. This proof, and the others that accompany it, are then mechanised within the dependently typed language and proof assistant Agda. The methodology of this is discussed in Chapter 5. Finally, Chapter 6 evaluates our type system. We look at how it might be refined, and find many areas of further work that would be interesting to explore.

Chapter 2

Background

2.1 Separation logic

Much of the original inspiration for this work came from separation logic. Separation logic [14, 16] is an extension to Hoare logic, a system for formalising properties about imperative programs, and has been used to prove properties about resource access in concurrent programs [13]. It is formalised through *specifications*, which comprise of some code alongside a *precondition* a *postcondition*.

$$\{precondition\} \text{ code } \{postcondition\}$$

The preconditions and postconditions make assertions about what points to what, inside the *heap*. In the example that O’Hearn gives, a cyclic list is constructed from two pointers. The precondition asserts that x and y point to 0 initially, and then the postcondition asserts that they point to each other afterwards.

$$\begin{aligned} &\{x \mapsto 0 * y \mapsto 0\} \\ &\quad [x] = y; \\ &\quad [y] = x; \\ &\{x \mapsto y * y \mapsto x\} \end{aligned}$$

The code sets the value of x to the location of y , and vice-versa. Separation logic extends Hoare logic with the notion of splitting the heap. In the example above, instead of asserting $x \mapsto y \wedge y \mapsto x$, the separation conjunction $*$ is used which says “ x points to y and separately y points to x ”. The heap must be able to split into two disjoint subheaps, such that $x \mapsto y$ holds in one and $y \mapsto x$ in the other. As a result x and y cannot point to the same location in the heap. These preconditions and postconditions can be derived from so-called small axioms. As an example, an imperative language might contain an axiom for writing to a pointer such as

$$\{x \mapsto -\}[x] = v \{x \mapsto v\}$$

amongst others for primitive operations like reading, allocating and freeing pointers.

The *frame rule* allows a specification to be extended by adding extra predicates that do not mention any variables used by the code.

$$\frac{\{p\} \text{code} \{q\}}{\{p * r\} \text{code} \{q * r\}}$$

where no variable occurring free
in r is modified by code

This rule ends up being quite significant, as its the reason why separation logic allows for programs to be reasoned about locally — the predicates involving the heap are shrunk so that they only contain variables that the code modifies.

There is also the *concurrency rule* for reasoning about running two programs side by side. It states that the preconditions and postconditions of the two programs must hold separately to each other.

$$\frac{\{p_1\} \text{code}_1 \{q_1\} \quad \{p_2\} \text{code}_2 \{q_2\}}{\{p_1 * p_2\} \text{code}_1 \parallel \text{code}_2 \{q_1 * q_2\}}$$

We will see later on in Chapter 6 how the system we end up developing parallels with this.

Krishnaswami built upon separation logic with a higher-order ML like language [9] with types. Like our system, it utilised monadic binding for sequencing, allowing imperative computation within an otherwise pure language. Unlike our system, this included separation logic predicates directly within the language. We will not work with predicates, and instead push all our assertions into the type level.

2.2 Monads categorically

As the language we will develop in Chapter 3 revolves around monads at its core, it is a good idea to take a brief look at what exactly they are from a mathematical standpoint, and dip our toes into the cool waters of category theory. There is an infamous definition [1] of a monad that goes along the lines of

A monad is just a monoid in the category of endofunctors, what's the problem?

Let's break this down. A *category* is a collection of objects (like a set), and morphisms (like functions between sets). In the same way, Haskell has `Hask` — the category of types. Its objects are types such as `Bool` and `Int`, and its morphisms are functions taking a type and returning another.

```
Bool, Int :: * -- Bool and Int are types
True :: Bool
42 :: Int
f : Bool -> Int -- a morphism between Bool and Int
```

¹JAMES IRY A Brief, Incomplete, and Mostly Wrong History of Programming Languages

A *functor* is a mapping between categories that preserves the structure. That is, a functor maps objects and morphisms from one category to objects and morphisms in another category. In the same way, a **Functor** in Haskell maps a type to another type, and its functions to other functions with `fmap`.

```
Maybe :: * -> *
Maybe Bool :: * -- Maps Bool to Maybe Bool
fmap :: (a -> b) -> (f a -> f b) -- Maps (a -> b) to (f a -> f b)
```

Functors can have mappings to other functors, and these are known as *natural transformations*. An *endofunctor* is a *functor* from one category to the same category. That is, it maps objects and morphisms from one category, to objects and morphisms in the same category. But what does that mean in Haskell? Well, since **Functor** maps types to types, we are mapping from `Hask` to `Hask` — all **Functors** in Haskell are endofunctors! They exist as type constructors with the kind `* -> *`.

A *monad* is then defined as:

1. An endofunctor $T : X \rightarrow X$
2. With a natural transformation to flatten two monads together
 $\mu : T(T(X)) \rightarrow T(X)$
3. And another natural transformation $\eta : X \rightarrow T(X)$
4. That satisfy some properties (the monad laws)

Or in Haskell,

```
class Functor m => Monad m a where
  join :: m (m a) -> m a
  return :: a -> m a
```

However as you might be aware, the actual definition in Haskell uses `bind >>=`, not the mathematical definition of `join`. It can be implemented in terms of `>>=` instead, and vice-versa using the fact that a monad is also a functor.

```
(>>=) :: m a -> (a -> m b) -> m b
join m = m >>= \x -> x
m >>= f = join (fmap f m)
```

The “monoid in the category of endofunctors” bit then comes about because a monoid is defined as:

1. A set X
2. With an associative binary operation $\circ : X \times X \rightarrow X$
3. And an identity element $e : X$
4. That satisfy some properties (the monoid laws)

A monoid is the same deal as a monad, except functor composition is replaced with the Cartesian product and the unit function is replaced with the identity element. And if monoids allow chaining together elements, then monads in programming allow chaining together computation. In our system we will be using the more pragmatic Haskell definition of a monad that uses `bind` rather than `join`. But the fundamental idea of sequencing remains the same.

2.3 Type-level programming within GHC

Before creating the system in Chapter 3, we originally attempted to embed a resourceful system within GHC’s type system. Whilst the original Haskell 2010 language specification was based on System F, there has been much work [7][19] carried out to add dependent types to GHC via a plethora of language extensions. Namely, `TypeFamilies` and `GADT` provide a lot of the power needed to carry out type-level programming, and we leveraged this expressiveness to emulate a resource-tracking version of the `IO` monad, dubbed `SubIO`. It carried around a heap with a phantom type parameter:

```
newtype SubIO (a :: [j]) b = SubIO (IO b)
deriving (Functor, Applicative, Monad)
```

We then defined a subclass of `Monad` that extended it with a sequencing operator and concurrency operator, both which merged the heap parameters together.

```
class Monad (m j) => HeapMonad (m :: [x] -> * -> *) j where
  (>>=) :: m j a -> (a -> m k b) -> m (j ** k) b
  (|||) :: m j a -> m k b -> m (j ** k) b
  (|||) x f = x >>= const f
```

It merged the two heaps together through the `**` type family, which threw a type error if they overlapped.

```
type family Overlap a b :: Bool where
  Overlap '[] b = 'False
  Overlap (a ': as) b = If (MemberP a b) 'True (Overlap as b)
```

```
type family a ** b :: [c] where
  (a ** b) = If (Overlap a b)
    (TypeError ('Text "Heaps overlap!"))
    (a :++ b)
```

An instance for `SubIO` was given, and with it some file operations from `System.IO`, transformed to accept the file argument as a type argument, so it could be tracked in the heap.

```
instance HeapMonad SubIO j where
  (SubIO x) >>= f = SubIO (x >>= \z -> let (SubIO y) = f z in y)
  (SubIO x) ||| (SubIO y) = SubIO (forkIO (x >> pure ()) >> y)
```

```

readFile :: forall filePath. KnownSymbol filePath => SubIO '[filePath]
  ↳ String
readFile = fileOp Prelude.readFile

appendFile :: forall filePath. KnownSymbol filePath => String -> SubIO
  ↳ '[filePath] ()
appendFile x = fileOp (`Prelude.appendFile` x)

writeFile :: forall filePath. KnownSymbol filePath => String -> SubIO
  ↳ '[filePath] ()
writeFile x = fileOp (`Prelude.writeFile` x)

fileOp :: forall filePath a. KnownSymbol filePath => (FilePath -> IO a)
  ↳ -> SubIO '[filePath] a
fileOp f = let fp = symbolVal (Proxy :: Proxy filePath)
           in SubIO (f fp)

```

All together, this meant that functions could be written with the familiar `do` syntax, and even better run concurrently with `|||` — provided their heaps didn't overlap.

```

someIO = do
  writeFile @"foo.txt" "hello"
  s <- readFile
  return s

someMoreIO x = do
  appendFile @"bar.txt" x
  readFile >>= SubIO . putStrLn

concExample = runSubIO $ someIO ||| someMoreIO "blah"
--- this *should* type error
-- badConcExample = runSubIO $ someIO ||| someIO

```

However we ultimately discovered that overlapping instances don't always throw a type error, as the phantom type parameter keeping track of the heap was evaluated **lazily**. One workaround involved lowering the heap down to the term level and forcing evaluation via `Data.Proxy` and `seq`. Still, we think there are most likely better ways to embed this within GHC's type system.

2.4 Hindley-Damas-Milner

One of our aims whilst developing this type system was to see how well it integrates with existing functional programming languages, specifically those in the ML family. Therefore we based our system on the Hindley-Damas-Milner

(HDM) type system [4]. Originally designed to formalise the type system of ML, it is one of the first formalisations of a polymorphic type system. It was heavily influential at the time and still continues to be so today, paving the way for newer type systems such as System F.

2.4.1 Syntax

Before we can talk about a type system, we need to talk about the language that it operates on. HDM uses an applicative language: A language in which you can apply abstractions.

$\langle \text{expression } e \rangle ::= x \mid \lambda x.e \mid e e' \mid \text{let } x = e' \text{ in } e$

$\langle \text{type } \tau \rangle ::= \square \mid \alpha \mid \tau' \rightarrow \tau$

$\langle \text{type scheme } \sigma \rangle ::= \tau \mid \forall \alpha.\sigma$

$\langle \text{context } \Gamma \rangle ::= \cdot \mid \Gamma, x : \sigma$

Expressions are an extension of the venerable lambda calculus, with the addition of a new let expression that binds an expression to a variable. This is notably different from abstraction as it allows for polymorphism, which we will see later.

Types are either unit types, type variables or function types. However expressions are not assigned types directly, instead they are given *type schemes* which quantify over type variables. The distinction between the two is necessary so that quantifiers can only appear at the top level.

A *context* is a map from variables to type schemes, represented as a kind of linked list. Whenever we need to work out the type of a variable expression like x , we traverse the context to find its type scheme.

There is also the notion of *free type variables*, which are type variables *inside* a *type scheme* which have not been bound (quantified over).

$$\begin{aligned} \text{FTV}(\forall \alpha.\tau) &= \text{FTV}(\tau) \setminus \{\alpha\} \\ \text{FTV}(\alpha) &= \{\alpha\} \\ \text{FTV}(\tau \rightarrow \tau') &= \text{FTV}(\tau') \cup \text{FTV}(\tau) \end{aligned}$$

2.4.2 Static semantics

A type system describes how we assign *valid* types to our expressions. Most of the time, it's through a typing relation like

$$\Gamma \vdash e : \tau$$

You can read this as “*In the context* Γ , e *has the type* τ ”. This is a relation between a context, an expression and a type, in the same way that $a \leq b$ is a relation between two numbers. With relations we can form *judgements*, such as $1 \leq 2$, or $\Gamma \vdash \lambda x.(y x) : \alpha \rightarrow \tau$. These are just statements that we can make —

they might be true, they might be false. Is $42 \leq 19$? Does $\Gamma \vdash z : \beta$? I don't know, you tell me. We need something else to be able to tell whether or not they make sense, so we have *typing rules* that dictate what typing judgements we can prove. A typing judgement that we can prove tells us that the expression in it is *well typed*, and that it does indeed have the type it claims to have.

$$\frac{x : \sigma \in \Gamma \quad \sigma > \tau}{\Gamma \vdash x : \tau} \text{VAR} \qquad \frac{\Gamma \vdash e : \tau' \rightarrow \tau \quad \Gamma \vdash e' : \tau'}{\Gamma \vdash e e' : \tau} \text{APP}$$

$$\frac{\Gamma, x : \tau' \vdash e : \tau}{\Gamma \vdash \lambda x. e : \tau' \rightarrow \tau} \text{ABS} \qquad \frac{\Gamma \vdash e' : \tau' \quad \Gamma, x : \bar{\Gamma}(\tau') \vdash e : \tau}{\Gamma \vdash \text{let } x = e' \text{ in } e : \tau} \text{LET}$$

These rules consist of some *premises* above a line, and a *conclusion* below it. Premises and conclusions are just other judgements, that we have been able to prove. The gist of these rules is that if you have proof of all the premises above, then you can infer the conclusion at the bottom. So for example, you can read APP as “If e has the type $\tau' \rightarrow \tau$ in Γ and e' has the type τ' also in Γ , then e' applied to e has the type τ in Γ .”

APP is one of the four typing rules in the syntax-directed HDM type system, and tells us what happens to the types when we apply an argument to a function. Let's take a look at what the others mean. ABS relates to functions, sometimes called abstractions. It says if e has the type τ in the context Γ , *extended* with x having the type scheme τ' , then $\lambda x. e$ has the type $\tau' \rightarrow \tau$ in Γ . In other words, if e can have type τ provided it has access to $x : \tau'$, then we can make an abstraction out of it.

If we have a variable expression x , then VAR tells us how we can get the type for it. First, we need to make sure x exists in the context Γ . It will have some type scheme σ , but we can't directly assign that to an expression — the typing relation assigns types to terms, not type schemes. Instead, $\sigma > \tau$ says there needs to be a type τ that σ can be *instantiated* to. This instantiation relation says that if σ is a type scheme $\forall \alpha_1 \dots \alpha_n. \tau'$, then there exists a mapping of type variables to types $\alpha_1 \mapsto \tau_1, \dots, \alpha_n \mapsto \tau_n$ which we can apply to those bound type variables in τ' to give us τ .

As it turns out, these mappings from type variables to types are really common, and they are called *substitutions*. The whole rule then means, if we can look up a type scheme σ for x , and then instantiate it to some type τ , then we can infer that x has the type τ in Γ .

The LET rule introduces polymorphism to the language. That is, it allows one expression to be used for multiple types. Consider the identity function

$$\lambda x. x$$

What type should it have? It could be used both as a $\text{Int} \rightarrow \text{Int}$. But the reality is that we don't care about the underlying type, and want it to work on *anything*.

$$\Gamma \vdash \lambda x. x : \alpha \rightarrow \alpha$$

With type variables instead, we have abstracted over the concrete types. Assuming we have a context like $\Gamma = \bullet, a : \text{Int}$, we might try to use it multiple times, each with different types, by passing it in as an argument to abstraction.

$$\Gamma \vdash (\lambda y.(y y) (y a)) (\lambda x.x) : \text{Int}$$

But this will not work, and it is in fact ill-typed. Our identity function will get passed around as $y : \alpha \rightarrow \alpha$ through ABS. When either $(y y)$ or $(y a)$ try to look up y in VAR, they will find that they can't instantiate it to the type they want, because α is not quantified over!

Instead, we need to pass around the identity function as $y : \forall \alpha. \alpha \rightarrow \alpha$. With a let expression, we can instead say

$$\Gamma \vdash \text{let } y = \lambda x.x \text{ in } (y y) (y a) : \text{Int}$$

If we look at the first premise for LET, we have proof that $\Gamma \vdash \lambda x.x : \alpha \rightarrow \alpha$. In the second premise, we put it into the context as $y : \bar{\Gamma}(\alpha \rightarrow \alpha)$. $\bar{\Gamma}(\alpha \rightarrow \alpha)$ is the close function, which is defined as

$$\bar{\Gamma}(\tau) = \forall \alpha_1 \dots \alpha_n. \tau \text{ where } \{\alpha_1, \dots, \alpha_n\} = \text{FTV}(\tau) \setminus \text{FTV}(\Gamma)$$

This will take the dangling free type variables in our type and create a new type scheme that binds them by quantifying over them: It *closes* over free type variables.

$$\bar{\Gamma}(\alpha \rightarrow \alpha) = \forall \alpha. \alpha \rightarrow \alpha$$

And now that the type variable α is quantified, VAR is able to instantiate y to *both* an $(\alpha \rightarrow \alpha) \rightarrow (\alpha \rightarrow \alpha)$ and an $\text{Int} \rightarrow \text{Int}$.

Syntax-directed

Note that we have been using the syntax-directed rules: a more modern treatment of the original system. In the syntax-directed rules, the typing relation assigns types to terms [17, p.15], not type schemes as in [4]. The original system had six rules:

$$\begin{array}{c} \frac{x : \sigma \in \Gamma}{\Gamma \vdash x : \sigma} \text{TAUT} \quad \frac{\Gamma \vdash e : \sigma \quad \sigma > \sigma'}{\Gamma \vdash e : \sigma'} \text{INST} \quad \frac{\Gamma \vdash e : \sigma \quad \alpha \notin \text{FTV}(\Gamma)}{\Gamma \vdash e : \forall \alpha. \sigma} \text{GEN} \\ \\ \frac{\Gamma \vdash e : \tau' \rightarrow \tau \quad \Gamma \vdash e' : \tau'}{\Gamma \vdash e e' : \tau} \text{COMB} \quad \frac{\Gamma, x : \tau' \vdash e : \tau}{\Gamma \vdash \lambda x. e : \tau' \rightarrow \tau} \text{ABS} \\ \\ \frac{\Gamma \vdash e : \sigma \quad \Gamma, x : \sigma \vdash e' : \tau}{\Gamma \vdash \text{let } x = e \text{ in } e' : \tau} \text{LET} \end{array}$$

As we can see, the syntax-directed rules have merged TAUT and INST into VAR, and GEN into LET via $\bar{\Gamma}$. It is possible to prove that these two systems are equivalent, but the benefit of having just four rules is that we now have exactly one rule for each syntactic form of expression. This means that *the shape of the proof is identical to the shape of the syntax*. This is discussed in more detail in Section 6.4

2.4.3 Dynamic Semantics

You might have noticed that the previous section was titled *Static Semantics*. Whenever we defined the typing relation and the typing rules, we assigned a sense of meaning to these types. For example, something of type $\tau' \rightarrow \tau$ is going to be a function, because our ABS rule dictates so. These semantics are known *statically*. Without knowing what any of our programs evaluate to, we can work out their types.

Usually want to be able to show that these static semantics are *sound*, which in essence means if a program has a type τ , it actually evaluates to something of type τ ². But to show this, we first need to define what it means to evaluate a program. This is called the *dynamic semantics* of a language.

There are multiple ways to define the dynamic semantics. In the original work, Milner and Damas defined a *denotational* semantics for the language [11, 3], which describes evaluation in terms of mathematical objects and functions operating on them. For brevity, we will omit the details here, and relegate them to Appendix A. The important part is that once the dynamic semantics are in place, it is possible to define semantic entailment:

$$\Gamma \vDash e : \tau$$

Which can be read as, “*In the context Γ , e evaluates to a value in τ* ”. Soundness then boils down to showing that if e has type τ , e does actually evaluate to a value in τ .

$$\Gamma \vdash e : \tau \rightarrow \Gamma \vDash e : \tau$$

²Type soundness is discussed in more detail in Section 4.1

Chapter 3

The Resourceful System

In this chapter we will formalise our resource-tracking, monadic IO language, based on the Hindley-Damas-Milner type system that we saw in the previous chapter. Just like before we will define the syntax, static semantics and dynamic semantics, but to avoid ambiguity we will also redefine concepts from HDM in more detail.

3.1 Syntax

We begin by giving the grammar as shown in figure [3.1](#). The expression language is an extension of the lambda calculus with let polymorphism. As before we use x , $\lambda x.e$ and $e e'$ to refer to variable lookup, abstraction and application respectively, and $\text{let } x = e \text{ in } e'$ for polymorphically binding e to x inside e' .

Our language also contains product types, $\tau \times \tau'$, which are the same as tuples (\mathbf{a}, \mathbf{b}) in Haskell. New product types can be introduced with $e \times e'$, and can similarly be eliminated with the projection expressions $\pi_1 e$ and $\pi_2 e$. \square is the unit type $()$ in Haskell), and has the eponymous constructor \square .

More interestingly we have the monadic additions. Unlike Krishnaswami [\[9\]](#) we do not separate the language into expression and computation languages. $\llbracket e \rrbracket$, lifts a regular expression into the IO monad (`return` in Haskell), and $e \gg e'$ is the standard monadic sequencing or binding operation.

Resources r represent something that we want to keep track of and prevent from being accessed concurrently, for example a file system or database. The resources we chose for the grammar are such examples.

Heaps ρ are then formed from resources. They differ from the notion of heaps in separation logic, and are constructed from either single resources or by merging other heaps. They are used to tag the resources that a computation in an IO monad might be accessing. There is also the notion of a **World** heap which encapsulates all possible resources and sub-heaps. It can be thought of as the ordinary **IO** monad in Haskell, where the outside world is one indivisible resource causing everything to be executed sequentially.

type τ	$::= \square \mid \alpha \mid \tau \rightarrow \tau' \mid \tau \times \tau' \mid \text{IO}_\rho \tau$
type scheme σ	$::= \forall \alpha. \sigma \mid \tau$
context Γ	$::= \cdot \mid \Gamma, x : \sigma$
expression e	$::= \square \mid e \times e' \mid \pi_1 e \pi_2 e$ $\mid x \mid \lambda x. e \mid e e' \mid \text{let } x = e \text{ in } e'$ $\mid \llbracket e \rrbracket \mid \llbracket e \rrbracket_r \mid e \ggg e' \mid e \vee e'$
resource r	$::= \text{File} \mid \text{Network} \mid \text{Database} \mid \dots$
heap ρ	$::= r \mid \rho \cup \rho' \mid \text{World}$

Figure 3.1: The grammar for our resourceful language.

$\llbracket e \rrbracket_r$, pronounced *e using r*, lifts a value into the IO monad similarly to $\llbracket e \rrbracket$, but creates a new heap with the resource r and tags the monad with it. Novel to our language is the $e \vee e'$ operator. It joins two monadic computations together into one that uses both their resources, returning a new heap made by merging e and e' 's heaps — provided that they do not overlap and are well formed. Heap well formedness will be defined later on.

Free type variables are defined on types, type schemes and contexts as follows:

$$\begin{aligned}
\text{FTV}(\square) &= \emptyset \\
\text{FTV}(\alpha) &= \{\alpha\} \\
\text{FTV}(\tau \rightarrow \tau') &= \text{FTV}(\tau) \cup \text{FTV}(\tau') \\
\text{FTV}(\text{IO}_\rho \tau) &= \text{IO}_\rho \text{FTV}(\tau) \\
\text{FTV}(\forall \alpha_1, \dots, \alpha_n. \tau) &= \text{FTV}(\tau) - \{\alpha_1, \dots, \alpha_n\} \\
\text{FTV}(\Gamma) &= \bigcup_{x:\sigma \in \Gamma} \text{FTV}(\sigma)
\end{aligned}$$

Free variables are defined on terms — be careful not to mix these up with free type variables!

$$\begin{aligned}
\text{FV}(x) &= \{x\} & \text{FV}(\lambda x.e) &= \text{FV}(e) \setminus \{x\} \\
\text{FV}(e \ e') &= \text{FV}(e) \cup \text{FV}(e') & \text{FV}(\text{let } x = e \text{ in } e') &= \text{FV}(e) \cup (\text{FV}(e') \setminus \{x\}) \\
\text{FV}(\llbracket e \rrbracket) &= \text{FV}(e) & \text{FV}(e \ggg e') &= \text{FV}(e) \cup \text{FV}(e') \\
\text{FV}(\square) &= \emptyset & \text{FV}(\llbracket e \rrbracket_r) &= \text{FV}(e) \\
\text{FV}(e_1 \times e_2) &= \text{FV}(e_1) \cup \text{FV}(e_2) & \text{FV}(\pi_i e) &= \text{FV}(e)_{i=1,2} \\
\text{FV}(e_1 \vee e_2) &= \text{FV}(e_1) \cup \text{FV}(e_2)
\end{aligned}$$

3.1.1 Substitution

Substitution may seem self-explanatory, but it is important we define it crystal clear. It plays a vital role in proving properties about our system, and there exists subtle differences with how it is defined on terms, types and type schemes. A substitution is a map from type variables to types, written as $\tau'[\tau/\alpha]$ to replace the type τ for the type variable α in the type τ' .

$$\begin{aligned}
\square[\tau/\alpha] &= \square \\
\alpha'[\tau/\alpha] &= \begin{cases} \tau & \text{if } \alpha' = \alpha \\ \alpha' & \text{otherwise} \end{cases} \\
(e \times e')[\tau/\alpha] &= e[\tau/\alpha] \times e'[\tau/\alpha] \\
(\tau_1 \rightarrow \tau_2)[\tau/\alpha] &= (\tau_1[\tau/\alpha] \rightarrow \tau_2[\tau/\alpha]) \\
(\text{IO}_\rho \tau')[\tau/\alpha] &= \text{IO}_\rho \tau'[\tau/\alpha]
\end{aligned}$$

Substitution is associative, and a substitution that substitutes multiple type variables at once can be written like $[\tau_1/\alpha_1, \dots, \tau_m/\alpha_m]$. It is also extended to type schemes, but note that this only substitutes *free* type variables.

$$\begin{aligned}
\forall \alpha'. \sigma[\tau/\alpha] &= \begin{cases} \forall \alpha'. \sigma & \text{if } \alpha' = \alpha \\ \forall \alpha'. \sigma[\tau/\alpha] & \text{otherwise} \end{cases} \\
\tau[\tau/\alpha] &= \tau[\tau/\alpha] \text{ (Substitution on type)}
\end{aligned}$$

It should not be confused with instantiation, where the *bound* type variables $\forall \alpha_1, \dots, \alpha_n$ are substituted inside the type of a type scheme. Instantiation is not a function though, it is a relation $\sigma > \tau$ in which we say σ can be instantiated to τ .

$$\boxed{\sigma > \tau}$$

$$\frac{\text{dom}(s) = \{\alpha_1, \dots, \alpha_n\} \quad \tau'[s] = \tau}{\forall \alpha_1, \dots, \alpha_n. \tau' > \tau}$$

The domain of a substitution, $\text{dom}(s)$, is the set of type variables that it will replace, for example:

$$\text{dom}([\tau_1/\alpha, \tau_2/\beta]) = \{\alpha, \beta\}$$

Instantiation can then be read as, if there exists a substitution s that substitutes exactly all the bound type variables in the type scheme $\forall\alpha_1, \dots, \alpha_n.\tau'$ to give the type τ , then $\forall\alpha_1, \dots, \alpha_n.\tau' > \tau$.

Furthermore we define the relation $\sigma > \sigma'$ on type schemes as well, and say that σ is *more general than* σ' if for all τ , $\sigma' > \tau \rightarrow \sigma > \tau$.

Substitution (not instantiation) is also defined on contexts. Substitution on contexts only substitutes *free* type variables.

$$\begin{aligned} \cdot[\tau/\alpha] &= \cdot \\ (\Gamma, x : \sigma)[\tau/\alpha] &= \Gamma[\tau/\alpha], x : \sigma[\tau/\alpha] \end{aligned}$$

There also exist term-level substitutions which map variables to other terms, and can be applied as follows:

$$\begin{aligned} \square[v/\alpha] &= \square \\ \alpha'[v/\alpha] &= \begin{cases} v & \text{if } \alpha' = \alpha \\ \alpha' & \text{otherwise} \end{cases} \\ (e \ e')[v/\alpha] &= e[v/\alpha] \ e'[v/\alpha] \\ \lambda x.e[v/\alpha] &= \begin{cases} \lambda x.e & \text{if } x = \alpha \\ \lambda x.(e[v/\alpha]) & \text{otherwise} \end{cases} \\ \text{let } x = e' \text{ in } e[v/\alpha] &= \begin{cases} \text{let } x = (e'[v/\alpha]) \text{ in } e & \text{if } x = \alpha \\ \text{let } x = (e'[v/\alpha]) \text{ in } (e[v/\alpha]) & \text{otherwise} \end{cases} \\ (e_1 \times e_2)[v/\alpha] &= e_1[v/\alpha] \times e_2[v/\alpha] \\ (\pi_i \ e)[v/\alpha] &= \pi_i(e[v/\alpha])_{i=1,2} \\ \llbracket e \rrbracket[v/\alpha] &= \llbracket e[v/\alpha] \rrbracket \\ \llbracket e \rrbracket_r[v/\alpha] &= \llbracket e[v/\alpha] \rrbracket_r \\ (e \gg\!\!\gg e')[v/\alpha] &= e[v/\alpha] \gg\!\!\gg e'[v/\alpha] \\ (e_1 \ \vee \ e_2)[v/\alpha] &= e_1[v/\alpha] \ \vee \ e_2[v/\alpha] \end{aligned}$$

3.1.2 Barendregt's variable convention

Sometimes, we will end up in a scenario with two separate expressions such as $\lambda x.x$ and $(y \ x)$. We know that the x inside the first expression is distinct from the x in the second expression, but when working with proofs we will need to show this somehow. We will use the Barendregt variable convention [\[1\]](#) to deal with this: If we want to show that the x inside $\lambda x.x$ is different from the x in $(y \ x)$, we can use *alpha conversion* to rename x to z and get $\lambda z.z$. This new expression is in fact equivalent to $\lambda x.x$, and we can always choose a new

unique name to avoid collisions with any *free variables*. The Barendregt variable convention says that whenever we have a bound variable (an x inside a $\lambda x.e$ or let $x = e'$ in e), we can just assume that we have performed alpha conversion on it so that the variable name is unique.

3.2 Static semantics

Our static semantics begin with the syntax-directed rules of the Hindley-Damas-Milner system for the typing relation $\Gamma \vdash e : \tau$.

$$\frac{x : \sigma \in \Gamma \quad \sigma > \tau}{\Gamma \vdash x : \tau} \text{VAR} \qquad \frac{\Gamma \vdash e : \tau' \rightarrow \tau \quad \Gamma \vdash e' : \tau'}{\Gamma \vdash e e' : \tau} \text{APP}$$

$$\frac{\Gamma, x : \tau' \vdash e : \tau}{\bar{\Gamma} \vdash \lambda x.e : \tau' \rightarrow \tau} \text{ABS} \qquad \frac{\Gamma \vdash e' : \tau' \quad \Gamma, x : \bar{\Gamma}(\tau') \vdash e : \tau}{\Gamma \vdash \text{let } x = e' \text{ in } e : \tau} \text{LET}$$

Like before, $\bar{\Gamma}$ is defined as

$$\bar{\Gamma}(\tau) = \forall \alpha_1 \dots \alpha_n. \tau \text{ where } \{\alpha_1, \dots, \alpha_n\} = \text{FTV}(\tau) \setminus \text{FTV}(\Gamma)$$

The typing rule for \square expressions introduces the monomorphic type, and the rule for products introduces $e \times e'$.

$$\frac{}{\Gamma \vdash \square : \square} \text{UNIT} \qquad \frac{\Gamma \vdash e : \tau \quad \Gamma \vdash e' : \tau'}{\Gamma \vdash e \times e' : \tau \times \tau'} \text{PRODUCT}$$

Product types also have two eliminators, which project out the inner type.

$$\frac{\Gamma \vdash e : \tau \times \tau'}{\Gamma \vdash \pi_1 e : \tau} \text{PROJ1} \qquad \frac{\Gamma \vdash e : \tau \times \tau'}{\Gamma \vdash \pi_2 e : \tau'} \text{PROJ2}$$

Now we introduce the monadic parts of the language. Our language only has one type of monad, the IO monad, which is parameterised by both its *heap* ρ and its encapsulated type. Monadic values can be introduced into the language with $\llbracket e \rrbracket$, which lifts a pure term into **any well formed** heap.

$$\frac{\Gamma \vdash e : \tau \quad \text{ok } \rho}{\Gamma \vdash \llbracket e \rrbracket : \text{IO}_\rho \tau} \text{LIFT}$$

We need to be careful what heaps we allow terms to be lifted into, as the entire point of this system is to avoid heaps containing duplicate resources. It would be all too easy to introduce a nonsensical type like $\text{IO}_{\text{File} \cup \text{File}} \tau$ with LIFT, if it were not for the premise $\text{ok } \rho$.

$\text{ok } \rho$ is a new relation we define in Figure 3.2 to establish what heaps we consider to be well formed, in a similar vein to Krishnaswami 9. All heaps

$$\boxed{\text{ok } \rho} \quad \overline{\text{ok World}} \quad \overline{\text{ok } r} \quad \frac{\text{ok } \rho \quad \text{ok } \rho' \quad \rho \cap \rho' = \emptyset}{\text{ok } \rho \cup \rho'}$$

Figure 3.2: Rules for heap well formedness.

consisting of a single resource are well formed, as well as heaps made by merging two other well formed heaps *that are distinct*, i.e. they do not share any of the same resources.

Whilst LIFT lets us construct monads in any heap, we will eventually want to have constructors for IO monads that use specific resources. When first designing the system, we used placeholder functions that just used fixed resources:

$$\overline{\Gamma \vdash \text{readFile} : \text{IO}_{\text{File}} \square} \quad \overline{\Gamma \vdash \text{readNetwork} : \text{IO}_{\text{Network}} \square}$$

`readFile` and `readNet` are examples of typical operations that can consume a specific resource — their heap consists of just a single resource. This was then generalised to $\llbracket e \rrbracket_r$, which lifts any pure term into an IO monad with a heap consisting of the resource r .

$$\frac{\Gamma \vdash e : \tau}{\Gamma \vdash \llbracket e \rrbracket_r : \text{IO}_r e} \text{USE}$$

Note that it only uses a single resource, not any arbitrary heap. A heap with just one resource is always well formed, so there is no need for an `ok` r premise. `USE` is meant to be used to annotate terms with that a computation needs for itself, whilst on the other hand `LIFT` is for bringing pure expressions into a monadic computation.

Once we have an IO value, we can sequence computation by binding it with a function that returns another IO.

$$\frac{\Gamma \vdash e : \text{IO}_{\rho} \tau' \quad \Gamma \vdash e' : \tau' \rightarrow \text{IO}_{\rho} \tau}{\Gamma \vdash e \gg e' : \text{IO}_{\rho} \tau} \text{BIND}$$

Note that the types of the two IOs must use the same resources. However, we want expressions such as

$$\llbracket \square \rrbracket_{\text{File}} \gg \lambda x. \llbracket \square \rrbracket_{\text{Net}}$$

to also be well typed. In particular, we want the above expression to be of type $\text{IO}_{\text{File} \cup \text{Net}} \square$. But if `BIND` requires the heaps to be the same, we must first somehow “cast” `readFile` and `readNet` to the same type. This is the purpose of the `SUB` (subsumption) rule:

$$\frac{\Gamma \vdash e : \text{IO}_{\rho'} \tau \quad \rho' \geq \rho \quad \text{ok } \rho}{\Gamma \vdash e : \text{IO}_{\rho} \tau} \text{SUB}$$

$$\boxed{\rho' \geq: \rho}$$

$$\frac{}{\text{World} \geq: \rho} \text{TOP} \quad \frac{}{\rho \geq: \rho} \text{REFL} \quad \frac{\rho' \geq: \rho}{\rho' \cup \rho'' \geq: \rho} \text{UNIONL} \quad \frac{\rho' \geq: \rho}{\rho'' \cup \rho' \geq: \rho} \text{UNIONR}$$

Figure 3.3: Subheap rules




Figure 3.4: An example of some heaps and their subheap orderings.

It lets something of type $\text{IO}_{\rho'}\tau$ become a $\text{IO}_{\rho}\tau$, provided that the heap ρ is a *subheap* of ρ' . The subheap rules for heaps, shown in figure 3.3, define the $\geq:$ relation.

Intuitively, a heap σ can thought of being a subheap of another heap σ' , if σ' subsumes σ , similarly to how subtyping works. For example, $\text{Net} \geq: \text{Net} \cup \text{File}$, since $\text{Net} \cup \text{File}$ “overlaps” with the heap Net .

If one views this relation as an ordering, then we have World defined as the least upper bound — this represents using all possible resources, and as mentioned earlier IO_{World} can be thought of as the *IO* monad in Haskell, where sequencing interacts with the state of the entire world. In a sense this could be viewed as a form of subtyping, but by constraining it to just resources and not actual types, we avoid the extra overhead and complexity subtyping would normally give us in the presence of type inference 5.

The subheaping relation is both reflexive (by definition), and transitive.

Theorem 1. *For all a, b, c , if $c \geq: b$ and $b \geq: a$ then $c \geq: a$.*

Proof. By induction on $c \geq: b$. See proof B.1 in the appendix. \square

The purpose of this type system is to allow programs to be run concurrently, but reject the ones that concurrently access the same resource. In this regard, the CONC rule is the heart and soul of the system. It takes two monadic expressions, merges their heaps and returns a product of their two inner types, inside IO .

$$\frac{\Gamma \vdash e_1 : \text{IO}_{\rho_1}\tau_1 \quad \Gamma \vdash e_2 : \text{IO}_{\rho_2}\tau_2 \quad \text{ok } \rho_1 \cup \rho_2}{\Gamma \vdash e_1 \Upsilon e_2 : \text{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2} \text{CONC}$$

$$\boxed{\Gamma \vdash e : \tau}$$

$$\frac{x : \sigma \in \Gamma \quad \sigma > \tau}{\Gamma \vdash x : \tau} \text{VAR} \quad \frac{\Gamma \vdash e : \tau' \rightarrow \tau \quad \Gamma \vdash e' : \tau'}{\Gamma \vdash e e' : \tau} \text{APP} \quad \frac{\Gamma, x : \tau \vdash e : \tau'}{\Gamma \vdash \lambda x. e : \tau \rightarrow \tau'} \text{ABS}$$

$$\frac{\Gamma \vdash e' : \tau' \quad \Gamma, x : \bar{\Gamma}(\tau') \vdash e : \tau}{\Gamma \vdash \text{let } x = e' \text{ in } e : \tau} \text{LET} \quad \frac{}{\Gamma \vdash \square : \square} \text{UNIT}$$

$$\frac{\Gamma \vdash e : \tau \quad \Gamma \vdash e' : \tau'}{\Gamma \vdash e \times e' : \tau \times \tau'} \text{PRODUCT} \quad \frac{\Gamma \vdash e : \tau \times \tau'}{\Gamma \vdash \pi_1 : \tau} \text{PROJ1} \quad \frac{\Gamma \vdash e : \tau \times \tau'}{\Gamma \vdash \pi_2 : \tau'} \text{PROJ2}$$

$$\frac{\Gamma \vdash e : \tau \quad \text{ok } \rho}{\Gamma \vdash \llbracket e \rrbracket : \text{IO}_\rho \tau} \text{LIFT} \quad \frac{\Gamma \vdash e : \tau}{\Gamma \vdash \llbracket e \rrbracket_r : \text{IO}_r \tau} \text{USE}$$

$$\frac{\Gamma \vdash e : \text{IO}_\rho \tau' \quad \Gamma \vdash e' : \tau' \rightarrow \text{IO}_\rho \tau}{\Gamma \vdash e \gg e' : \text{IO}_\rho \tau} \text{BIND}$$

$$\frac{\Gamma \vdash e_1 : \text{IO}_{\rho_1} \tau_1 \quad \Gamma \vdash e_2 : \text{IO}_{\rho_2} \tau_2 \quad \text{ok } \rho_1 \cup \rho_2}{\Gamma \vdash e_1 \curlywedge e_2 : \text{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2} \text{CONC}$$

$$\frac{\Gamma \vdash e : \text{IO}_{\rho'} \tau \quad \rho' \geq \rho \quad \text{ok } \rho}{\Gamma \vdash e : \text{IO}_\rho \tau} \text{SUB}$$

Figure 3.5: The typing rules for our resourceful language.

The premises ensure that the new merged heap must be well formed, as we do not want to allow programs that try to use the same resource concurrently, such as

$$\llbracket \square \rrbracket_{\text{File}} \curlywedge \llbracket \square \rrbracket_{\text{File}}$$

We do however, want to allow programs that run two expressions that do not share any resources:

$$\llbracket \square \rrbracket_{\text{File}} \curlywedge \llbracket \square \rrbracket_{\text{Net}} : \text{IO}_{\text{File} \cup \text{Net}}$$

All together these rules define the typing relation, and are displayed in full in Figure 3.5. To give a better idea of how they are used in practice, we will look at some examples of various programs and their types.

3.2.1 Examples

Monadic binding

Here we show a proof tree for the expression

$$\text{let } x = \lambda y. \llbracket \square \rrbracket \text{ in } (x \square) \gg \lambda z. \llbracket z \rrbracket_{\text{File}}$$

The lift operator in $\lambda y. \llbracket \square \rrbracket$ can lift into any heap. But the smallest heap we can use for this overall program is `File`, which is needed by $\lambda z. \llbracket z \rrbracket_{\text{File}}$. So when we choose a heap to lift \square into, we lift it into `File`.

$$\frac{\frac{\frac{\bullet, y : \alpha \vdash \square : \square}{\bullet, y : \alpha \vdash \llbracket \square \rrbracket : \text{IO}_{\text{File}} \square} \text{ok File}}{\bullet \vdash \lambda y. \llbracket \square \rrbracket : \alpha \rightarrow \text{IO}_{\text{File}} \square} \text{ (1)}}{\frac{\frac{\frac{\bullet, x : \alpha \rightarrow \text{IO}_{\text{File}} \square \vdash x : \alpha \rightarrow \text{IO}_{\text{File}} \square}{\bullet, x : \alpha \rightarrow \text{IO}_{\text{File}} \square \vdash \square : \square}}{\bullet, x : \alpha \rightarrow \text{IO}_{\text{File}} \square \vdash x \square : \text{IO}_{\text{File}} \square} \quad \frac{\frac{z : \alpha \in \bullet, x : \alpha \rightarrow \text{IO}_{\text{File}} \square, z : \alpha \quad \alpha > \square}{\bullet, x : \alpha \rightarrow \text{IO}_{\text{File}} \square, z : \alpha \vdash z : \square}}{\bullet, x : \alpha \rightarrow \text{IO}_{\text{File}} \square, z : \alpha \vdash \llbracket z \rrbracket_{\text{File}} : \text{IO}_{\text{File}} \square}}{\bullet, x : \alpha \rightarrow \text{IO}_{\text{File}} \square \vdash \lambda z. \llbracket z \rrbracket_{\text{File}} : \text{IO}_{\text{File}} \square} \text{ (2)}}{\bullet, x : \alpha \rightarrow \text{IO}_{\text{File}} \square \vdash (x \square) \gg \lambda z. \llbracket z \rrbracket_{\text{File}} : \text{IO}_{\text{File}} \square} \text{ (2)}} \text{ (1) (2)}$$

$$\frac{}{\bullet \vdash \text{let } x = \lambda y. \llbracket \square \rrbracket \text{ in } (x \square) \gg \lambda z. \llbracket z \rrbracket_{\text{File}} : \text{IO}_{\text{File}} \square}$$

Let polymorphism

This example doesn't contain any resourceful elements, but is just an example of how let polymorphism allows the same variable to be given different types at each call site. Note how the premises in (1) and (2) assign x different types.

$$\frac{\frac{\bullet, x : \alpha \rightarrow \alpha \vdash x : (\square \rightarrow \square) \rightarrow (\square \rightarrow \square)}{\bullet, x : \alpha \rightarrow \alpha \vdash x x : \square \rightarrow \square} \text{ (1)}}{\bullet, x : \alpha \rightarrow \alpha \vdash x : \square \rightarrow \square} \quad \frac{\bullet, x : \alpha \rightarrow \alpha \vdash \square : \square}{\bullet, x : \alpha \rightarrow \alpha \vdash x \square : \square} \text{ (2)}}{\bullet \vdash \lambda y. y : \alpha \rightarrow \alpha} \quad \frac{\frac{y : \alpha \in \bullet, y : \alpha \quad \alpha > \alpha}{\bullet, y : \alpha \vdash y : \alpha} \text{ (1) (2)}}{\bullet \vdash \lambda y. y : \alpha \rightarrow \alpha} \quad \frac{\bullet, x : \alpha \rightarrow \alpha \vdash (x x) (x \square) : \square}{\bullet \vdash \text{let } x = \lambda y. y \text{ in } (x x) (x \square) : \square}$$

Concurrency

Here is our first example of accessing two resources concurrently — with a bit of imagination one can think of this as reading a file from a disk whilst

simultaneously fetching data over the network.

$$\frac{\frac{\frac{\cdot, x : \alpha \vdash x : \alpha \quad \text{ok File}}{\cdot, x : \alpha \vdash \llbracket x \rrbracket_{\text{File}} : \text{IO}_{\text{File}} \alpha}}{\cdot \vdash \lambda x. \llbracket x \rrbracket_{\text{File}} : \alpha \rightarrow \text{IO}_{\text{File}} \alpha} \quad \cdot \vdash \square : \square}{\cdot \vdash (\lambda x. \llbracket x \rrbracket_{\text{File}}) \square : \text{IO}_{\text{File}} \square} \quad \frac{\cdot \vdash \square : \square \quad \text{ok Net}}{\cdot \vdash \llbracket \square \rrbracket_{\text{Net}} : \text{IO}_{\text{Net}} \square} \quad \text{File} \cap \text{Net} = \emptyset}{\cdot \vdash (\lambda x. \llbracket x \rrbracket_{\text{File}}) \square \curlywedge \llbracket \square \rrbracket_{\text{Net}} : \text{IO}_{\text{File} \cup \text{Net}} \square \times \square}$$

Subsumption

Assume there is a function inside our context called `writeFile` with the type $\square \rightarrow \text{IO}_{\text{File}} \square$. We can subsume its heap to be part of a larger heap, like $\text{File} \cup \text{Net}$. We will show this with the expression

$$\llbracket \square \rrbracket_{\text{File}} \curlywedge \llbracket \square \rrbracket_{\text{Net}} \gg \gg \lambda x. \text{writeFile} (\pi_1 x)$$

$$\Gamma = \cdot, \text{writeFile} : \square \rightarrow \text{IO}_{\text{File}} \square$$

$$\frac{\frac{\Gamma \vdash \square : \square}{\Gamma \vdash \llbracket \square \rrbracket_{\text{File}} : \text{IO}_{\text{File}} \square} \quad \frac{\Gamma \vdash \square : \square}{\Gamma \vdash \llbracket \square \rrbracket_{\text{Net}} : \text{IO}_{\text{Net}} \square} \quad \text{File} \cap \text{Net} = \emptyset}{\Gamma \vdash \llbracket \square \rrbracket_{\text{File}} \curlywedge \llbracket \square \rrbracket_{\text{Net}} : \text{IO}_{\text{File} \cup \text{Net}} \square} \quad (1)$$

$$\frac{\Gamma, x : \square \times \square \vdash \text{writeFile} : \square \rightarrow \text{IO}_{\text{File}} \square \quad \frac{\Gamma, x : \square \times \square \vdash x : \square \times \square}{\Gamma, x : \square \times \square \vdash \pi_1 x : \square}}{\Gamma, x : \square \times \square \vdash \text{writeFile} (\pi_1 x) : \text{IO}_{\text{File}} \square} \quad (2)$$

$$\frac{(1) \quad \frac{(2) \quad \text{File} \geq : \text{File} \cup \text{Net} \quad \text{ok File} \cup \text{Net}}{\Gamma, x : \square \times \square \vdash \text{writeFile} (\pi_1 x) : \text{IO}_{\text{File} \cup \text{Net}} \square} \text{SUB}}{\Gamma \vdash \lambda x. \text{writeFile} (\pi_1 x) : \text{IO}_{\text{File} \cup \text{Net}} \square}}{\Gamma \vdash \llbracket \square \rrbracket_{\text{File}} \curlywedge \llbracket \square \rrbracket_{\text{Net}} \gg \gg \lambda x. \text{writeFile} (\pi_1 x) : \text{IO}_{\text{File} \cup \text{Net}} \square}$$

3.3 Dynamic semantics

As shown in Chapter 2, the dynamic semantics in the original Hindley-Damas-Milner system was based on denotational semantics. In our type system, we will use operational semantics. Operational semantics are similar to what we have seen before in the definition of the static semantics. We define a bunch of inference rules, and from these build up proofs. Tofte [17] had the idea of using operational semantics for not just the typing rules, but also for the dynamic semantics. We choose this approach over denotational semantics as it unifies our approach to types and evaluation, and as Tofte said, “*it seems a bit unfortunate that we should have to understand domain theory to be able to investigate whether a type inference system admits faulty programs*”.




Figure 3.6: Terminal values

3.3.1 Values

Before we can talk about how we evaluate a program, we need to define what constitutes a fully evaluated program. That is, what terms are a result of a completed computation. We define a unary relation called **value**, and give rules describing what terms can be considered values in Figure 3.6.

For example, we cannot evaluate the unit type any further, therefore all \square s are considered values. The same goes for product types, but only if both inner components are values themselves: $\square \times \square$ is a finished value, but $f e \times \square$ might still have evaluation left to do. A lambda on its own is a value too. Without being applied to an argument it cannot be evaluated any further: The computation inside of it is suspended. In a similar fashion a lifted computation cannot be computed any further *on its own*. We will see later how binding can run this computation, but by itself it will not evaluate to anything.

3.3.2 Small-step semantics

The approach to operational semantics we will be taking is *small-step operational semantics*. In small-step operational semantics we define a step relation $a \rightsquigarrow b$ which says that in one “step”, a reduces to b . b might then go onto reduce further if it is able to, or it could be a finished value. So reduction can be thought of as evaluation of a program, bit by bit. Small-step semantics differs from big-step semantics, where the relation $a \Downarrow b$ says that at the end of the day, a will reduce to b , and b will not reduce any further.

As an example, if an expression e_1 reduces to e'_1 , i.e. $e_1 \rightsquigarrow e'_1$, then we want the application $e_1 e_2$ to reduce as well. We can write this as

$$\frac{e_1 \rightsquigarrow e'_1}{e_1 e_2 \rightsquigarrow e'_1 e_2}$$

We call this type of reduction rule which takes smaller reductions and updates it within a bigger structure, ξ -reduction. There are ξ -reductions for other expressions with structure inside, namely product types, the concurrent operator and the bind operator.

Another type of reduction rule is β -reduction, which comes from the lambda calculus. When we apply an argument to an abstraction, we substitute the bound variable inside the abstraction with the argument. The β -reduction rule

defines this.

$$\frac{\text{value } e_2}{(\lambda x.e_1) e_2 \rightsquigarrow e_1[e_2/x]}$$

It is important to note the premise here that states the argument being applied must be a value. This enforces a strict evaluation order, since in order for the argument to be a value it must be completely reduced. A lazily evaluated semantics might forgo this extra requirement, so that the argument can be reduced after substitution.

As mentioned earlier, lifted expressions $\llbracket e \rrbracket$ are suspended much like lambdas, and so are values since they cannot reduce any further *on their own*. However, with the bind operator, the value inside them can be extracted out and fed into an abstraction.

$$\overline{\llbracket e \rrbracket \gg e' \rightsquigarrow e' e}$$

Unlike the semantic rule for β -reduction, there is no premise enforcing that $\llbracket e \rrbracket$ is a value, since all lifted terms are values anyway.

For an expression lifted into a resourceful IO monad with $\llbracket e \rrbracket_\rho$, one might be tempted to just reduce this to a $\llbracket e \rrbracket$.

$$\overline{\llbracket e \rrbracket_r \rightsquigarrow \llbracket e \rrbracket}$$

And we could then also define the reduction for concurrency like so:

$$\overline{\llbracket v \rrbracket \curlywedge \llbracket w \rrbracket \rightsquigarrow \llbracket v \times w \rrbracket}$$

However, the intermediate $\llbracket e \rrbracket$ can be confusing. The purpose of the lift operator is to lift a pure value into any possible resource bound monad. When we see a lift, we think of the typing judgement `LIFT` that allows it to fit any heap, when in fact a use should restrict what heaps it can go into. For these reasons, we instead define its reduction identically to lift.

$$\overline{\llbracket e \rrbracket_r \gg e' \rightsquigarrow e' e}$$

Concurrency is then defined as chaining together two binds, and returning the lifted product of the two results.

$$\overline{v \curlywedge w \rightsquigarrow v \gg \lambda v.(w \gg \lambda w.\llbracket v \times w \rrbracket)}$$

This might seem like the opposite of concurrency — executing the computation in sequence — but because our monad does not have any state (see Section [6.5](#)),

$$v \gg \lambda v.(w \gg (\lambda w.\llbracket v \times w \rrbracket))$$

is identical to

$$w \gg \lambda w.(v \gg (\lambda v.\llbracket v \times w \rrbracket))$$

In fact the previous definition for concurrency as $\llbracket v \rrbracket \curlywedge \llbracket w \rrbracket \rightsquigarrow \llbracket v \times w \rrbracket$ will have the same reduction steps at the end of the day. In a real-world program-

$$\boxed{e \rightsquigarrow e'}$$

$$\frac{e_1 \rightsquigarrow e'_1}{e_1 e_2 \rightsquigarrow e'_1 e_2} \quad \frac{e_2 \rightsquigarrow e'_2}{e_1 e_2 \rightsquigarrow e_1 e'_2} \quad \frac{\text{value } e'}{(\lambda x.e) e' \rightsquigarrow e[e'/x]}$$

$$\overline{\text{let } x = e' \text{ in } e \rightsquigarrow e[e'/x]}$$

$$\frac{e_1 \rightsquigarrow e'_1}{e_1 \times e_2 \rightsquigarrow e'_1 \times e_2} \quad \frac{e_2 \rightsquigarrow e'_2}{e_1 \times e_2 \rightsquigarrow e_1 \times e'_2} \quad \frac{e \rightsquigarrow e'}{\pi_1 e \rightsquigarrow \pi_1 e'} \quad \frac{e \rightsquigarrow e'}{\pi_2 e \rightsquigarrow \pi_2 e'}$$

$$\overline{\pi_1(e_1 \times e_2) \rightsquigarrow e_1} \quad \overline{\pi_2(e_1 \times e_2) \rightsquigarrow e_2} \quad \frac{e_1 \rightsquigarrow e'_1}{e_1 \gg\gg e_2 \rightsquigarrow e'_1 \gg\gg e_2}$$

$$\overline{\llbracket v \rrbracket \gg\gg e_2 \rightsquigarrow e_2 v} \quad \overline{\llbracket e \rrbracket_r \gg\gg e' \rightsquigarrow e' e}$$

$$\overline{v \Upsilon w \rightsquigarrow v \gg\gg \lambda v.(w \gg\gg \lambda w.\llbracket v \times w \rrbracket)}$$

Figure 3.7: Dynamic Semantics

ming language implementation, the evaluation would actually be implemented concurrently.

Although the $v \Upsilon w$ might evaluate to two separate $\gg\gg$ s, the static semantics of concurrency and binding are different. For example, given $\Gamma \vdash \llbracket v \rrbracket : \text{IO}_\rho \tau$ we are allowed to bind *liftv* with itself:

$$\Gamma \vdash \llbracket v \rrbracket \gg\gg \lambda v.(\llbracket v \rrbracket \gg\gg (\lambda v.\llbracket v \times v \rrbracket)) : \text{IO}_\rho(\tau \times \tau)$$

But with the concurrent operator, this is a type error, since the premise $\text{ok } \rho \cup \rho$ does not hold: There is no type in any context that can be given to $v \Upsilon v$.

3.3.3 The reflexive and transitive closure

We can go a *step* further and define \rightarrow as the reflexive, transitive closure of \rightsquigarrow . What does that mean? If \rightsquigarrow is a binary relation of two terms, then the transitive closure \rightarrow is a new relation that maintains all the relations of \rightsquigarrow and is transitive, i.e. if $a \rightarrow b$ and $b \rightarrow c$, then $a \rightarrow c$. A reflexive transitive closure extends this so that for any a , $a \rightarrow a$.

Intuitively speaking, if the small-step inference rule $a \rightsquigarrow b$ says that a reduces to b in exactly one step, then $a \rightarrow b$ says a reduces to b in zero or more steps. Instead of reducing one step at a time, it can go all the way.



Figure 3.8: Left: the \rightsquigarrow relation. Right: The reflexive, transitive closure of \rightsquigarrow , the reduction relation \twoheadrightarrow .

For example, if we have the expression $(\lambda x.x \square) (\lambda y.y)$ then have the following relations:

$$(\lambda x.x \square) (\lambda y.y) \rightsquigarrow (\lambda y.y \square) \rightsquigarrow \square$$

$$(\lambda x.x \square) (\lambda y.y) \twoheadrightarrow \square$$

We do not use this relation to prove any properties about our system, but it sheds light on how our expression language eventually reduces to a value in the end.

Chapter 4

Properties

A type system is no good unless we can trust that the programs actually evaluate to the type they are given. In this section we will prove a number of lemmas and theorems which will eventually show that the type system is sound.

Before we begin, we need to show a couple of properties about contexts that crop up later on in our proofs. The first states that if we have two variables in the context with the same names but different type schemes, then we can ignore the first one as it is overshadowed by the second — any reference to x will result in $x : \sigma'$.

Lemma 1 (Drop). *If $\Gamma, x : \sigma, x : \sigma' \vdash e : \tau$, then $\Gamma, x : \sigma' \vdash e : \tau$*

We can also say that if two variables x and y are indeed different, we are free to permute them, swapping them about.

Lemma 2 (Swap). *If $x \neq y$ and $\Gamma, x : \sigma, y : \sigma' \vdash e : \tau$, then $\Gamma, y : \sigma', x : \sigma \vdash e : \tau$*

And if we have a typing judgement in an empty context, we can weaken the judgement to extend this to any other context.

Lemma 3 (Weaken). *If $\cdot \vdash e : \tau$, then for Γ , $\Gamma \vdash e : \tau$.*

We will also need to show this property about instantiation and the close function.

Lemma 4. *If $\sigma > \sigma'$ then $\overline{\Gamma, x : \sigma}(\tau) > \overline{\Gamma, x : \sigma'}(\tau)$.*

Proof. We provide an informal proof as follows.

1. If $\sigma > \sigma'$, then for every τ that $\sigma > \tau$, $\sigma' > \tau$.
2. So σ must parameterise over at least the same number of type variables if not *more* than σ' .
3. So σ has at least the same number of free type variables if not *less* than σ' .

4. By definition of the close function, $\overline{\Gamma, x : \sigma}(\tau)$ will then have at least the same number of bound type variables if not *more* than $\overline{\Gamma, x : \sigma'}(\tau)$.
5. So if $\overline{\Gamma, x : \sigma'}(\tau)$ can be instantiated to some τ , then $\overline{\Gamma, x : \sigma}(\tau)$ can also be instantiated to that τ , as it has enough bound type variables to handle everything the former can — if it had less type variables than the former, then it would not be able to instantiate these types, but for the converse excess type variables can be mapped to whatever.

□

Now we prove the generalisation theorem, which states that if a type scheme inside the context is an instantiation of another type scheme, then we can use the more general type scheme and preserve how expressions are typed. This is an adaptation of a lemma from Wright and Felleisen [20, Lemma 4.6], which in turn is an adaptation of a lemma from Damas and Milner.

Theorem 2 (Generalisation). *If $\Gamma, x : \sigma' \vdash e : \tau$ and $\sigma > \sigma'$, then $\Gamma, x : \sigma \vdash e : \tau$.*

Proof. Begin with induction on the proof for $\Gamma, x : \sigma' \vdash e : \tau$.

VAR From the premises we have $x : \sigma' \in \Gamma$ and $\sigma' > \tau$. We also have $\sigma > \sigma'$ but by definition of $\sigma > \sigma'$, if $\sigma' > \tau$ then $\sigma > \tau$. And we also have by definition $x : \sigma \in \Gamma, x : \sigma$. So putting the pieces together, we can use VAR to get

$$\frac{x : \sigma \in \Gamma, x : \sigma \quad \sigma > \tau}{\Gamma, x : \sigma \vdash x : \tau}$$

ABS We have $\Gamma, x : \sigma' \vdash \lambda y. e : \tau$. We want to be able to apply the swap lemma, Lemma 2, so we can get $\Gamma, y : \sigma'', x : \sigma' \vdash e : \tau$. But for that we need proof that $x \neq y$. It is possible to consider the case where $x = y$ and prove it separately, but to make life easier we will use the Barendregt variable convention and instead assume $\lambda y. e$ has undergone some α -conversion such that $x \neq y$.

Now with our $\Gamma, y : \sigma'', x : \sigma' \vdash e : \tau$, using the induction hypothesis results in $\Gamma, y : \sigma'', x : \sigma \vdash e : \tau$. We can then swap x back again to get $\Gamma, x : \sigma, y : \sigma'' \vdash e : \tau$, and then use ABS for $\Gamma, x : \sigma \vdash \lambda y. e : \tau$ as needed.

LET For let $y = e'$ in e , we have $\Gamma, x : \sigma' \vdash e' : \tau'$ and $\Gamma, x : \sigma', y : \overline{\Gamma, x : \sigma'}(\tau') \vdash e : \tau$ and we aim to show

$$\Gamma, x : \sigma \vdash \text{let } y = e' \text{ in } e : \tau$$

First off, we need to convert the $y : \overline{\Gamma, x : \sigma'}(\tau')$ to a $y : \overline{\Gamma, x : \sigma}(\tau')$ somehow. But it can be shown that if $\sigma > \sigma'$, then $\overline{\Gamma, x : \sigma}(\tau) > \overline{\Gamma, x : \sigma'}(\tau)$ due to Lemma 4. So by the inductive hypothesis, we are able to get $\Gamma, x : \sigma', y : \overline{\Gamma, x : \sigma}(\tau') \vdash e : \tau$.

Because y is bound, we know that $x \neq y$ from Barendregt's variable convention. We take the following steps:

$$\begin{array}{ll}
\Gamma, x : \sigma', y : \overline{\Gamma, x : \sigma(\tau')} \vdash e : \tau & \\
\Gamma, y : \overline{\Gamma, x : \sigma(\tau')}, x : \sigma' \vdash e : \tau & \text{by swapping, Lemma 2} \\
\Gamma, y : \overline{\Gamma, x : \sigma(\tau')}, x : \sigma \vdash e : \tau & \text{by inductive hypothesis} \\
\Gamma, x : \sigma, y : \overline{\Gamma, x : \sigma(\tau')} \vdash e : \tau & \text{by swapping again}
\end{array}$$

And proceed to construct the proof for LET as previously.

APP From the premises we have $\Gamma, x : \sigma' \vdash e : \tau' \rightarrow \tau$ and $\Gamma, x : \sigma' \vdash e' : \tau'$. We wish to show $\Gamma, x : \sigma \vdash e e' : \tau$. We apply the induction hypothesis to get $\Gamma, x : \sigma \vdash e : \tau' \rightarrow \tau$ and $\Gamma, x : \sigma \vdash e' : \tau$. Then use APP to build up a proof of $\Gamma, x : \sigma \vdash e e' : \tau$.

The remaining cases The rest of the possible proofs for $\Gamma, x : \sigma' \vdash e : \tau$ can all be proved by applying the induction hypothesis on their structure, much like the case for APP, and so are omitted for brevity.

□

If a context has all the same type schemes for each free variable in an expression as another context, then we can make the same typing judgements with that other context. In other words, we are free to ignore extra variables not used by the expression.

Lemma 5. *If for all $x \in \text{FV}(e)$ where $x : \sigma \in \Gamma$, $x : \sigma \in \Delta$, and if $\Gamma \vdash e : \tau$ then $\Delta \vdash e : \tau$.*

We can also show that we can apply a substitution to both the context and type given in a judgement.

Lemma 6. *If $\Gamma \vdash e : \tau$, then for any substitution s , $\Gamma[s] \vdash e : \tau[s]$.*

An important lemma that we need to show is the substitution lemma, which relates to substituting a variable for an expression, whose type we can prove. The following proof is taken from Wright and Felleisen [20].

Lemma 7 (Substitution). *If $\Gamma \vdash e : \tau$ and $\Gamma, x : \forall \alpha_1, \dots, \alpha_n. \tau \vdash e' : \tau'$, and $x \notin \text{dom}(\Gamma)$ and $\alpha_1, \dots, \alpha_n \cap \text{FTV}(\Gamma) = \emptyset$, then $\Gamma \vdash e'[e/x] : \tau'$.*

Note that the domain of a context $\text{dom}(\bullet, x_1 : \sigma_1, \dots, x_n : \sigma_n)$ is defined as $\{x_1, \dots, x_n\}$.

Proof. Begin with induction on the proof of $\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau \vdash e' : \tau'$.

VAR We have $\Gamma, x : \forall \alpha_1 \dots \alpha_n \vdash y : \tau'$ and want to show $\Gamma \vdash y[e/x] : \tau'$.

If $y \neq x$, then by definition $y[e/x] = y$, so we just need to show $\Gamma \vdash y : \tau'$. From the premises we also have $y : \tau' \in \Gamma, x : \forall \alpha_1 \dots \alpha_n$, but since $y \neq x$ we have $y : \tau' \in \Gamma$. And from here we can use VAR to get $\Gamma \vdash y : \tau'$.

If $y = x$, then we need to show $\Gamma \vdash x : \tau'$. With the premise

$$\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau \vdash x : \tau'$$

we also know $\forall \alpha_1 \dots \alpha_n. \tau > \tau'$, so there exists a substitution s which replaces exactly $\alpha_1, \dots, \alpha_n$, such that $\tau[s] = \tau'$.

We can use Lemma 6 to get $\Gamma[s] \vdash e : \tau[s]$, or $\Gamma[s] \vdash e : \tau'$. And furthermore, because $\alpha_1, \dots, \alpha_n \cap \text{FTV}(\Gamma) = \emptyset$, applying s to Γ will do nothing since Γ 's free type variables are distinct, i.e. $\Gamma = \Gamma[s]$. Leaving us with $\Gamma \vdash e : \tau'$ as required.

ABS We have $\Gamma, x : \forall \alpha_1 \dots \alpha_n \vdash \lambda y. e' : \tau_1 \rightarrow \tau_2$ and want to show

$$\Gamma \vdash (\lambda y. e')[e/x] : \tau_1 \rightarrow \tau_2$$

y is bound, so by convention $y \neq x$. From the premises we have

$$\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau, y : \tau_1 \vdash e' : \tau_2$$

We start by choosing a substitution s , which maps $\alpha_1, \dots, \alpha_n$ to fresh type variables $\alpha'_1, \dots, \alpha'_n$. Additionally, they are distinct as follows.

$$\text{FTV}(\Gamma) \cap \alpha_1, \dots, \alpha_n \cap \alpha'_1, \dots, \alpha'_n = \emptyset$$

Now we manipulate the premise in the following order.

$$\begin{aligned} \Gamma, y : \tau_1, x : \forall \alpha_1 \dots \alpha_n. \tau \vdash e' : \tau_2 & \quad \text{by swapping, Lemma 2} \\ (\Gamma, y : \tau_1, x : \forall \alpha_1 \dots \alpha_n. \tau)[s] \vdash e' : \tau_2 & \quad \text{by Lemma 6} \\ \Gamma, y : \tau_1[s], x : (\forall \alpha_1 \dots \alpha_n. \tau)[s] \vdash e' : \tau_2 & \quad \text{since } \Gamma[s] = \Gamma \end{aligned}$$

And since the range of s is exactly $\alpha_1, \dots, \alpha_n$, we have $(\forall \alpha_1 \dots \alpha_n. \tau)[s] = \forall \alpha_1 \dots \alpha_n. \tau$ — substitution only substitutes free type variables, not bound type variables, and here any possible free type variables are being shadowed.

$$\Gamma, y : \tau_1[s], x : \forall \alpha_1 \dots \alpha_n. \tau \vdash e' : \tau_2 \tag{4.1}$$

We can apply Lemma 5 to $\Gamma \vdash e : \tau$, giving us

$$\Gamma, y : \tau_1[s] \vdash e : \tau \tag{4.2}$$

And when we combine $x \notin \text{dom}(\Gamma)$ with the fact that $x \neq y$, we get

$$x \notin \text{dom}(\Gamma, y : \tau_1[s]) \tag{4.3}$$

Because of how we chose s , we also have¹

$$\text{FTV}(\Gamma, y : \tau_1[s]) \cap \alpha_1, \dots, \alpha_n = \emptyset \tag{4.4}$$

¹ y is distinct here because of the Barendregt variable convention.

Eventually, we use (4.2), (4.1), (4.3) and (4.4) with the induction hypothesis to arrive at

$$\Gamma, y : \tau_1[s] \vdash e'[e/x] : \tau_2[s]$$

But s is a bijection due to how we chose it: That means an inverse, s^{-1} exists. We “apply it to both sides” with Lemma 6

$$\begin{aligned} (\Gamma, y : \tau_1[s])[s^{-1}] &\vdash e'[e/x] : (\tau_2[s])[s^{-1}] \\ (\Gamma, y : \tau_1[s])[s^{-1}] &\vdash e'[e/x] : \tau_2 \\ \Gamma[s^{-1}], y : \tau_1[s][s^{-1}] &\vdash e'[e/x] : \tau_2 \\ \Gamma[s^{-1}], y : \tau_1 &\vdash e'[e/x] : \tau_2 \end{aligned}$$

And because $\alpha'_1, \dots, \alpha'_n \cap \text{FTV}(\Gamma) = \emptyset$, we can get rid of that last substitution and arrive at $\Gamma, y : \tau_1 \vdash e'[e/x] : \tau_2$. From here, we build our way back up with ABS.

$$\begin{aligned} \Gamma &\vdash \lambda y. (e'[e/x]) : \tau_1 \rightarrow \tau_2 \\ \Gamma &\vdash (\lambda y. e')[e/x] : \tau_1 \rightarrow \tau_2 \quad \text{because } x \neq y \end{aligned}$$

LET We have the proof $\Gamma, x : \sigma \vdash \text{let } y = e_1 \text{ in } e_2 : \tau'$ and want to show

$$\Gamma \vdash (\text{let } y = e_1 \text{ in } e_2)[e/x] : \tau'$$

The first premise can be fed directly into the induction hypothesis

$$\begin{aligned} \Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau \vdash e_1 : \tau_1 \\ \Gamma \vdash e_1[e/x] : \tau_1 \end{aligned} \tag{4.5}$$

Since y is bound and thus $y \neq x$, we use the swap lemma with the second premise to get

$$\Gamma, y : \overline{\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau}(\tau_1), x : \forall \alpha_1 \dots \alpha_n. \tau \vdash e_2 : \tau' \tag{4.6}$$

And since we have $\Gamma \vdash e : \tau$, from Lemma 5 there is

$$\Gamma, y : \overline{\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau}(\tau_1) \vdash e : \tau \tag{4.7}$$

Now we want to call the inductive hypothesis on (4.6) and (4.7), but that means we first need to show

$$\{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\Gamma, y : \overline{\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau}(\tau_1)) = \emptyset$$

We do this as follows:

$$\begin{aligned}
& \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\Gamma, y : \overline{\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau}(\tau_1)) \\
& \subseteq \{\alpha_1, \dots, \alpha_n\} \cap (\text{FTV}(\Gamma) \cup \overline{\text{FTV}(\overline{\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau}(\tau_1))}) \\
& = \{\alpha_1, \dots, \alpha_n\} \cap (\text{FTV}(\overline{\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau}(\tau_1))) \\
& = \{\alpha_1, \dots, \alpha_n\} \cap (\text{FTV}(\tau_1) \setminus (\text{FTV}(\tau_1) \setminus \text{FTV}(\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau))) \\
& = \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\tau_1) \cap \text{FTV}(\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau) \\
& \subseteq \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau) \\
& = \{\alpha_1, \dots, \alpha_n\} \cap (\text{FTV}(\Gamma) \cup \text{FTV}(\forall \alpha_1 \dots \alpha_n. \tau)) \\
& = \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\forall \alpha_1 \dots \alpha_n. \tau) \\
& = \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\tau) \setminus \{\alpha_1, \dots, \alpha_n\} \\
& = \emptyset
\end{aligned}$$

And then we can use it to arrive at

$$\Gamma, y : \overline{\Gamma, x : \forall \alpha_1 \dots \alpha_n. \tau}(\tau_1) \vdash e_2[e/x] : \tau'$$

The next part involves an observation, that for any x, τ and σ in any Γ , $\overline{\Gamma}(\tau) > \overline{\Gamma, x : \sigma}(\tau)$, as closing over a context with a *smaller* domain will result in the type scheme having *more* bound type variables — i.e. it is more general. Therefore we can use the generalisation lemma, Lemma 2, to get

$$\begin{array}{ll}
\Gamma, y : \overline{\Gamma}(\tau_1) \vdash e_2[e/x] : \tau' & \\
\Gamma \vdash \text{let } y = e_1[e/x] \text{ in } e_2[e/x] : \tau' & \text{with LET and (4.5)} \\
\Gamma \vdash (\text{let } y = e_1 \text{ in } e_2)[e/x] : \tau' & \text{by definition of substitution}
\end{array}$$

The remaining cases The rest of the proofs for $\Gamma \vdash e'[e/x] : \tau'$ are proved by applying the induction hypothesis on their structure. □

4.1 Type soundness

Now that we have the prerequisite lemmas out of the way, we can move onto proving that the system is **sound**. Soundness is a property of a logical system, stating that every judgement that can be proved within the system is *valid*. In the type theory realm, a typing judgement is valid if the expression actually evaluates to a member of that type. So a type system is considered to be sound if and only if, for every typing judgement that says an expression is well-typed, the expression evaluates to a member of that type — or as Milner put it, “*well-typed programs can't go wrong*” [11][§3.7].

$$\Gamma \vdash e : \tau \rightarrow \Gamma \vDash e : \tau$$

We saw how this was shown with denotational semantics in Chapter 2, but here we have modelled our dynamic semantics with operational semantics. Instead, we prove the soundness of our type system *syntactically*. This approach was first introduced by Wright and Felleisen in 1994 [20], and has since become the de facto method for proving soundness with operational semantics.

Syntactic type soundness involves proving two properties of a type system, *type progress* and *type preservation*. The former states that for any typing judgement we can prove, the expression is either a terminal value, or can be reduced further. The latter property says that when an expression has a type, reducing it to another expression does not change the type.

We will prove the closed variants of these properties, that is they hold in the empty context. The proofs for these properties within our system are as follows.

Theorem 3 (Progress). *If $\bullet \vdash e : \tau$, then either value e or there exists a e' such that $e \rightsquigarrow e'$.*

Proof. Begin by induction on $\bullet \vdash e : \tau$.

VAR This case is not possible — it is impossible to construct a proof for $x \in \bullet$, and thus impossible to arrive at $\bullet \vdash x : \tau$.

ABS Any lambda abstraction is a value due to the rule

$$\frac{}{\text{value } \lambda x.e}$$

APP We have $\bullet \vdash e_1 e_2 : \tau$, so by the premises of APP, we have $\bullet \vdash e_1 : \tau' \rightarrow \tau$ and $\bullet \vdash e_2 : \tau'$.

Consider the induction hypothesis on the first premise. Either value e_1 , or $e_1 \rightsquigarrow e'_1$. In the latter case where e_1 reduces, we can then construct a proof that $e_1 e_2$ reduces, and we're done.

$$\frac{e_1 \rightsquigarrow e'_1}{e_1 e_2 \rightsquigarrow e'_1 e_2}$$

In the former case where e_1 is a value, we then apply the induction hypothesis to $\bullet \vdash e_2 : \tau'$. Again, if e_2 reduces further then we can also construct a proof that $e_1 e_2$ reduces.

But if value e_2 , we cannot reduce in the same way again. Instead, since we know e_1 is a value, and since we also know that $\bullet \vdash e_1 : \tau' \rightarrow \tau$, we can deduce that e_1 must be an abstraction: it has the form $\lambda x.e$. Furthermore, because e_2 is a value, we can then apply the beta reduction rule.

$$\frac{\text{value } e_2}{(\lambda x.e) e_2 \rightsquigarrow e[e_2/x]}$$

We are calling a function in this case, and now have proof that it reduces.

LET Our dynamic semantics for LET are lazy: $\text{let } x = e' \text{ in } e$ can reduce without evaluating e' first. And so for every let statement, we can immediately reduce.

$$\overline{\text{let } x = e' \text{ in } e \rightsquigarrow e[e'/x]}$$

UNIT All unit expressions are values.

$$\overline{\text{value } \square}$$

PRODUCT For $\bullet \vdash e_1 \times e_2 : \tau_1 \times \tau_2$, we have the premises $\bullet \vdash e_1 : \tau_1$ and $\bullet \vdash e_2 : \tau_2$. Using the induction hypothesis on either, if either of them reduce further, then we can also reduce $e_1 \times e_2$.

$$\frac{e_1 \rightsquigarrow e'_1}{e_1 \times e_2 \rightsquigarrow e'_1 \times e_2} \qquad \frac{e_2 \rightsquigarrow e'_2}{e_1 \times e_2 \rightsquigarrow e_1 \times e'_2}$$

If both of them are values, then the product itself is a value.

$$\frac{\text{value } e_1 \quad \text{value } e_2}{\text{value } e_1 \times e_2}$$

PROJ1 For $\bullet \vdash \pi_1 e : \tau$, we have the premise $\bullet \vdash e : \tau \times \tau'$. Applying the induction hypothesis to the premise, if it reduces further, then we can reduce the expression as well.

$$\frac{e \rightsquigarrow e'}{\pi_1 e \rightsquigarrow \pi_1 e'}$$

However, if e is a value then we know it must be of the form $e_1 \times e_2$, as that is the only expression that can have a product type and be a value. Thus we can use the reduction rule

$$\overline{\pi_1 (e_1 \times e_2) \rightsquigarrow e_1}$$

PROJ2 The proof for PROJ2 is pretty much identical to that of PROJ1, so we will omit it for brevity.

LIFT All lifted expressions are also values by definition.

$$\overline{\text{value } \llbracket e \rrbracket}$$

USE And likewise, all use expressions are values by definition.

$$\overline{\text{value } \llbracket e \rrbracket_r}$$

BIND For $\bullet \vdash e_1 \gg e_2 : \tau$, we have two premises

$$\bullet \vdash e_1 : \text{IO}_\rho \tau' \qquad \bullet \vdash e_2 : \tau' \rightarrow \text{IO}_\rho \tau$$

Applying the induction hypothesis to the first premise, if e_1 reduces further then we can also say $e_1 \gg e_2$ reduces further

$$\frac{e_1 \rightsquigarrow e'_1}{e_1 \gg e_2 \rightsquigarrow e'_1 \gg e_2}$$

If e_1 turns out to be a value instead, then because it has the type $\text{IO}_\rho \tau'$, it must be either of the form $\llbracket e \rrbracket$ or $\llbracket e \rrbracket_r$. These are the only expressions that can both be values and have the type $\text{IO}_\rho \tau'$. In either case, there exist reduction rules for both of these.

$$\overline{\llbracket e \rrbracket \gg e_2 \rightsquigarrow e_2 e} \qquad \overline{\llbracket e \rrbracket_r \gg e_2 \rightsquigarrow e_2 e}$$

CONC We can immediately reduce any expression of the form $v \Upsilon w$.

$$\overline{v \Upsilon w \rightsquigarrow v \gg \lambda v.(w \gg \lambda w.\llbracket v \times w \rrbracket)}$$

SUB If we have a subsumption judgement resulting in $\bullet \vdash e : \text{IO}'_\rho \tau$, then we also have the premise $\bullet \vdash e : \text{IO}_\rho \tau$ for some other ρ . But we can just call the induction hypothesis on this premise to show that either e reduces or e is a value. In any case, it is the same e — so we are done.

□

Theorem 4 (Preservation). *If $\bullet \vdash e : \tau$ and $e \rightsquigarrow e'$, then $\bullet \vdash e' : \tau$.*

Proof. As usual, start with induction on the proof for $\bullet \vdash e : \tau$.

VAR Like with the proof for progress, it is not possible to have a typing judgement for a variable in the empty context.

ABS Whilst it is possible to have a typing judgement of $\bullet \vdash \lambda x.e$, it is not possible to have a reduction of the form $(\lambda x.e) \rightsquigarrow e'$. A lambda abstraction cannot reduce any further, and so we do not need to consider this case.

APP We have $\bullet \vdash e_1 e_2 : \tau$. Consider the possible cases for $e_1 e_2 \rightsquigarrow e'$: either e_1 or e_2 must have reduced. When e_1 was reduced

$$\frac{e_1 \rightsquigarrow e'_1}{e_1 e_2 \rightsquigarrow e'_1 e_2}$$

We can use the induction hypothesis with this alongside the first premise $\bullet \vdash e_1 : \tau' \rightarrow \tau$ to get

$$\bullet \vdash e'_1 : \tau' \rightarrow \tau$$

And then using APP and the second premise, we build up a proof for

$$\bullet \vdash e'_1 e_2 : \tau$$

If e_2 reduced to e'_2 instead, the steps are the same: Use the induction hypothesis to get $\bullet \vdash e_2 : \tau'$, and then $\bullet \vdash e_1 e'_2 : \tau$. However, there is one other scenario in which an application can reduce, and that is through beta reduction:

$$\frac{\text{value } e_2}{(\lambda x.e) e_2 \rightsquigarrow e[e_2/x]}$$

We need to show that $\bullet \vdash e[e_2/x] : \tau$. And in this case e_1 must be a lambda abstraction, so the typing rule for the first premise must be an ABS.

$$\text{APP} \frac{\text{ABS} \frac{\bullet, x : \tau' \vdash e : \tau}{\bullet \vdash \lambda x.e : \tau' \rightarrow \tau} \quad \bullet \vdash e_2 : \tau'}{\bullet \vdash (\lambda x.e) e_2 : \tau}$$

Furthermore, we know that $\text{dom}(\bullet)$ and the bound type variables in τ' are distinct — the domain of \bullet is empty, and there are no bound type variables in the type scheme τ' . This allows us to use the substitution lemma, Lemma 7, with $\bullet \vdash e_2 : \tau'$ and $\bullet, x : \tau' \vdash e : \tau$, giving us

$$\bullet \vdash e[e_2/x] : \tau$$

LET There is only way a let expression can reduce, and that is by

$$\text{let } x = e' \text{ in } e \rightsquigarrow e[e'/x]$$

The premises are

$$\begin{array}{l} \bullet \vdash e' : \tau' \\ \bullet, x : \tau' \vdash e : \tau \end{array}$$

And we know that there are no bound type variables in the type scheme τ' , and so they are naturally distinct from $\text{dom}(\bullet)$. Therefore we can just use the substitution lemma again to show that

$$\bullet \vdash e[e'/x] : \tau$$

UNIT In this case, the judgement is $\bullet : \square \vdash \square$. But there is no way for a unit \square to reduce, i.e. $\square \not\rightsquigarrow e$, so we do not need to consider this case.

PRODUCT An expression of the form $\bullet \vdash e_1 \times e_2 : \tau_1 \times \tau_2$ can be reduced in either two ways. Either $e_1 \rightsquigarrow e'_1$ and $e_1 \times e_2 \rightsquigarrow e'_1 \times e_2$, in which case

$$\begin{array}{ll} \bullet \vdash e_1 : \tau_1 & \text{by the premises} \\ \bullet \vdash e'_1 : \tau_1 & \text{by induction hypothesis} \\ \bullet \vdash e'_1 \times e_2 : \tau_1 \times \tau_2 & \text{by PRODUCT} \end{array}$$

Or if $e_2 \rightsquigarrow e'_2$, then we repeat the same steps but acting on the second premise.

PROJ1 $\cdot \vdash \pi_1 e : \tau$ can reduce in two ways as well. If $e \rightsquigarrow e'$ and $\pi_1 e \rightsquigarrow \pi_1 e'$ then

$$\begin{array}{ll} \cdot \vdash e : \tau \times \tau' & \text{by the premises} \\ \cdot \vdash e' : \tau \times \tau' & \text{by induction hypothesis} \\ \cdot \vdash \pi_1 e' : \tau & \text{by PROJ1} \end{array}$$

However it can also be reduced if e happens to be a product, namely $\pi_1(e_1 \times e_2) \rightsquigarrow e_1$. So we just need to show that $\cdot \vdash e_1 : \tau$, which can be done by climbing up the proof tree.

$$\text{PRODUCT} \frac{\cdot \vdash e_1 : \tau \quad \cdot \vdash e_2 : \tau}{\cdot \vdash e_1 \times e_2 : \tau \times \tau'} \\ \text{PROJ1} \frac{}{\cdot \vdash \pi_1 (e_1 \times e_2) : \tau}$$

PROJ2 Again, the proof for PROJ2 is pretty much identical to that of PROJ1 and so is omitted for brevity.

LIFT For $\cdot \vdash \llbracket e \rrbracket : \text{IO}_\rho \tau$, just like \square there is no way for $\llbracket e \rrbracket$ to reduce any further — they are both values. There is nothing needed to be done for this case.

USE The same also applies for $\llbracket e \rrbracket_r$.

BIND The proof we are handling looks like this

$$\text{BIND} \frac{\cdot \vdash e_1 : \text{IO}_\rho \tau' \quad \cdot \vdash e_2 : \tau' \rightarrow \text{IO}_\rho \tau}{\cdot \vdash e_1 \gg e_2 : \text{IO}_\rho \tau}$$

The first thing to note is that there are multiple possible ways in which reduction could have occurred: We will look at each of them one by one.

$e_1 \gg e_2 \rightsquigarrow e'_1 \gg e_2$ The premise for this reduction is $e_1 \rightsquigarrow e'_1$, so we take the following steps

$$\begin{array}{ll} \cdot \vdash e_1 : \text{IO}_\rho \tau' & \text{by the premises} \\ \cdot \vdash e'_1 : \text{IO}_\rho \tau' & \text{by induction hypothesis} \\ \cdot \vdash e'_1 \gg e_2 : \text{IO}_\rho \tau & \text{by BIND} \end{array}$$

$\llbracket e \rrbracket \gg e_2 \rightsquigarrow e_2 e$ We know $\cdot \vdash \llbracket e \rrbracket : \text{IO}_\rho \tau'$, but that does not necessarily mean that the proof for it came from the LIFT rule — it could have also been derived from subsumption, SUB. To help when this scenario arises, we use a small helper lemma listed in the appendix, Lemma [10](#), that says if $\Gamma \vdash \llbracket e \rrbracket : \text{IO}_\rho \tau'$, then $\Gamma \vdash e : \tau'$. Now we can use ABS to construct $\cdot \vdash e_2 e : \text{IO}_\rho \tau$ as needed.

$\llbracket e \rrbracket_r \gg e_2 \rightsquigarrow e_2 e$ This is identical to the case above, except we use Lemma [11](#) instead.

CONC We have the proof and premises

$$\text{CONC} \frac{\bullet \vdash v : \text{IO}_{\rho_1} \tau_1 \quad \bullet \vdash w : \text{IO}_{\rho_2} \tau_2 \quad \text{ok } \rho_1 \cup \rho_2}{\bullet \vdash v \Upsilon w : \text{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2}$$

There is also only one possible reduction that could have occurred

$$v \Upsilon w \rightsquigarrow v \gg\equiv \lambda v. (w \gg\equiv \lambda w. \llbracket v \times w \rrbracket)$$

So we need to show $\bullet \vdash v \gg\equiv \lambda v. (w \gg\equiv \lambda w. \llbracket v \times w \rrbracket) : \text{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2$ — not exactly the prettiest judgement — which we can construct eventually from a BIND.

To begin, we will first show that both v and w can be subsumed into the larger combined heap, $\rho_1 \cup \rho_2$.

$$\begin{array}{ll} \bullet \vdash v : \text{IO}_{\rho_1} \tau_1 & \text{from premises} \\ \rho_1 \geq: \rho_1 \cup \rho_2 & \text{by UNIONL} \\ \text{ok } \rho_1 \cup \rho_2 & \text{from premises} \\ \bullet \vdash v : \text{IO}_{\rho_1 \cup \rho_2} \tau_1 & \text{by SUB} \end{array}$$

$$\begin{array}{ll} \bullet \vdash w : \text{IO}_{\rho_2} \tau_2 & \text{from premises} \\ \bullet, v : \tau_1 \vdash w : \text{IO}_{\rho_2} \tau_2 & \text{by Lemma 3} \\ \rho_2 \geq: \rho_1 \cup \rho_2 & \text{by UNIONR} \\ \text{ok } \rho_1 \cup \rho_2 & \text{from premises} \\ \bullet, v : \tau_1 \vdash w : \text{IO}_{\rho_1 \cup \rho_2} \tau_2 & \text{by SUB} \end{array}$$

Then we show that we can make a lambda, taking a *pure* τ_2 , and making a $\tau_1 \times \tau_2$ inside that new IO monad.

$$\frac{\frac{\frac{\bullet, v : \tau_1, w : \tau_2 \vdash v : \tau_1 \quad \bullet, v : \tau_1, w : \tau_2 \vdash w : \tau_2}{\bullet, v : \tau_1, w : \tau_2 \vdash v \times w : \tau_1 \times \tau_2}}{\bullet, v : \tau_1, w : \tau_2 \vdash \llbracket v \times w \rrbracket : \text{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2}}{\bullet, v : \tau_1 \vdash \lambda w. \llbracket v \times w \rrbracket : \tau_2 \rightarrow \text{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2}}$$

We then take this lambda and bind our *monadic* w to it

$$\frac{\bullet, v : \text{IO}_{\rho_1} \tau_1 \vdash w : \text{IO}_{\rho_1 \cup \rho_2} \tau_2 \quad \bullet, v : \tau_1 \vdash \lambda w. \llbracket v \times w \rrbracket : \tau_2 \rightarrow \text{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2}{\bullet, v : \tau_1 \vdash w \gg\equiv (\lambda w. \llbracket v \times w \rrbracket) : \text{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2}$$

Now we can make another lambda, this time taking in the *pure* τ_1

$$\frac{\bullet, v : \tau_1 \vdash w \gg\equiv (\lambda w. \llbracket v \times w \rrbracket) : \text{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2}{\bullet \vdash \lambda v. w \gg\equiv (\lambda w. \llbracket v \times w \rrbracket) : \tau_1 \rightarrow \text{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2}$$

And then we finally bind it with our *monadic* v

$$\frac{\bullet \vdash v : \mathbb{IO}_{\rho_1 \cup \rho_2} \tau_1 \quad \bullet \vdash \lambda v. w \gg= \lambda w. \llbracket v \times w \rrbracket : \tau_1 \rightarrow \mathbb{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2}{\bullet \vdash v \gg= (\lambda v. w \gg= (\lambda w. \llbracket v \times w \rrbracket)) : \mathbb{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2}$$

The proof is quite long, but since our typing system is syntax directed, it closely matches the expression it reduces to, making it easier to construct.

SUB Last but not least, we have a subsumption typing judgement

$$\text{SUB} \frac{\bullet \vdash e : \mathbb{IO}_{\rho'} \tau \quad \rho \geq: \rho' \quad \text{ok } \rho}{\bullet \vdash e : \mathbb{IO}_{\rho} \tau}$$

However because we know nothing of the syntax of e , e could reduce to anything and we are left with some e' , such that $e \rightsquigarrow e'$. Nevertheless, applying the induction hypothesis with $\bullet \vdash e : \mathbb{IO}_{\rho'} \tau$ and $e \rightsquigarrow e'$ gives us

$$\bullet \vdash e' : \mathbb{IO}_{\rho'} \tau$$

Which we can stick back into SUB to get

$$\bullet \vdash e' : \mathbb{IO}_{\rho} \tau$$

□

Together these two properties can ensure that no expression can go wrong, which within operational semantics means it cannot get *stuck*. A stuck expression is an expression which cannot be reduced any further (in its normal form) and is not a value. The progress theorem has shown that evaluation of any well-typed expression cannot end up in this state, so it does indeed evaluate to something in the end. And the preservation theorem guarantees us that the type of the final value will remain the same. Therefore we have shown the operational equivalent of the denotational

$$\Gamma \vdash e : \tau \rightarrow \Gamma \Vdash e : \tau$$

Chapter 5

Mechanisation in Agda

In this chapter, we will look at how the proofs in Chapter 4 were mechanically formalised — in other words, proved within a proof assistant. The proofs were mechanised within Agda [12], a dependently typed programming language with an ML syntax, similar to that of Haskell’s. It can be used for general purpose programming, but since it is rooted in Martin-Löf intuitionistic type theory [10], it can also be used as a prove formal properties.

Properties are proven in Agda by writing programs that satisfy types. This approach takes advantage of the Curry-Howard correspondence, which is that propositions are analogous to types, and proofs are analogous to programs that fulfil that type. We first create types that represent our propositions. Then we can then prove our propositions by constructing a value for the program that satisfies the type — hence why intuitionistic logic is also known as constructive logic.

As a brief example, a proposition of the form $\forall a.a \rightarrow b$, translates into a program of type $\forall a \rightarrow b$. If you read the proposition as “*if a, then b*”, then you can read the type as “*given any proof of a, then I can construct a proof of b*”. And that is indeed what the program should do: It is a function that takes an argument of a , and returns something of type b .

The vast majority of the work in mechanising was in the proofs related to polymorphism in the Hindley-Damas-Milner system, not the resourceful parts. The syntax-directed Hindley-Damas-Milner system has been successfully formalised in Coq [6], and a formalisation for System F was recently created in Agda [2]. However as it stands, we are not aware of any formalizations of Hindley-Damas-Milner within Agda. Regardless, the techniques and approaches here are similar, and we base much of the work on the proofs in [20].

5.1 Definitions

The framework within Agda for working with the type system is built upon the formalisation of the simply typed lambda calculus in Agda by Kokke et al. [8]

```

data Term : Set where
  \_      : Id → Term
  λ_⇒_    : Id → Term → Term
  _'      : Term → Term → Term
  !t_⇒_in'_ : Id → Term → Term → Term
  [ ]     : Term → Term
  _>=>_   : Term → Term → Term
  □       : Term
  use     : Resource → Term → Term
  _x_     : Term → Term → Term
  π1    : Term → Term
  π2    : Term → Term
  _Υ_     : Term → Term → Term

data Heap : Set where
  World : Heap
  \_     : Resource → Heap
  _U_    : Heap → Heap → Heap

data Type : Set where
  \_ : Id → Type
  _⇒_ : Type → Type → Type
  IO  : Heap → Type → Type
  □   : Type
  _x_ : Type → Type → Type

data TypeScheme : Set where
  V_·_ : Id → TypeScheme → TypeScheme
  \_   : Type → TypeScheme

```

Listing 1: Grammar definitions. Note that some of the notation (e.g. \forall) had to be substituted due to overlaps with keywords in Agda.

It begins with the grammar, shown in Listing 1.

For relations, such as the heap well-formed relation $\text{ok } \rho$, we use a pattern of defining a new data type. The type then acts as a proposition that the relation holds, and creating an instance of the data type constitutes as providing proof. Furthermore, the constructors for the data type directly correspond to the rules for the relation.

```

data Ok : Heap → Set where
  OkZ : ∀ {r}
    -----
    → Ok (` r)
  OkS : ∀ {a b}
    → Ok a
    → Ok b
    → a n b =∅
    -----
    → Ok (a u b)
  OkWorld : -----
    Ok World

```

We also define a data type to represent proof that a heap is a subheap of another heap. Since subheaping is a binary relation over two heaps, the data type is also parameterised over two heaps, which appear in the type.

```

data _≥:_ : Heap → Heap → Set where
  ≥:World : ∀ {ρ} → World ≥: ρ
  ≥:refl : ∀ {ρ} → ρ ≥: ρ
  ≥:Ul : ∀ {ρ ρ' ρ''}
    → ρ' ≥: ρ
    -----
    → ρ' ∪ ρ'' ≥: ρ
  ≥:Ur : ∀ {ρ ρ' ρ''}
    → ρ' ≥: ρ
    -----
    → ρ'' ∪ ρ' ≥: ρ

```

If we wanted to show that $\text{Net} \cup \text{File} \geq: \text{Net}$, we would follow the same steps that we would carry out in a proof tree, in order to construct a value that inhabits the type $\text{Net} \cup \text{File} \geq: \text{Net}$.

$\begin{array}{l} _ : \text{Net} \cup \text{File} \geq: \text{Net} \\ _ = \geq:U^l \geq:refl \end{array}$	$\frac{\text{REFL} \overline{\text{Net} \geq: \text{Net}}}{\text{UNIONL} \overline{\text{Net} \cup \text{File} \geq: \text{Net}}}$
---	--

The most important relation however, is the typing relation. With Agda's Unicode support we are able to define the rules, shown in Listing 2, with a notation similar to what we used in Chapter 3. Typing judgements can be constructed as follows:

$\begin{array}{l} _ : \emptyset \vdash \lambda \text{"x"} \Rightarrow \text{"x"} \text{ } (\text{"\alpha"} \Rightarrow \text{"\alpha"}) \\ _ = \vdash \lambda (\vdash \text{Z (Inst SZ refl refl)}) \end{array}$	$\begin{array}{l} \text{VAR} \frac{x : \alpha \in \bullet, x : \alpha \quad \alpha > \alpha}{\bullet, x : \alpha \vdash x : \alpha} \\ \text{ABS} \frac{}{\bullet \vdash \lambda x. x : \alpha \rightarrow \alpha} \end{array}$
--	---

Where Inst SZ refl refl is proof that the type scheme α can be instantiated to the type α (by not instantiating anything type variables at all).

```

data _>_ : TypeScheme → Type → Set where
  Inst : ∀ {σ τ}
    → (s : Substitution)
    → subDomain s ≡ TSvars σ
    → subT s (TStype σ) ≡ τ
    -----
    → σ > τ

```

The small-step relation is defined in a similar way to the typing relation.

```

data _≈_ : Term → Term → Set where
  ξ→1 : ∀ {e1 e2 e1'}
    → e1 ≈ e1'
    -----
    → e1 · e2 ≈ e1' · e2

  ξ→2 : ∀ {e1 e2 e2'}
    → e2 ≈ e2'
    -----
    → e1 · e2 ≈ e1 · e2'

  β→λ : ∀ {x e e'}
    → Value e'
    -----
    → (λ x ⇒ e) · e' ≈ e [ x := e' ]
-- and so on ...

```

5.2 Type schemes and type variables

One of the main design decisions made early on was how to represent type schemes. Quantified type variables in type schemes are often represented as a sequence $\forall \alpha_1, \dots, \alpha_n \cdot \tau$. This expands out to $\forall \alpha_1 \cdot \forall \dots \cdot \forall \alpha_n \cdot \tau$ in the end, and is how type schemes are ultimately defined in the grammar. But we need to be able to reason about type variables in the sequence format for some of the proofs. A helper function `WV` was created for this reason. It allows for type schemes to be created and manipulated in terms of lists.

```

WV : List Id → Type → TypeScheme
WV (α :: αs) τ = V α · (WV αs τ)
WV [] τ = ` τ

```

Now we can rewrite propositions, such as the substitution lemma (Lemma [7](#)), in a way that lets us bind the list of quantified type variables and use it elsewhere.

```

subst : ∀ {Γ x e e' as τ τ'}
  → Γ ⊢ e § τ
  → Γ , x § WV as τ ⊢ e' § τ'
  → Disjoint as (FTVC Γ)
  -----
  → Γ ⊢ e' [ x := e ] § τ'

```

There are also functions to extract the quantified type variables and type from a type scheme, and equivalence proofs that can be used to convince the type checker that they are equivalent to their `WV` form.


```

data _⊢_&_ : Context → Term → Type →
↳ Set where

⊢` : V {Γ x σ τ}
  → x & σ ∈ Γ
  → σ > τ
  -----
  → Γ ⊢ ` x & τ

⊢λ : V {Γ x τ' τ e}
  → Γ , x & ` τ' ⊢ e & τ
  -----
  → Γ ⊢ λ x ⇒ e & (τ' ⇒ τ)

⊢· : V {Γ e e' τ τ'}
  → Γ ⊢ e & τ' ⇒ τ
  → Γ ⊢ e' & τ'
  -----
  → Γ ⊢ e · e' & τ

⊢lt : V {Γ e e' τ τ' x}
  → Γ ⊢ e' & τ'
  → Γ , x & close Γ τ' ⊢ e & τ
  -----
  → Γ ⊢ lt x = e' in' e & τ

⊢× : V {Γ e e' τ τ'}
  → Γ ⊢ e & τ
  → Γ ⊢ e' & τ'
  -----
  → Γ ⊢ e × e' & τ × τ'

⊢π1 : V {Γ e τ τ'}
  → Γ ⊢ e & τ × τ'
  -----
  → Γ ⊢ π1 e & τ

⊢π2 : V {Γ e τ τ'}
  → Γ ⊢ e & τ × τ'
  -----
  → Γ ⊢ π2 e & τ'

⊢□ : V {Γ}
  -----
  → Γ ⊢ □ & □

-- Monadic rules

⊢[] : V {Γ e τ ρ}
  → Γ ⊢ e & τ
  → Ok ρ
  -----
  → Γ ⊢ [] e & IO ρ τ

⊢use : V {Γ e τ r}
  → Γ ⊢ e & τ
  -----
  → Γ ⊢ use r e & IO (` r) τ

⊢>>= : V {Γ e e' τ τ' ρ}
  → Γ ⊢ e & (IO ρ τ')
  → Γ ⊢ e' & (τ' ⇒ IO ρ τ)
  -----
  → Γ ⊢ e >>= e' & IO ρ τ

⊢γ : V {Γ e1 e2 τ1 τ2 ρ1 ρ2}
  → Γ ⊢ e1 & IO ρ1 τ1
  → Γ ⊢ e2 & IO ρ2 τ2
  → Ok (ρ1 ∪ ρ2)
  -----
  → Γ ⊢ e1 γ e2 & IO (ρ1 ∪ ρ2) (τ1
↳ x τ2)

⊢sub : V {Γ e τ ρ ρ'}
  → Γ ⊢ e & IO ρ' τ
  → ρ ≥: ρ'
  → Ok ρ
  -----
  → Γ ⊢ e & IO ρ τ

```

Listing 2: The typing rules as they are defined in Agda.

```

TStype : TypeScheme → Type
TStype (V _ · σ) = TStype σ
TStype (` τ) = τ

TSvars : TypeScheme → List Id
TSvars (V α · σ) = α :: TSvars σ
TSvars (` τ) = []

TStype≡ : ∀ {αs τ} → TStype (VW αs τ) ≡ τ
TStype≡ {[]} {τ} = refl
TStype≡ {α :: αs} {τ} = TStype≡ {αs}

TSvars≡ : ∀ {αs τ} → TSvars (VW αs τ) ≡ αs
TSvars≡ {[]} = refl
TSvars≡ {α :: αs} = cong (λ_::_ α) TSvars≡

```

5.3 Properties

The type preservation theorem (Theorem [4](#)) is proven with the following function — given proof that e has type τ in the empty context, and proof that e reduces to e' in one step, then we can provide proof that e' also has the type τ in the empty context.

```

preservation : ∀ {e e' τ}
  → ∅ ⊢ e ∋ τ
  → e ~ e'
  -----
  → ∅ ⊢ e' ∋ τ

```

For progress (Theorem [3](#)), we need to be able to say either the expression is a value or it reduces to something else. We create a new data type to represent the two cases where this can happen.

```

data Progress (e : Term) : Set where
  step : ∀ {e'}
    → e ~ e'
    -----
    → Progress e
  done : Value e
  -----
  → Progress e

progress : ∀ {e τ}
  → ∅ ⊢ e ∋ τ
  -----
  → Progress e

```

5.4 Postulations

There are number of properties that are postulated throughout, and these should eventually be proven at some point. Postulations do not require any proofs and are just assumed by Agda to be true, so they are used sparingly here. They are mainly used in cases where we know something to be true that is usually needed to prove another property, but the proof for them is large and non-trivial.

One such example is when dealing with the Barendregt variable and convention. On paper it is fine to implicitly assume all bound variables are unique, but Agda is not so lenient. We need to explicitly perform the alpha conversion and show that the new variable is distinct from any free variable and that the two expressions are equivalent

```
α-conv : ∀ {x e}
        → ∃[ y ] ((∀ {e'} → y ∉ FV e') → (λ x ⇒ e ≡ λ y ⇒ (e [ x := ` y ])))
```

Working with existential quantification and these non-trivial equivalences can get pretty unwieldy, so in some proofs we simply just make the postulation

```
postulate x∉FVv : x ∉ FV v
```

Chapter 6

Evaluation and further work

6.1 Separation logic

In Section [2.1](#) we briefly touched separation logic, but now that we have defined our system we can look at some parallels between them. Take the frame rule again, which allows additional predicates to be inferred in a specification, under the condition that the code does not use any of the variables in the predicate.

$$\frac{\{p\} \text{code } \{q\}}{\{p * r\} \text{code } \{q * r\}} \text{ where no variable occurring free in } r \text{ is modified by } \text{code}$$

Our well-formed heap judgement has a rule that allows heaps to be merged, under the condition that they do not overlap:

$$\frac{\text{ok } \rho \quad \text{ok } \rho' \quad \rho \cap \rho' = \emptyset}{\text{ok } \rho \cup \rho'}$$

This mirrors the idea in separation logic that we only need to worry about the variables relevant to our code, or in our terminology, the resources relevant to our type. So our subsumption rule, SUB, that lets heaps get promoted to a larger heap, ends up being our version of the frame rule. And the concurrency rule from separation logic

$$\frac{\{p_1\} \text{code}_1 \{q_1\} \quad \{p_2\} \text{code}_2 \{q_2\}}{\{p_1 * p_2\} \text{code}_1 \parallel \text{code}_2 \{q_1 * q_2\}}$$

Ends up becoming our concurrency typing rule, CONC. The preconditions and postconditions are the heaps in the IO type, and the separating conjunction $*$ is our union \cup .

$$\frac{\Gamma \vdash e_1 : \text{IO}_{\rho_1} \tau_1 \quad \Gamma \vdash e_2 : \text{IO}_{\rho_2} \tau_2 \quad \text{ok } \rho_1 \cup \rho_2}{\Gamma \vdash e_1 \text{ } \Upsilon \text{ } e_2 : \text{IO}_{\rho_1 \cup \rho_2} \tau_1 \times \tau_2}$$

6.2 Well formed heaps

We have proved the system defined in Chapter 3 is sound, and moreover we can show that every concurrent expression has a well formed heap — that is a resource is not used more than once in the heap, and thus not accessed concurrently.

Theorem 5 (Concurrent heap is well formed). *If $\Gamma \vdash e_1 \vee e_2 : \text{IO}_\rho \tau$, then $\text{ok } \rho$*

Proof. By induction on the possible proofs for $\Gamma \vdash e_1 \vee e_2 : \text{IO}_\rho \tau$.

CONC We need to show $\text{ok } \rho_1 \cup \rho_2$, and from the premises we have $\text{ok } \rho_1 \cup \rho_2$.

SUB We need to show $\text{ok } \rho'$ and from the premises, we also have $\text{ok } \rho'$.

□

As you can see, this just follows from the definition of the concurrency and subsumption rules. What we would really like to prove is that *all* monadic expressions have well formed heaps.

Conjecture 1. *If $\cdot \vdash e : \text{IO}_\rho \tau$, then $\text{ok } \rho$*

This is a stronger version of Theorem 5. However this is not made easy in the system as it stands, for one simple yet annoying reason: Types are not necessarily well formed, and ABS allows any type to be introduced. This can be easily shown with the identity function.

$$\frac{\frac{x : \text{IO}_{\text{FileUFile}\mathcal{T}} \in \cdot, x : \text{IO}_{\text{FileUFile}\mathcal{T}} \quad \text{IO}_{\text{FileUFile}\mathcal{T}} > \text{IO}_{\text{FileUFile}\mathcal{T}}}{\cdot, x : \text{IO}_{\text{FileUFile}\mathcal{T}} \vdash x : \text{IO}_{\text{FileUFile}\mathcal{T}}} \text{VAR}}{\cdot \vdash \lambda x.x : \text{IO}_{\text{FileUFile}\mathcal{T}} \rightarrow \text{IO}_{\text{FileUFile}\mathcal{T}}} \text{ABS}$$

ABS acts as a mechanism to generate types with malformed heaps, such as an IO monad that contains two Files in its heap. In reality this is not an issue, as there is no way to construct a value to pass into this lambda. Because of this I believe that Conjecture 1 is still provable, but only when the context only contains well formed type schemes, and hence why it only applies for the empty context.

We can extend this to include a relation for well formed contexts and well formed types, and use this in our definition.

$$\begin{array}{c}
\boxed{\text{ok } \Gamma} \\
\frac{}{\text{ok } \cdot} \quad \frac{\text{ok } \Gamma \quad \text{ok } \sigma}{\text{ok } \Gamma, x : \sigma} \\
\boxed{\text{ok } \sigma} \\
\frac{\text{ok } \tau}{\text{ok } \forall \alpha_1, \dots, \alpha_n \cdot \tau} \\
\boxed{\text{ok } \tau} \\
\frac{}{\text{ok } \square} \quad \frac{}{\text{ok } \alpha} \quad \frac{\text{ok } \tau \quad \text{ok } \tau'}{\text{ok } \tau \rightarrow \tau'} \quad \frac{\text{ok } \tau \quad \text{ok } \tau'}{\text{ok } \tau \times \tau'} \quad \frac{\text{ok } \tau \quad \text{ok } \rho}{\text{ok } \text{IO}_\rho \tau}
\end{array}$$

Conjecture 2. *If $\text{ok } \Gamma$ and $\Gamma \vdash e : \text{IO}_\rho \tau$, then $\text{ok } \rho$*

So far two lemmas have been proven to aid in the proof of the above conjecture.

Lemma 8. *If $\text{ok } \tau$ then for any substitution s , $\text{ok } \tau[s]$*

Lemma 9. *If $\text{ok } \sigma$ and $\sigma > \tau$, then $\text{ok } \tau$*

This idea of well formed types adds quite a bit of extra overhead to the system, so it is left as further work.

6.3 Let polymorphism

All this theory was motivated by a very practical cause: to provide some simple concurrency guarantees at the type level in a programming language. Thus the intention is that some of this may eventually make its way into a type system of an existing or novel language.

So instead of starting with the simply typed lambda calculus, the system was based off of the Hindley-Damas-Milner system with its polymorphic type discipline. It cannot be understated the amount of extra complexity and work this introduced into the proofs for soundness, but polymorphism is an essential feature in today's modern functional programming languages. The let polymorphism does not interact much with the monadic and resourceful parts of the type system, but it is desirable to show that the two are compatible with each other regardless.

$$\begin{array}{c}
\frac{x : \tau \in \Gamma}{\tau > \tau} \text{VAR} \quad \frac{z : \tau \rightarrow \tau \in \Gamma}{\tau \rightarrow \tau > \tau \rightarrow \tau} \text{VAR} \\
\frac{\Gamma, x : \tau \vdash x : \tau}{\Gamma \vdash \lambda x.x : \tau \rightarrow \tau} \text{ABS} \quad \frac{\Gamma, z : \tau \rightarrow \tau \vdash z : \tau \rightarrow \tau}{\Gamma, z : \tau \vdash z z : \tau \rightarrow \tau} \text{APP} \\
\hline
\Gamma \vdash \text{let } z = \lambda x.x \text{ in } z z : \tau \rightarrow \tau \quad \text{LET}
\end{array}$$




Figure 6.1: How syntax directed typing rules mean that the proof tree is isomorphic to the syntax tree.

6.4 Almost syntax directed

As mentioned earlier, the system was based off a syntax directed treatment of Hindley-Damas-Milner. That meant that there were only four rules instead of the usual six, where instantiation and generalisation are rolled into VAR and LET respectively. The main benefit of this was that it meant there was exactly one typing judgement for each form of syntax, so constructing proofs for a program was a breeze — you just follow the corresponding rules for each term. For example, in the expression $\text{let } z = \lambda x.x \text{ in } z z$, a let consisting of an abstraction and an application, the proof begins with a LET, then the premises ABS and APP as shown in figure [6.1](#).

In our system however, the subsumption rule SUB prevents this. For every typing judgement of the form $\Gamma \vdash e : \text{IO}_\rho \tau$, there are two possible proofs that can be constructed for it: the proof corresponding to the syntax of e , and SUB. Furthermore, because the subheap relation is reflexive, its possible to have an infinite proof tree with repeated applications of SUB as shown in figure [6.2](#). When working with proofs this did not turn out to be too big of an issue, but it is undoubtedly a little bit unsatisfying to lose such a nice property of the system.

6.5 Modelling state within the monad

The modelling of the IO monad is extremely simplified, and does nothing more than sequence computation. To illustrate this point, the reduction rule for

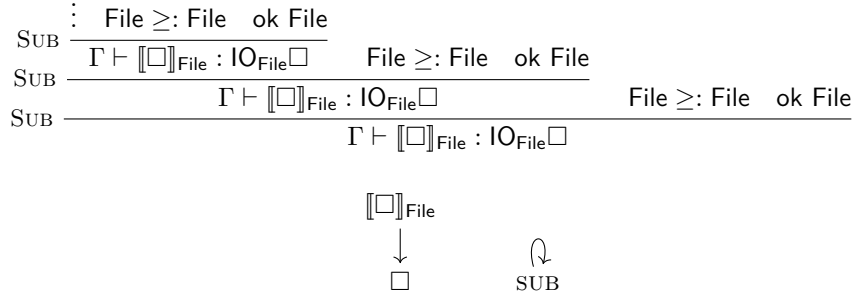


Figure 6.2: How an infinite chain of SUB can be produced, leading nowhere.

concurrency is completely sequential.

$$\overline{v \Upsilon w \rightsquigarrow v \gg= \lambda v. (w \gg= \lambda w. \llbracket v \times w \rrbracket)}$$

In the real world, this would be pointless. We want to be able to compute these expressions concurrently, so rewriting everything in terms of $\gg=$ won't cut it. By modelling the state of the world, like what actually happens in the **IO** monad of Haskell, we can perform more advanced reasoning about things such as concurrency. One such possible method would be to create a new rule for reduction of monadic computations, that passes about state. Then the concurrency rule could be written something like this:

$$\boxed{\langle e, s \rangle \rightsquigarrow \langle e', s' \rangle} \quad \frac{\langle e_1, s \rangle \rightsquigarrow \langle e'_1, s'_1 \rangle \quad \langle e_2, s \rangle \rightsquigarrow \langle e'_2, s'_2 \rangle}{\langle e_1 \Upsilon e_2, s \rangle \rightsquigarrow \langle e'_1 \times e'_2, s'_1 \cup s'_2 \rangle}$$

We could then reason further about how the two states $s'_1 \cup s'_2$ are merged — something that separation logic could be used for.

6.6 Extending the inference algorithm

One of the main contributions of the Hindley-Damas-Milner system was the inclusion of a type inference algorithm, Algorithm W [11]. A type system on its own only lets us ascertain that a program is well-typed, given that we supply it an expression and a type. But the design of Hindley-Damas-milner made it possible for the type to be inferred algorithmically, from the program alone. Algorithm W also always infers the *principal type*: the most general, polymorphic type an expression could have. Algorithm W freed ML programmers from having to explicitly write type annotations, and is one of its most iconic features. Naturally, we should explore extending Algorithm W to handle our new resourceful types.

6.7 Heap polymorphism

It would be convenient if we were able to have expressions like the following well-typed:

$$\begin{aligned} \Gamma = \bullet, f : \square \rightarrow \mathbf{IO}_{\text{File}}\square, g : \square \rightarrow \mathbf{IO}_{\text{Net}}\square \\ \Gamma \vdash \text{let } x = \llbracket \square \rrbracket \text{ in } (x \ggg f) \curlywedge (x \ggg g) : \mathbf{IO}_{\text{File} \cup \text{Net}}\square \times \square \end{aligned}$$

This is not possible in the current type system: x needs to be both $\mathbf{IO}_{\text{File}}$ and \mathbf{IO}_{Net} simultaneously! This is very much the same question that Hindley, Damas and Milner set out to solve, except instead of polymorphic *types* we want to have polymorphic *heaps*. The basic idea for this would involve extending the definition of a heap to allow for variable lookup, \mathbf{h} :

$$\text{heap } \rho ::= r \mid \sigma \cup \sigma' \mid \text{World} \mid \mathbf{h}$$

And then redefining substitutions to be pairs of a map from type variables to types, and a map from heap variables to heaps.

$$S : (\alpha \mapsto \tau \times h \mapsto \rho)$$

Then in the definition of substitution, whenever a mapping from a heap variable to heap is encountered we would apply the substitution.

$$\begin{aligned} (\mathbf{IO}_{\rho'}\tau)[h/\rho] &= \mathbf{IO}_{\rho'[h/\rho]}\tau \\ r[h/\rho] &= r \\ (\rho_1 \cup \rho_2)[h/\rho] &= \rho_1[h/\rho] \cup \rho_2[h/\rho] \\ \text{World}[h/\rho] &= \text{World} \\ h'[h/\rho] &= \begin{cases} \rho & \text{if } h' = h \\ h' & \text{otherwise} \end{cases} \end{aligned}$$

6.8 Dependently typed resources

So far our resources have just been simple static placeholders that represented system resources, for example the filesystem and the network. This is just for the sake of argument: resources can be as general or as specific as is needed for the programmer's use case.

Suppose a programmer wanted to perform multiple concurrent operations on various files. A resource representing the entire filesystem would be too coarse-grained for this purpose. Instead they might want to have a separate resource for each file.

$$\text{resource } r ::= \text{foo.txt} \mid \text{bar.txt} \mid \text{baz.txt} \dots$$

But this is very static. If the programmer writes some code that operates on a new file, then the new file would need to be appended to the list of these

predefined resources. And what if they don't know what files they will be touching at typecheck time? Instead, it would be much more natural if the programmer could write programs like

```
readFile "foo.txt" : IOfoo.txtString
writeFile "bar.txt" : String → IObar.txt□
```

and have the file automatically *lifted* from the expression and into the heap of the type. This is a job for **dependent types**: having the type depend on the term. However in our system, our resources and types are not unified. In a dependently typed lambda calculus, the Π rule substitutes values into types like so¹

$$\frac{\Gamma \vdash \Pi(x : \tau').\tau : \star \quad \Gamma, x : \tau' \vdash e : \tau}{\Gamma \vdash \lambda x.e : \Pi(x : \tau').\tau}$$

$$\frac{\Gamma \vdash e : \Pi(x : \tau').\tau \quad e' : \tau'}{\Gamma \vdash e e' : \tau[e'/\tau']}$$

We would need to combine the syntaxes for resources and types together for this rule, or flesh out types to include type constructors through which resources and heaps could be embedded in. Alternatively, we could be lazy and avoid this by making it dependent only in the heap. We could introduce a type of lambda for IO monads that tags it with a placeholder resource, which gets substituted when an argument is applied.

$$\frac{\Gamma, x : \tau' \vdash e : \tau}{\Gamma \vdash (\lambda_{\text{IO}} x.e) : \tau' \rightarrow_{\text{IO}} \tau}$$

$$\frac{\Gamma \vdash e : \tau' \rightarrow_{\text{IO}} \tau \quad \text{Value } e' \quad \Gamma \vdash e' : \tau'}{\Gamma \vdash e e' : \text{IO}_{e'}\tau}$$

It would be wise to restrict this to values, so we aren't going about tagging resources with half-evaluated computations. We would also need a notion of equality for values so that we can safely tell what values are disjoint within a heap. This will pose some interesting challenges in practice, like knowing statically what variable expressions might represent the same resource. Does the typechecker here know that $x = y$?

```
x ← readLine
readFile x
let y = x in writeFile y "hello"
```

¹ $\Gamma \vdash \tau : \star$ is a judgement saying that τ is a type in some universe.

6.9 Casting heaps

We have only been passing about the primitive unit type, but realistically we are going to want to have functions that actually have functionality, like the aforementioned

$$\text{readFile} : \text{String} \rightarrow \text{IO}_{\rho}\text{String}$$

A function like this is not going to be built-in, instead it will probably be part of a standard library. So how will it be implemented? The easiest solution would be to wrap around an existing function that reads files in a plain IO monad. Remember that plain IO is equivalent to IO_{World} :

$$\text{readFile}' : \text{String} \rightarrow \text{IO}_{\text{World}}\text{String}$$

How do we bring the IO monad into a File heap? If it were a pure computation, then we could have just used the use operator:

$$\lambda s \rightarrow \llbracket \text{readFile}' s \rrbracket_{\text{File}}$$

But it's a monadic computation, so we would end up with a doubly nested monad like $\text{IO}_{\text{File}}(\text{IO}_{\text{World}}\text{String})$. Can we subsume it? No! Because *World* is at the top of the sub-heap chain. It already uses all possible resources, and we can't claim that it uses less than what it actually does! The type is overly cautious in this case.

In an ideal world, standard library functions would be built up from a handful of built-ins for extremely low level primitives, that keep track of things like file descriptors. Then when they are put together, we could subsume them to keep track of higher-level resources. For example, after opening several file descriptors, we could say we are using them to access a database. But practically, we are going to have to port functions that use the full IO_{World} monad to, a monad that uses only the resources that we *know* the function is accessing.

Haskell's `base` library contains the `unsafePerformIO` function, which has the blasphemous type signature `IO a -> a`. But it provides an escape mechanism for library authors to tell the type checker that something really doesn't have any side effects, and is safe to use as such. If we are to allow library authors to annotate their heaps correctly, one such solution would be to provide an unsafe function like this, except instead of discarding the IO part, it would discard heap annotations.

$$\text{unsafeCastHeap} : \text{IO}_{\rho}\alpha \rightarrow \text{IO}_{\rho'}\alpha$$

There is an implicit element of trust here that this will be used selectively by library authors, only when they know that a function accesses a specific set of resources, and only those resources.

Another more radical approach to this would be to flip *World* on its head: Introduce a notion of an “empty” heap which contains no resources, and as such is distinct from all other heaps. Existing code in a plain IO monad would have this empty heap by default, and when a library author wants to tag resources a function might use, they would subsume it upwards into the heap they want.

Chapter 7

Conclusion

We have presented a pragmatic approach to reasoning about concurrent access of resources within a type system. We build upon the Hindley-Damas-Milner type system, incorporating the idea of modelling effects with monads, and extend it with our resourceful constructs. We can prove it retains type soundness, and in the process of doing so we have also created a framework for mechanising the proof of HDM's soundness within Agda. Furthermore, what we have built here serves an excellent basis for future work. There are many different approaches yet to explore that could extend the system in interesting directions.

By allowing heaps to be merged and subsumed into using more resources than necessary, we draw parallels to the local reasoning of separation logic's frame rule. The frame rule is what allows separation logic to scale well, and so we also have good reason to believe that our system will scale for the same reasons.

And although it is simple, the system occupies a point in the design space that fits well into existing functional programming languages, providing a lot of utility for very little complexity. Concurrent programming is a notoriously difficult task, so a type system that can help whittle down cases of concurrent resource access means its one less thing the programmer needs to worry about.

Bibliography

- [1] BARENDREGT, H. P. *The lambda calculus*, vol. 3. North-Holland Amsterdam, 1984.
- [2] CHAPMAN, J., KIREEV, R., NESTER, C., AND WADLER, P. System f in agda, for fun and profit. In *System F in Agda, for fun and profit* (2019), vol. International Conference on Mathematics of Program Construction, Springer, pp. 255–297.
- [3] DAMAS, L. *Type assignment in programming languages*. PhD thesis, The University of Edinburgh, 1984.
- [4] DAMAS, L., AND MILNER, R. Principal type-schemes for functional programs. In *Principal Type-Schemes for Functional Programs*. (1982), vol. POPL 82, ACM, pp. 207–212.
- [5] DOLAN, S., AND MYCROFT, A. Polymorphism, subtyping, and type inference in msub. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages* (2017), pp. 60–72.
- [6] DUBOIS, C. Proving ml type soundness within coq. In *Proving ML type soundness within Coq* (2000), vol. International Conference on Theorem Proving in Higher Order Logics, Springer, pp. 126–144.
- [7] EISENBERG, R. A. *Dependent types in haskell: Theory and practice*. University of Pennsylvania, 2016.
- [8] KOKKE, W., SIEK, J. G., AND WADLER, P. Programming language foundations in agda. *Science of Computer Programming* (2020), 102440.
- [9] KRISHNASWAMI, N. Separation logic for a higher-order typed language. In *Separation logic for a higher-order typed language* (2006), pp. 73–82. Presented at Workshop on Semantics, Program Analysis and Computing Environments for Memory Management, SPACE 2006.
- [10] MARTIN-LÖF, P., AND SAMBIN, G. *Intuitionistic type theory*, vol. 9. Bibliopolis Naples, 1984.
- [11] MILNER, R. A theory of type polymorphism in programming. *Journal of computer and system sciences* 17, 3 (1978), 348–375.

- [12] NORELL, U. Dependently typed programming in agda. In *Advanced Functional Programming: Lecture Notes in Computer Science*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 230–266.
- [13] O’HEARN, P. Resources, concurrency, and local reasoning. *Theoretical computer science* 375, 1-3 (2007), 271–307.
- [14] O’HEARN, P. Separation logic. *Communications of the ACM* 62, 2 (2019), 86–95.
- [15] PEYTON JONES, S. L., AND WADLER, P. Imperative functional programming. In *Imperative functional programming* (1993), vol. Proceedings of the 20th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, pp. 71–84.
- [16] REYNOLDS, J. C. Separation logic: A logic for shared mutable data structures. In *Separation logic: A logic for shared mutable data structures* (2002), vol. Proceedings 17th Annual IEEE Symposium on Logic in Computer Science, IEEE, pp. 55–74.
- [17] TOFTE, M. *Operational semantics and polymorphic type inference*. PhD thesis, The University of Edinburgh, 1988.
- [18] WADLER, P. Monads for functional programming. In *Monads for functional programming* (1995), vol. International School on Advanced Functional Programming, Springer, pp. 24–52.
- [19] WEIRICH, S., VOIZARD, A., DE AMORIM, P. H. A., AND EISENBERG, R. A. A specification for dependent types in haskell. *Proceedings of the ACM on Programming Languages* 1, ICFP (2017), 31.
- [20] WRIGHT, A. K., AND FELLEISEN, M. A syntactic approach to type soundness. *Information and computation* 115, 1 (1994), 38–94.

Appendix A

Denotational semantics of Hindley-Damas-Milner

The *semantic domain* of the expression language defines the possible values an expression can have. It is a *complete partial order*, defined as

$$\begin{aligned}\mathbb{V} &= \mathbb{B}_0 + \dots + \mathbb{B}_n + \mathbb{F} + \mathbb{W} \\ \mathbb{F} &= \mathbb{V} \rightarrow \mathbb{V} \\ \mathbb{W} &= \{\cdot\}\end{aligned}$$

\mathbb{F} is the set of functions, \mathbb{W} is the set of error values and \mathbb{V} is the set of possible values. $\mathbb{B}_0, \dots, \mathbb{B}_n$ represent the sets of “basic” values, such as booleans and naturals.

A complete partial order (cpo) is a pair (D, \sqsubseteq) consisting of a set D and a partial order \sqsubseteq (a function that orders elements in D , but not necessarily all of them, hence the term partial), such that

1. there is a least element \perp
2. each directed subset $x_0 \sqsubseteq \dots \sqsubseteq x_n \sqsubseteq \dots$ has a least upper bound (*lub*)

Since \mathbb{V} represents all possible data values, we can extract an *ideal* of it to model the values of certain types. A subset I of our cpo \mathbb{V} is called an ideal, iff it satisfies the following properties:

1. it is downwards closed: $\forall v_0 \in \mathbb{V}, v_1 \in \mathbb{V}, v_0 \sqsubseteq v_1 \rightarrow v_0 \in I \rightarrow v_1 \in I$.
2. it is closed under lubs of ω -chains $(v_0 \sqsubseteq v_1 \sqsubseteq v_2 \sqsubseteq \dots)$

Bottom \perp is useful for representing programs that do not terminate, like below:

$$\text{let } x = \lambda y.y \text{ in } x \ x$$

What type does should this program have? It should assume the type of whatever is needed. For instance, we would expect this program to have a type of \square as the argument is not used.

$$(\lambda z. \square) (\text{let } x = \lambda y. y \text{ in } x \ x) : \square$$

An environment η is defined as a map from identifiers to values

$$\text{Environment} : \text{Id} \rightarrow \mathbb{V}$$

The semantic equation for the expression language describes how the syntax (enclosed in $\llbracket \ \rrbracket$) is evaluated in a given environment.

$$\begin{aligned} \mathcal{E} &: \text{Expression} \rightarrow \text{Environment} \rightarrow \mathbb{V} \\ \mathcal{E}[\llbracket x \rrbracket] \eta &= \eta[x] \\ \mathcal{E}[\llbracket e_1 e_2 \rrbracket] \eta &= \begin{cases} \perp & \text{if } v_1 = \perp \\ (v_1 | \mathbb{F}) v_2 & \text{if } v_1 \in \mathbb{F} \\ \text{wrong} & \text{otherwise} \end{cases} \\ &\quad \text{where } v_i = \mathcal{E}[\llbracket e_i \rrbracket] \eta, \ i = \{1, 2\} \\ \mathcal{E}[\llbracket \lambda x. e \rrbracket] \eta &= (\lambda v. \mathcal{E}[\llbracket e \rrbracket] \eta[v/x]) \text{ in } \mathbb{V} \\ \mathcal{E}[\llbracket \text{let } x = e_1 \text{ in } e_2 \rrbracket] \eta &= \mathcal{E}[\llbracket e_2 \rrbracket] \eta[\mathcal{E}[\llbracket e_1 \rrbracket] \rho/x] \end{aligned}$$

The notation used here is described in [11, 3]. There is also a type evaluation function $\mathcal{T} : \text{Type} \rightarrow \text{Valuation} \rightarrow \overline{\mathbb{V}}$

$$\begin{aligned} \mathcal{T}[\llbracket \square \rrbracket] \psi &= \mathbb{B}_{\square} \\ \mathcal{T}[\llbracket \text{Bool} \rrbracket] \psi &= \mathbb{B}_{\text{Bool}} \\ \mathcal{T}[\llbracket \alpha \rrbracket] \psi &= \psi[\alpha] \\ \mathcal{T}[\llbracket \tau \rightarrow \tau' \rrbracket] \psi &= \mathcal{T}[\llbracket \tau \rrbracket] \psi \rightarrow \mathcal{T}[\llbracket \tau' \rrbracket] \psi \end{aligned}$$

Where $\overline{\mathbb{V}}$ is the set of all ideals in \mathbb{V} that do not contain **wrong**. Given a valuation, we can define a relation between values and types.

$$v :_{\psi} \tau \iff v \in \mathcal{T}[\llbracket \tau \rrbracket] \psi$$

With it, we can also define a relation between environments η and type environments Γ (contexts), that says an environment respects a type environment if each binding in the context has the same type in the environment.

$$\eta :_{\psi} \Gamma \iff \forall x : \tau \in \Gamma \ \eta[x] :_{\psi} \tau$$

Finally, we can then define what semantic entailment means given a type environment, expression and type.

$$\Gamma \vDash e : \tau \iff \forall \psi \forall \eta \ \eta :_{\psi} \Gamma \rightarrow \mathcal{E}[\llbracket e \rrbracket] \eta :_{\psi} \tau$$

Appendix B

Proofs

B.1 Transitivity of the subheap relation

Proof. We want to show $c \geq a$. Proceed with induction on $c \geq b$.

TOP c must be World, so from TOP we have $c \geq a$.

REFL By definition of REFL, $b = c$, and so we get $c \geq a$ from $b \geq a$.

UNIONL c is of the form $\rho' \cup \rho''$, and from the premise we have $\rho' \geq b$. Use the induction hypothesis with $b \geq a$ and $\rho' \geq b$ to get $\rho' \geq a$, and then UNIONL gives us $\rho' \cup \rho'' \geq a$.

UNIONR c is of the form $\rho'' \cup \rho'$, and from the premise we have $\rho'' \geq b$. Use the induction hypothesis with $b \geq a$ and $\rho'' \geq b$ to get $\rho'' \geq a$, and then UNIONR gives us $\rho'' \cup \rho' \geq a$.

□

B.2 Helper lemmas

Lemma 10. *If $\Gamma \vdash \llbracket e \rrbracket : \text{IO}_\rho \tau$, then $\Gamma \vdash e : \tau$.*

Proof. This might seem obvious, but because of subsumption we need to unravel for the proof for it first. Begin with induction on $\Gamma \vdash \llbracket e \rrbracket : \text{IO}_\rho \tau$.

LIFT Straight from the premises, $\Gamma \vdash e : \tau$.

SUB We have the premise $\Gamma \vdash \llbracket e \rrbracket : \text{IO}_{\rho'} \tau$. Just put this back into the induction hypothesis to get $\Gamma \vdash e : \tau$.

□

Lemma 11. *If $\Gamma \vdash \llbracket e \rrbracket_r : \text{IO}_\rho \tau$, then $\Gamma \vdash e : \tau$.*

Proof. Identical to that of Lemma 10, except we handle the case USE instead of LIFT. □