Tackling the Awkward Squad: monadic input/output, concurrency, exceptions, and foreign-language calls in Haskell

Simon Peyton Jones
Microsoft Research, Cambridge
simonpj@microsoft.com
http://research.microsoft.com/users/simonpj

1st August 2000

Abstract

Functional programming may be beautiful, but to write real applications we must grapple with awkward real-world issues: input/output, robustness, concurrency, and interfacing to programs written in other languages.

These draft lecture notes, for the 2000 Marktoberdorf Summer School, give an overview of the techniques that have been developed by the Haskell community to address these problems. I introduce various proposed extensions to Haskell along the way, and I offer an operational semantics that explains what these extensions mean.

I have to deliver camera-ready copy of the final version of these notes in early November 2000. If you have any suggestions for how I might improve them before then, I would be absolutely delighted to receive them. (My email address is above.) I am particularly interested to know whether there are any passages in which you got lost, or felt that you missed the point, so that I can try to make them clearer.

The Powerpoint slides I used for the lectures are also available at the above URL.

1 Introduction

There are lots of books about functional programming [Thompson, 1999, Hudak, 2000, Bird and Wadler, 1988]. They tend to concentrate on the beautiful core of functional programming: higher order functions, algebraic data types, polymorphic type systems, and so on. These lecture notes are about the bits that usually *aren't* written about. To write programs that are *useful* as well as *beautiful*, the programmer must, in the end, confront the Awkward Squad, a range of un-beautiful but crucial issues, generally concerning interaction with the external world:

- Input and output.
- Error detection and recovery; for example, perhaps the program should time out if something does not happen in time.
- Concurrency, when the program must react in a timely way to independent input sources.
- Interfacing to libraries or components written in some other language.

The call-by-value (or strict) family of functional languages have generally taken a pragmatic approach to these questions, mostly by adopting a similar approach to that taken by imperative

languages. You want to print something? No problem; we'll just have a function `printChar` that has the side effect of printing a character. Of course, `printChar` isn't really a function any more (because it has a side effect), but in practice this approach works just fine, provided you are prepared to specify order of evaluation as part of the language design — and that is just what almost all other programming languages do, from FORTRAN and Java to mostly-functional ones like Lisp, and Standard ML.

Call-by-need (or lazy) languages, such as Haskell, wear a hair shirt because their evaluation order is deliberately unspecified. Suppose that we were to extend Haskell by adding side-effecting "functions" such as `putChar`. Now consider this list

```
xs = [printChar 'a', printChar 'b']
```

(The square brackets and commas denote a list in Haskell.) What on earth might this mean? In SML, evaluating this binding would print `'a'` followed by `'b'`. But in Haskell, the calls to `printChar` will only be executed if the elements of the list are evaluated. For example, if the only use of `xs` is in the call `(length xs)`, then nothing at all will be printed, because `length` does not touch the elements of the list.

The bottom line is that *laziness* and *side effects* are, from a practical point of view, incompatible. If you want to use a lazy language, it pretty much has to be a *purely* functional language; if you want to use side effects, you had better use a strict language.

For a long time this situation was rather embarassing for the lazy community: even the input/output story for purely-functional languages was weak and unconvincing, let alone error recovery, concurrency, etc. Over the last few years, a surprising solution has emerged: the monad[1]. Using monads we have found how to structure programs that perform input/output so that we can, in effect, do imperative programming where that is what we want, and only where we want.

The "standard" version of Haskell is Haskell 98, which comes with an I/O library that uses the monadic approach. However, Haskell 98 is not rich enough to deal with rest of the Awkward Squad (exceptions, concurrency, etc), so we have extended Haskell 98 in a number of experimental ways, adding support for concurrency [Peyton Jones et al., 1996], exceptions [Peyton Jones et al., 1999, Marlow et al., 2000], and a foreign-language interface [Peyton Jones et al., 1998, Finne et al., 1999]. So far, these developments have mostly been documented in scattered research papers; my purpose in these lectures is to gather some of it together into a coherent account. In what follows, when I refer to "Haskell", I will always mean Haskell 98, rather than earlier versions of the language, unless otherwise specified.

2 A web server

As a motivating example, we will explore the issues involved in writing a web server in Haskell. It makes an interesting case study because it involves every member of the Awkward Squad:

- It is I/O intensive.
- It requires concurrency.
- It requires interaction with pre-existing low-level I/O libraries.
- It requires robustness. Dropped connections must time out; it must be possible to reconfigure the server without dropping running connections; errors must be logged.

[1] I say "surprising" because anything with as exotic a name as "monad" — derived from category theory, one of the most abstract branches of mathematics — is unlikely to be very useful to red-blooded programmers. But one of the joys of functional programming is the way in which apparently-exotic theory can have a direct and practical application, and the monadic story is a good example.

The Haskell web server we use as a case study is remarkably small. It uses only 1500 lines of Haskell to implement (more than) the HTTP/1.1 standard. It is robust enough to run continuously for weeks at a time, and its performance is broadly comparable with the widely-used Apache server. (Apache handles 950 connections/sec on the machine we used, while the Haskell web server handles 700 connections/sec. But this is a bit of an apples-and-oranges comparison: on the one hand Apache has much more functionality while, on the other, the Haskell web server has had very little performance tuning applied.)

3 Input and output

The first member of the Awkward Squad is input/output, and that is what we tackle first.

3.1 The problem

We begin with an apparently fundamental conflict. A purely functional program implements a *function*; it has no side effect. Yet the ultimate purpose of running a program is invariably to cause some side effect: a changed file, some new pixels on the screen, a message sent, or whatever. Indeed it's a bit cheeky to call input/output "awkward" at all. I/O is the raison d'être of every program. — a program that had no observable effect whatsoever (no input, no output) would not be very useful.

Well, if the side effect can't be in the functional program, it will have to be outside it. For example, perhaps the functional program could be a function mapping an input character string to an output string:

```
main :: String -> String
```

Now a "wrapper" program, written in (gasp!) C, can get an input string from somewhere (a specified file, for example, or the standard input), apply the function to it, and store the result string somewhere (another file, or the standard output). Our functional programs must remain pure, so we locate all sinfulness in the "wrapper".

The trouble is that one sin leads to another. What if you want to read more than one file? Or write more than one file? Or delete files, or open sockets, or sleep for a specified time, ... ? The next alternative, and one actually adopted by the first version of Haskell, is to enrich the argument and result type of the main function:

```
main :: [Response] -> [Request]
```

Now the program takes as its argument a (lazy) list of `Response` values and produces a (lazy) list of `Request` values. Informally a `Request` says something like "please get the contents of file `/etc/motd`", while a `Response` might say "the contents you wanted is `No email today`". More concretely, `Request` and `Response` are both ordinary algebraic data types, something like this:

```
type Filename = String

data Request = ReadFile Filename
             | WriteFile Filename String
             | ....

data Response = RequestFailed
              | ReadSucceeded String
              | WriteSucceeded
              | ...
```

There is still a wrapper program, as before. It repeatedly takes a request off the result list, acts

on the request, and attaches an appropriate response to the argument list. There has be be some clever footwork to deal with the fact that the function has to be applied to a list of responses before there *are* any responses in the list, but that isn't a problem in a lazy setting.

This request/response story is expressive enough that it was adopted as the main input/output model in the first version of Haskell, but it has several defects:

- It is hard to extend. New input or output facilities can be added only by extending the `Request` and `Response` types, and by changing the "wrapper" program. Ordinary users are unlikely to be able to do this.
- There is no very close connection between a request and its corresponding response. It is extremely easy to write a program that gets one or more "out of step".
- Even if the program remains in step, it is easy to accidentally evaluate the response stream too eagerly, and thereby block emitting a request until the response to that request has arrived – which it won't.

Rather than elaborate on these shortcomings, we move swiftly on to a better solution.

3.2 Monadic I/O

Here is the key idea behind monadic I/O:

> A value of type `IO a` is an "action" that, when performed, may do some input/output, before delivering a value of type `a`.

This is an admirably abstract statement, and I would not be surprised if it means almost nothing to you at the moment. So here is another, more concrete way of looking at these "actions":

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

This type definition says that a value of type `IO a` is a function that, when applied to an argument of type `World`, delivers a new `World` together with a result of type `a`. The idea is rather program-centric: the program takes the state of the entire world as its input, and delivers a modified world as a result, modified by the effects of running the program. I will say in Section 4.1 why I don't think this view of `IO` actions as functions is entirely satisfactory, but it generates many of the right intuitions, so we will use it unashamedly for a while. Indeed we may visualise a value of type `IO a` like this:

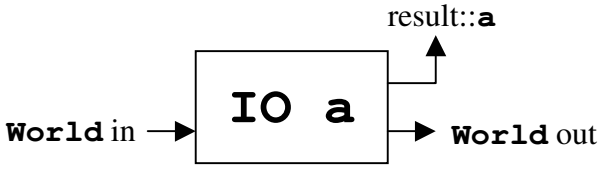

The `World` is fed in on the left, while the new `World`, and the result of type `a`, emerge on the right. In general, we will call a value of type `IO a` an *I/O action* or just *action*. In the literature you will often also find them called *computations*.

We can give `IO` types to some familiar operations, which are supplied as primitive:

```
getChar :: IO Char
putChar :: Char -> IO ()
```

`getChar` is an I/O action that, when performed, reads a character from the standard input (thereby having an effect on the world outside the program), and returns it to the program as the result

of the action. `putChar` is a function that takes a character and returns an action that, when performed, prints the character on the standard output (its effect on the external world), and returns the trivial value `()`. The pictures for these actions look like this (the box for `putChar` takes an extra input for the `Char` argument):

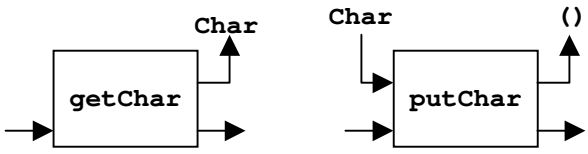

Suppose we want to read a character, and print the character we have read. Then we need to glue together `putChar` and `getChar` into a compound action, like this:

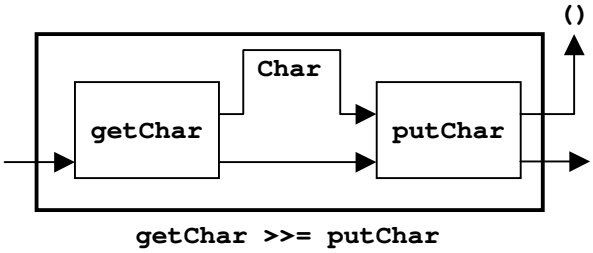

getChar >>= putChar

To achieve this we use a glue function, or combinator, also provided as primitive:

```
(>>=) :: IO a -> (a -> IO b) -> IO b

echo :: IO ()
echo = getChar >>= putChar
```

When the action `(getChar >>= putChar)` is performed, it first performs the action `getChar`, applies `putChar` to the result, and performs the resulting action.

Suppose that we wanted to perform `echo` twice in succession. We can't say `(echo >>= echo)`, because `(>>=)` expects a *function* as its second argument, not an action. Indeed, we want to throw away the result, `()`, of the first `echo`. It is convenient to define a second glue combinator, `(>>)`, in terms of the first:

```
(>>) :: IO a -> IO b -> IO b
(>>) a1 a2 = a1 >>= (\x -> a2)
```

The term `(\x -> a2)` is Haskell's notation for a lambda abstraction. This particular abstraction simply consumes the argument, `x`, throws it away, and returns `a2`. Now we can write

```
echoTwice :: IO ()
echoTwice = echo >> echo
```

In practice, it is very common for the second argument of `(>>=)` to be an explicit lambda abstraction. For example, here is how we could read a character and print it twice:

```
echoDup :: IO ()
echoDup = getChar >>= (\c -> (putChar c >> putChar c))
```

All the parentheses in this example are optional, because a lambda abstraction extends as far to the right as possible, and you will often see this laid out like this:

```
echoDup :: IO ()
echoDup = getChar    >>= \c ->
          putChar c  >>
          putChar c
```

The fact that this looks a bit like a sequence of imperative actions is no coincidence — that is exactly what we wish to specify. Indeed, in Section 3.3 we will introduce special syntax to mirror an imperative program even more closely.

How could we write an I/O action that reads two characters, and returns both of them? We can start well enough:

```
getTwoChars :: IO (Char,Char)
getTwoChars = getChar    >>= \c1 ->
              getChar    >>= \c2 ->
              ???
```

But what are we to put for the "`???`" part? It must be of type `IO (Char,Char)`, but we have done all the input/output required. What we need is one more combinator:

```
return :: a -> IO a
```

The action `(return v)` is an action that does no I/O, and immediately returns `v` without having any side effects. We may draw its picture like this:

Now we can easily complete `getTwoChars`:

```
getTwoChars :: IO (Char,Char)
getTwoChars = getChar    >>= \c1 ->
              getChar    >>= \c2 ->
              return (c1,c2)
```

Here is a more realistic action that reads a whole line of input:

```
getLine :: IO [Char]
getLine = getChar     >>= \c ->
          if c == '\n' then
              return []
          else
              getLine     >>= \cs ->
              return (c : cs)
```

In this example, the "`[]`" is Haskell's notation for the empty list, while the infix constructor "`:`" is the list constructor.

A complete Haskell program defines a single big I/O action, called `main`, of type `IO ()`. The program is executed by performing the action. Here, for example, is a program that reads a complete line from the input, reverses it, and prints it on the output:

```
main :: IO ()
main = getLine        >>= \ cs ->
       putLine (reverse cs)
```

We have not yet defined `putLine :: [Char] -> IO ()`; we leave it as an exercise.

3.3 do notation

Rather than make you write programs in the stylised form of the last section, Haskell provides a special syntax, dubbed "the do notation", for monadic computations. Using the do notation we can write getTwoChars as follows:

```
getTwoChars :: IO (Char,Char)
getTwoChars = do { c1 <- getChar ;
                   c2 <- getChar ;
                   return (c1,c2)
                 }
```

You can leave out the "c <-" part when you want to throw away the result of the action:

```
putTwoChars :: (Char,Char) -> IO ()
putTwoChars (c1,c2) = do { putChar c1; putChar c2 }
```

The syntax is much more convenient than using (>>=) and lambdas, so in practice everyone uses do notation for I/O-intensive programs in Haskell. But it is just notation! The compiler translates the do notation into calls to (>>=), just as before. The translation rules are simple[2]:

$$\begin{array}{rcl} \texttt{do \{ } x \texttt{ <- } e\texttt{; } s \texttt{ \}} & = & e \texttt{ >>= \textbackslash} x \texttt{-> do \{ } s \texttt{ \}} \\ \texttt{do \{ } e\texttt{; } s \texttt{ \}} & = & e \texttt{ >> do \{ } s \texttt{ \}} \\ \texttt{do \{ } e \texttt{ \}} & = & e \end{array}$$

A do expression can appear anywhere that an expression can. Here, for example, is getLine in do notation; it uses a nested do expression:

```
getLine :: IO [Char]
getLine = do { c <- getChar ;
               if c == '\n' then
                 return []
               else
                 do { cs <- getLine ;
                      return (c:cs)
          }      }
```

3.4 Control structures

If monadic I/O lets us do imperative programming, what corresponds to the control structures of imperative languages: for-loops, while-loops, and so on? In fact, we do not need to add anything further to get them: we can build them out of functions.

For example, after some initialisation our web server goes into an infinite loop, awaiting service requests. We can easily express an infinite loop as a combinator:

```
forever :: IO () -> IO ()
forever a = a >> forever a
```

So (forever a) is an action that repeats a forever. Suppose instead that we want to repeat a given action a specified number of times. That is, we want a function:

```
repeatN :: Int -> IO a -> IO ()
```

So (repeatN n a) is an action that, when performed, will repeat a n times. It is easy to define:

```
repeatN 0 a = return ()
repeatN n a = a >> repeat (n-1) a
```

[2] Haskell also allows a let form in do notation, but we omit that for brevity.

Notice that `forever` and `repeatN`, like `(>>)` and `(>>=)`, take an action as one of their arguments. It is this ability to treat an action as a first class value that allows us to define own control structures. Next, a `for` loop:

```
for :: [a] -> (a -> IO ()) -> IO ()
```

The idea is that `(for ns fa)` will apply the function `fa` to each element of `ns` in turn, in each case giving an action; these actions are then combined in sequence.

```
for []     fa = return ()
for (n:ns) fa = fa n >> for ns fa
```

We can use `for` to print the numbers between 1 and 10, thus:

```
printNums = for [1..10] (\n -> print n)
```

(Here, `[1..10]` is Haskell notation for the list of integers between 1 and 10; and `print` has type `Int -> IO ()`.) Another way to define `for` is this:

```
for ns fa = sequence_ (map fa ns)
```

Here, `map` applies `fa` to each element of `ns`, giving a list of actions; then `sequence_` combines these actions together in sequence. So `sequence_` has the type

```
sequence_ :: [IO a] -> IO ()
```

We call this function "`sequence_`" because it has a close cousin, with an even more beautiful type:

```
sequence :: [IO a] -> IO [a]
```

It takes a list of actions, each returning a result of type `a`, and glues them together into a single compound action returning a result of type `[a]`. It is easily defined:

```
sequence []     = return []
sequence (a:as) = do { r  <- a;
                       rs <- sequence as ;
                       return (r:rs) }
```

The "`_`" in "`sequence_`" reminds us that it throws away the results of the sub-actions, returning only `()`.

Notice what is happening here. Instead of having a fixed collection of control structures provided by the language designer, we are free to invent new ones, perhaps application-specific, as the need arises.

3.5 References

The `IO` operations so far allow us to write programs that do input/output in strictly-sequentialised, imperative fashion. It is natural to ask whether we can also model another pervasive feature of imperative languages, namely mutable variables. Taking inspiration from ML's `ref` types, we can proceed like this:

```
data IORef a  -- A built-in type
newIORef   :: a -> IO (IORef a)
readIORef  :: IORef a -> IO a
writeIORef :: IORef a -> a -> IO ()
```

A value of type `IORef a` is a reference to a mutable cell holding a value of type `a`. A new cell can be allocated using `newIORef`, supplying an initial value. Cells can be read and written using `readIORef` and `writeIORef`. Here is a small loop to compute the values between `1` and `n` in an imperative style:

```
count :: Int -> IO Int
count n = do { r <- newIORef 0 ;
               loop 1 r }
        where
          loop i r | i==n      = readIORef r
                   | otherwise = do { v <- readIORef r ;
                                      writeIORef r (v+i) ;
                                      loop (i+1) 1 }
```

This is an absolutely terrible example! For a start, the program is much longer and clumsier than it would be in a purely-functional style. Moreover, it purports to need the `IO` monad but does not really require any side effects at all. Thus, the `IO` monad enables us to transliterate an imperative program into Haskell, but if that's what you want to do, it would be better to use an imperative language in the first place!

Nevertheless, an `IORef` is often useful to "track" the state of some external-world object. For example, Haskell 98 provides a direct analogy of the Standard C library functions for opening, reading, and writing a file:

```
openFile :: [Char] -> IOMode -> IO Handle
hPutStr  :: Handle -> [Char] -> IO ()
hGetLine :: Handle -> IO [Char]
hClose   :: Handle -> IO ()
```

Now, suppose you wanted to record how many characters were read or written to a file. A convenient way to do this is to arrange that `hPutStr` and `hGetLine` each increment a mutable variable suitably. The `IORef` can be held in a modified Handle:

```
type HandleK = (Handle, IORef Int)
```

Now we can define variants of `hPutStr` and `hGetLine` that take a `HandleK` instead of a `Handle` and modify the mutable variable apropriately. For example:

```
hPutStrK :: HandleK -> [Char] -> IO ()
hPutStrK (h,r) cs = do { v <- readIORef r ;
                         writeIORef r (v + length cs) ;
                         hPutStr h cs }
```

In this example, the mutable variable models (part of) the state of the file being written to, by tracking the number of characters written to the file. Since the file itself is, in effect, an external mutable variable, it is not surprising that an internal mutable variable is appropriate to model its state.

3.6 A quick review

Let us summarise what we have learned so far:

- A complete Haskell program is a single (perhaps large) I/O action called `main`.
- Big I/O actions are built by gluing together smaller actions using `(>>=)`.
- An I/O action is a first-class value. It can be passed to a function as an argument, or returned as the result of a function call (consider `(>>)`, for example). Similarly, it can be stored in a data structure — consider the argument to `sequence`, for example.
- The fact that I/O actions can be passed around so freely makes it easy to define new "glue" combinators in terms in existing ones.

Monads were originally invented in a branch of mathematics called category theory, which is increasingly being applied to describe the semantics of programming languages. Eugenio Moggi first identified the usefulness of monads to describe composable "computations" [Moggi, 1989]. Moggi's work, while brilliant, is not for the faint hearted. For practical programmers the breakthrough came in Phil Wadler's paper "Comprehending monads" [Wadler, 1992a], in which he described the usefulness of monads in a programming context. Wadler wrote several more very readable papers about monads, which I highly recommend [Wadler, 1992b, Wadler, 1995, Wadler, 1997]. He and I built directly on this work to write the first paper about monadic I/O [Peyton Jones and Wadler, 1993].

In general, a *monad* is a triple of a type constructor M, and two functions, `return` and `>>=`, with types

$$\begin{array}{rcl} \texttt{return} & :: & \forall\alpha.\ \alpha \to M\ \alpha \\ \texttt{>>=} & :: & \forall\alpha\beta.\ M\ \alpha \to (\alpha \to M\ \beta) \to M\ \beta \end{array}$$

That is not quite all: these three must satisfy the following algebraic laws:

$$\texttt{return}\ x\ \texttt{>>=}\ f \quad = \quad f\ x \qquad (LUNIT)$$

$$m\ \texttt{>>=}\ \texttt{return} \quad = \quad m \qquad (RUNIT)$$

$$\frac{x \notin fv(m_3)}{m_1\ \texttt{>>=}\ (\lambda x.m_2\ \texttt{>>=}\ (\lambda y.m_3)) \quad = \quad (m_1\ \texttt{>>=}\ (\lambda x.m_2))\ \texttt{>>=}\ (\lambda y.m_3)} \qquad (BIND)$$

The last of these rules is much easier to understand when written in `do` notation:

```
do { x <- do { y <- m2;          =    do { y <- m2;
             m1 }                       x <- m1;
     m3 }                               m3
```

The `return` and `(>>=)` of the `IO` monad do indeed satisfy these properties. In these notes I present only one monad, the `IO` monad, but a single program may make use of many different monads, each with its own type constructor, *return* and *bind* operators. Haskell's type class mechanism allows one to overload the functions `return` and `(>>=)` so they can be used in any monad, and the `do` notation can likewise be used for any monad. Wadler's papers, cited above, give many examples of other monads, but we do not have space to pursue that topic here.

3.7 Implementation notes

How difficult is it for a compiler-writer to implement the `IO` monad? There seem to be two main alternatives.

Keep the monad right through. The first technique carries the `IO` monad right through the compiler to the code generator. Most funcional-language compilers translate the source program to an intermediate form based closely on the lambda calculus, apply optimising transformations to that intermediate form, and then generate code. It is entirely possible to extend the intermediate form by adding monadic constructs. One could simply add `(>>=)` and `return` as primitives, but it makes transformation much easier if one adds the `do`-notation directly, instead of a primitive `(>>=)` function. (Compare the two forms of the (BIND) rule given in the previous section.) This is the approach taken by Benton and Kennedy in MLj, their implementation of ML [Benton et al., 1998].

The functional encoding. The second approach, and the one use in the Glasgow Haskell Compiler, is to adopt the functional viewpoint of the `IO` monad, which formed the basis of our earlier pictorial descriptions:

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

If we represent the "world" argument by an un-forgeable token, of type `World`, then we can directly implement `return` and `(>>=)` like this:

```
return :: a -> IO a
return a = \w -> (a,w)

(>>=) :: IO a -> (a -> IO b) -> IO b
(>>=) m k = \w -> case m w of
                    (r,w') -> k r w'
```

Here `w` is the un-forgeable token that stands for the world. In the definition of `(>>=)` we see that the world returned by the first action is passed to the second, just as in the picture in Section 3.2. We must also implement the primitive `IO` operations, such as `getChar`, but that is now no different to implementing other primitive operations, such as addition of two integers.

Notice that neither `return` nor `(>>=)` duplicate the world token: that is, the world token is *single-threaded*. It is this property that allows us to implement `getChar` (and other `IO` primitives) by performing the operation right away — a sort of "update in place".

You might worry that there is an unbounded number of possible I/O "primitives" and you would be right. Indeed users might want to add new I/O operations that cannot be defined in terms of existing ones. What is needed, of course, is a way to call arbitrary I/O libraries supplied by the operating system, a topic we discuss in detail in Section 7.

So which of these two approach is better? Keeping the `IO` monad explicit is principled, but it means that every optimisation pass must deal explicitly with the new constructs. GHC's approach is more economical. For example, the three laws in Section 3.6, regarded as optimisations, are simple consequences and need no special attention. All the same, I have to say that I think the GHC approach is a bit of a hack. Why? Because it relies for its correctness on the fact that the compiler never duplicates a redex. Consider this expression:

```
getChar >>= \c -> (putChar c >> putChar c)
```

If we use GHC's definitions of `(>>=)` we can translate this to:

```
\w -> case getChar w of
        (c,w1) -> case putChar c w1 of
                    (_,w2) -> putChar c w2
```

The compiler would be entirely justified in replacing this code with:

```
\w -> case getChar w of
        (c,w1) -> case putChar c w1 of
                    (_,w2) -> putChar (fst (getChar w)) w2
```

Here I have replaced the second use of `c` with another call to `getChar w`. Two bad things have happened: first, the incoming world token, `w`, has been duplicated; and second, there will now be two calls to `getChar` instead of one. If this happens, our assumption of single-threadedness no longer holds, and neither does our efficient "update-in-place" implementation of `getChar`. Catastrophe!

In the functional language Clean, the whole I/O system is built on an explicit world-passing style. The single-threadedness of the world is ensured by Clean's uniqueness-type system, which

verifies that values which should be single-threaded (notably the world) are indeed used in single threaded way [Barendsen and Smetsers, 1996]. In Haskell, the `IO` monad maintains the world's single-threadedness by construction; so the programmer cannot err, but it is in principle possible for the compiler to do so.

In practice, GHC is careful never to duplicate an expression whose duplication might give rise to extra work (a redex), so it will never duplicate the call to `getChar` in this way. Indeed, Ariola and Sabry have shown formally that if the compiler never duplicates redexes, then indeed our implementation strategy is safe [Ariola and Sabry, 1998]. So GHC's approach is sound, but it is uncomfortable that an apparently semantics-preserving transformation, such as that above, does not preserve the semantics at all. This observation leads us neatly to the next question I want to discuss, namely how to give a semantics to the Awkward Squad.

4 What does it all mean?

It's always a good thing to give a precise semantics to a language feature. How, then, can we give a semantics for the `IO` monad?

4.1 A denotational semantics?

One approach to semantics is to take the functional viewpoint I described earlier:

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

In this view, the meaning of an action is just a function. One can make this story work, but it is a bit unsatisfactory:

- A program that loops for ever without performing any I/O would have the same meaning, namely $\bot$, as a program that deletes a file and then loops forever, or a program that prints "Hello world" repeatedly, although these programs would be regarded as highly distinguishable by a user. Nor is the problem restricted to erroneous programs: Some programs (server processes, for example) may be designed to run essentially forever, and it seems wrong to say that their meaning is simply $\bot$!
- Consider two Haskell programs running in parallel, each sending output to the other — a Web server and a Web browser, for example. The output of each must form part of the `World` given as the input to the other. Maybe it would be possible to deal with this through a fixpoint operator, but it seems complicated and un-intutitive (to me anyway!).
- The approach does not scale well when we add concurrency, which we will do in Section 5.

These problems may be soluble while remaining in a denotational framework, perhaps by producing a sequence of `World`s, or by returning a set of *trace* rather than a new `World`. To give the idea of the trace approach, we model `IO` like this:

```
type IO a = (a, Set Trace)
type Trace = [Event]
data Event = PutChar Char | GetChar Char | ...
```

A program that reads one character, and echoes it back to the screen, would have semantics

```
((), { [GetChar 'a', PutChar 'a'],
       [GetChar 'b', PutChar 'b'],
       [GetChar 'c', PutChar 'c'],
       ... })
```

		x, y	$\in$	*Variable*
		k	$\in$	*Constant*
		c	$\in$	*Constructor*
		ch	$\in$	*Char*
		d	$\in$	*Integer*
	Values	V	::=	$\backslash x$ `->` $M \mid k \mid c\, M_1 \cdots M_n \mid ch \mid d \mid$ `return` $M \mid M$ `>>=` $N \mid$ `putChar` $V \mid$ `getChar`
	Terms	M, N, H	::=	$x \mid V \mid M\, N \mid \cdots$
	Evaluation contexts	E	::=	$[\cdot] \mid$ E `>>=` M

Figure 1: The syntax of values and terms.

We return a *set* of traces, because the trace contains details of *inputs* as well as *outputs*, so there must be a trace for each possible input. The set of traces describes all the behaviours the program can have, and no others. For example `[GetChar 'x', PutChar 'y']` is excluded.

This approach is used to give the semantics of CSP by Roscoe [Roscoe, 1998]. However we will instead adopt an *operational* semantics, based on standard approaches to the semantics of process calculi. Ultimately, I think the two approaches are isomorphic, but I personally find the operational approach simpler and easier to understand.

4.2 An operational semantics

Our semantics is stratified in two levels: an *inner denotational semantics* that describes the behaviour of pure terms, while an *outer monadic transition semantics* that describes the behaviour of `IO` computations. We consider a simplified version of Haskell: our language has the usual features of a lazy functional language (lambda abstraction, application, data structures, case expressions, *etc.*), augmented with constants corresponding to `IO` operations. We will only present those elements of the syntax that are relevant to the semantics; other aspects (such as how we represent lists, or how to write a case expression) wouldn't aid comprehension of the semantics, and are not presented.

In this section, M and N range over terms in our language, and V ranges over all terms that the inner semantics treats as values, such as lambda abstractions and constants (Figure 1). What is less usual is that *we also treat the primitive* `IO` *operations as values.* For example, (M `>>=` N) is a value. It is just as if `(>>=)` was a lazy data constructor. By treating `IO` operations in this way in the inner semantics, we reinforce the distinction between an `IO` action and its effect on the environment when performed. The inner semantics describes the evaluation (to weak head normal form) of an `IO` action; the outer semantics describes its effect when performed.

We will give the semantics by describing how one *program state* evolves into a new program state by making a *transition.* For now, we model a program state simply as a term, but we write it in braces, thus $\{M\}$, to remind us that it is a program state.

$$\{\text{E}[\texttt{putChar}\ ch]\} \xrightarrow{!\,ch} \{\text{E}[\texttt{return}\ ()]\} \quad (PUTC)$$

$$\{\text{E}[\texttt{getChar}]\} \xrightarrow{?\,ch} \{\text{E}[\texttt{return}\ ch]\} \quad (GETC)$$

$$\{\text{E}[\texttt{return}\ N\ \texttt{>>=}\ M]\} \rightarrow \{\text{E}[M\ N]\} \quad (RET/BIND)$$

$$\frac{\mathcal{E}[\![M]\!] = V \quad M \not\equiv V}{\{\text{E}[M]\} \rightarrow \{\text{E}[V]\}}\ (FUN)$$

Figure 2: The basic transition rules

4.3 Labelled transitions

The transition from one program state to the next may or may not be *labelled* by an *event*, α. So we write a transition like this:

$$P \xrightarrow{\alpha} Q$$

The events α represent communication with the external environment; that is, input and output. Initially we will use just two events:

- $P \xrightarrow{!\,ch} Q$ means "program state P can move to Q, by writing the character ch to the standard output".
- $P \xrightarrow{?\,ch} Q$ means "program state P can move to Q, by reading the character ch from the standard input".

Here, then, are our first two transition rules.

$$\{\texttt{putChar}\ ch\} \xrightarrow{!\,ch} \{\texttt{return}\ ()\}$$

$$\{\texttt{getChar}\} \xrightarrow{?\,ch} \{\texttt{return}\ ch\}$$

The first rule says that a program consisting only of `putChar` ch can make a transition, labelled by $!\,ch$, to a progran consisting of `return ()`. The second rule is similar. But most programs consist of more than a single I/O action! What are we to do then? To answer that question we introduce *evaluation contexts*.

4.4 Evaluation contexts

The `getChar` transition rule is all very well, but what if the program consists of more than a single `getChar`? For example, consider the program:

```
main = getChar >>= \c -> putChar c
```

Which is the first I/O action that should be performed? The `getChar`, of course! We need a way to say "the first I/O action to perform is to the left of the `(>>=)`". Sometimes we may have to look to the left of more than one `(>>=)`. Consider the slightly artificial program

```
main = (getChar >>= \c -> getChar) >>= \d -> return ()
```

Here, the first I/O action to be performed is the leftmost `getChar`. In general, to find the first I/O action we "look down the left branch of the tree of `(>>=)` nodes".

We can formalise all this armwaving by using the now well-established notion of an *evaluation context* [Felleisen and Hieb, 1992, Wright and Felleisen, 1994]. The syntax of evaluation contexts

is this (Figure 1):

$$\mathrm{E} \quad ::= \quad [\cdot] \mid \mathrm{E}\ \texttt{>>=}\ M$$

An evaluation context E is a term with a hole, written $[\cdot]$, in it. We write $\mathrm{E}[M]$ to denote the result of filling the hole with the term M. The idea is that the hole is carefully placed to be the site where the next evaluation step must take place. Now we can give the real rules for `putChar` and `getChar`, in Figure 2. In general we will give each transition rule in a Figure, and give it a name — such as (PUTC) and (GETC) — for easy reference.

The rule for (PUTC), for example, should be read: "if a `putChar` occurs as the next I/O action, in a context $\mathrm{E}[\cdot]$, the program can make a transition, emitting a character and replacing the call to `putChar` by `return ()`". This holds for any evaluation context $\mathrm{E}[\cdot]$.

Let us see how to make transitions using our example program:

```
main = getChar >>= \c -> putChar c
```

Using rule (GETC) and the evaluation context (`[·] >>= \c -> putChar c`), we can make the transition:

$$\{\texttt{getChar >>= \c -> putChar c}\} \xrightarrow{?\,ch} \{\texttt{return}\ ch\ \texttt{>>= \c -> putChar c}\}$$

What next? We use the return/bind law of Section 3.6, expressed as a new transition rule:

$$\{\mathrm{E}[\texttt{return}\ N\ \texttt{>>=}\ M]\} \quad \rightarrow \quad \{\mathrm{E}[M\ N]\} \quad (RET/BIND)$$

Using this rule, we make the transition

$$\{\texttt{return}\ ch\ \texttt{>>= \c -> putChar c}\} \rightarrow \{\texttt{(\c -> putChar c)}\ ch\}$$

Now we need to do some ordinary, purely-functional evaluation work. We express this by "lifting" the inner denotational semantics into our transition system, like this:

$$\frac{\mathcal{E}[\![M]\!] = V \qquad M \not\equiv V}{\{\mathrm{E}[M]\} \ \rightarrow\ \{\mathrm{E}[V]\}} \ (FUN)$$

That is, if the term M has value V, as computed by the denotational semantics function $\mathcal{E}[\![]\!]$, then we can replace M by V at the active site[3]. The side condition $M \not\equiv V$ is just there to prevent the rule firing repeatedly without making progress, because $\mathcal{E}[\![V]\!] = V$ for any V. Rule (FUN) allows us to make the following transition, using normal beta reduction:

$$\{\texttt{(\c -> putChar c)}\ ch\} \rightarrow \{\texttt{putChar}\ ch\}$$

Now we can use the `putChar` rule to emit the character:

$$\{\texttt{putChar}\ ch\} \xrightarrow{!\,ch} \{\texttt{return ()}\}$$

And now the program is finished.

Referring back to the difficulties identified in Section 4.1, we can now distinguish a program A that simply loops forever, and a program B that repeatedly prints `'x'` forever. These programs both have denotation $\perp$ in a (simple) denotational semantics (Section 4.1), but they have different behaviours in our operational semantics. Program B will repeatedly make a transition with the label `!x`. But what happens to program A? To put it another way, what happens in rule (FUN)

[3] I am being a bit sloppy here, because a denotational semantics yields a mathematical value, not a term in the original language, but in fact nothing important is being swept under the carpet here. From a technical point of view it may well be simpler, in the end, to adopt an operational semantics for the inner purely-functional part too, but that would be a distraction here. Notice, too, that the valuation function of a denotational semantics would usually have an environment, ρ. But the rule (FUN) only requires the value of a closed term, so the environment is empty.

$$
\begin{array}{rcll}
r & \in & \texttt{IORef} & \\
\\
V & ::= & \ldots \mid \texttt{writeIORef}\ V\ N \mid \texttt{readIORef}\ V \mid \texttt{newIORef}\ M \mid r & \\
\\
P, Q, R & ::= & \{M\} & \text{The main program} \\
& \mid & \langle M \rangle_r & \text{An } \texttt{IORef} \text{ named } r, \text{ holding } M \\
& \mid & P \mid Q & \text{Parallel composition} \\
& \mid & \nu x.P & \text{Restriction}
\end{array}
$$

$$
\begin{array}{rcl}
\{\mathrm{E}[\texttt{readIORef}\ r]\} \mid \langle M \rangle_r & \rightarrow & \{\mathrm{E}[\texttt{return}\ M]\} \mid \langle M \rangle_r \quad (READIO) \\
\{\mathrm{E}[\texttt{writeIORef}\ r\ N]\} \mid \langle M \rangle_r & \rightarrow & \{\mathrm{E}[\texttt{return}\ ()]\} \mid \langle N \rangle_r \quad (WRITEIO)
\end{array}
$$

$$
\frac{r \notin fn\,(\mathrm{E}, M)}{\{\mathrm{E}[\texttt{newIORef}\ M]\} \ \rightarrow\ \nu r.(\{\mathrm{E}[\texttt{return}\ r]\} \mid \langle M \rangle_r)}\ (NEWIO)
$$

Figure 3: Extensions for `IORefs`

if $\mathcal{E}[\![M]\!] = \bot$? The simplest thing to say is that then there is no value V such that $\mathcal{E}[\![M]\!] = V$, and so (FUN) cannot fire. So no rule appplies, and the program is stuck. This constitutes an observably different sequence of transitions than program A[4].

Lastly, before we leave the topic of evaluation contexts, let us note that the term M in rule (FUN) always has type `IO` τ for some type τ; that is, an evaluation context E[·] always has an I/O action in its hole. So there is no need to explain how the program (say) {`True`} behaves, because it is ill-typed.

4.5 Dealing with `IORefs`

Let us now add `IORefs` to our operational semantics. The modifications we need are given in Figure 3:

- We add a new sort of *value* for each `IORef` primitive; namely `newIORef`, `readIORef`, and `writeIORef`.
- We add a new sort of value for `IORef` identifiers, r. An `IORef` identifier is the value returned by `newIORef` — you can think of it as the address of the mutable cell.
- We extend a program state to be a main thread $\{M\}$, as before, together with zero or more `IORefs`, each associated with a reference identifier r.

The syntax for program states in Figure 3 might initially be surprising. We use a vertical bar to join the main thread and the `IORefs` into a program state. For example, here is a program state for a program that has (so far) allocated two `IORefs`, called r_1 and r_2 respectively:

$$\{M\} \mid \langle N_1 \rangle_{r_1} \mid \langle N_2 \rangle_{r_2}$$

If you like, you can think of running the (active) program M in parallel with two (passive) containers r_1 and r_2, containing N_1 and N_2 respectively.

Here are the rules for reading and writing `IORefs`:

$$
\begin{array}{rcl}
\{\mathrm{E}[\texttt{readIORef}\ r]\} \mid \langle M \rangle_r & \rightarrow & \{\mathrm{E}[\texttt{return}\ M]\} \mid \langle M \rangle_r \quad (READIO) \\
\{\mathrm{E}[\texttt{writeIORef}\ r\ N]\} \mid \langle M \rangle_r & \rightarrow & \{\mathrm{E}[\texttt{return}\ ()]\} \mid \langle N \rangle_r \quad (WRITEIO)
\end{array}
$$

[4] You may argue that we should not say that A gets stuck when actually it is in an infinite loop. If you prefer, one can instead add a variant of (FUN) that makes an un-labelled transition to an unchanged state if $\mathcal{E}[\![M]\!] = \bot$. Then A would make an infinite sequence of un-labelled transitions. It's just a matter of taste.

$$
\begin{array}{rcll}
P \mid Q & \equiv & Q \mid P & (COMM) \\
P \mid (Q \mid R) & \equiv & (P \mid Q) \mid R & (ASSOC) \\
\nu x.\nu y.P & \equiv & \nu y.\nu x.P & (SWAP) \\
(\nu x.P) \mid Q & \equiv & \nu x.(P \mid Q), \quad x \notin fn\,(Q) & (EXTRUDE) \\
\nu x.P & \equiv & \nu y.P[y/x], \quad y \notin fn\,(P) & (ALPHA)
\end{array}
$$

$$
\frac{P \xrightarrow{\alpha} Q}{P \mid R \xrightarrow{\alpha} Q \mid R}\ (PAR) \qquad \frac{P \xrightarrow{\alpha} Q}{\nu x.P \xrightarrow{\alpha} \nu x.Q}\ (NU)
$$

$$
\frac{P \equiv P' \quad P' \xrightarrow{\alpha} Q' \quad Q' \equiv Q}{P \xrightarrow{\alpha} Q}\ (EQUIV)
$$

Figure 4: Structural congruence, and structural transitions.

The rule for `readIORef` says that if the next I/O action in the main program is `readIORef` r, and the main program is parallel with an `IORef` named r containing M, then the action `readIORef` r can be replaced by `return` M. This transition is quite similar to that for `getChar`, except that the transition is unlabelled because it is internal to the program — remember that only labelled transitions represent interaction with the external environment.

We have several tiresome details to fix up. First, we originally said that the transitions were for whole program states, but these two are for only *part* of a program state; there might be other `IORef`s, for example. Second, what if the main program was not adjacent to the relevant `IORef`? We want to say somehow that it can become adjacent to whichever `IORef` it pleases. To formalise these matters we have to give several "structural" rules, given in Figure 4. Rule (PAR), for example, says that if P can move to Q, then P in parallel with anything (R) can move to Q in parallel with the same anything — in short, non-participating pieces of the program state are unaffected. The equivalence rules (COMM), (ASSOC) say that $\mid$ is associative and commmutative, while (EQUIV) says that we are free to use these equivalence rules to bring parts of the program state together. In these rules, we take α to range over both events, such as $!\,ch$ and `?`ch, and also over the empty label. (In the literature, you will often see the empty event written τ.)

It's all a formal game. If you read papers about operational semantics you will see these rules over and over again, so it's worth becoming comfortable with them. They aren't optional though; if you want to conduct water-tight proofs about what can happen, it's important to specify the whole system in a formal way.

Here is the rule for `newIORef`:

$$
\frac{r \notin fn\,(\mathrm{E}, M)}{\{\mathrm{E}[\texttt{newIORef}\ M]\} \ \rightarrow\ \nu r.(\{\mathrm{E}[\texttt{return}\ ()]\} \mid \langle M \rangle_r)}\ (NEWIO)
$$

If the next I/O action in the main program is to create a new `IORef`, then it makes a transition to a new state in which the main program is in parallel with a newly-created (and suitably initialised) `IORef` named r. What is r? It is an arbitrary name whose only constraint is that it must not already be used in M, or in the evaluation context E. That is what the side condition $m \notin fn\,(\mathrm{E}, M)$ means — $fn\,(\mathrm{E}, M)$ means "the free names of E and M".

It should be clear that choosing a name that is already in use would be a Bad Thing. The side condition says that r cannot be mentioned in E or M. But what if r was in use somewhere *else* in the program state? That is the purpose of the "νr" part: it restricts the scope of r. Having introduced ν in this way, we need a number of structural rules (Figure 4) to let us move ν around. Notably, (EXTRUDE) lets us move all the ν's to the outside. Before we can use (EXTRUDE), though, we may need to use (ALPHA) to change our mind about the name we chose if we come

across a name-clash. Once all the ν's are at the outside, they don't get in the way at all.

5 Concurrency

A web server works by listening for connection requests on a particular socket. When it receives a request, it establishes a connnection and engages in a bi-directional conversation with the client. Early versions of the HTTP protocol limited this conversation to one utterance in each direction ("please send me this page"; "ok, here it is"), but more recent versions of HTTP allow multiple exchanges to take place, and that is what we do here.

If a web server is to service multiple clients, it must deal concurrently with each client. It is simply not acceptable to deal with clients one at a time. The obvious thing to do is to fork a new *thread* of some kind for each new client. The server therefore must be a *concurrent* Haskell program.

I make a sharp distinction between *parallelism* and *concurrency*:

- A *parallel* functional program uses multiple processors to gain performance. For example, it may be faster to evaluate $e_1 + e_2$ by evaluating e_1 and e_2 in parallel, and then add the results. Parallelism has no semantic impact at all: the meaning of a program is unchanged whether it is executed sequentially or in parallel. Furthermore, the results are deterministic; there is no possiblity that a parallel program will give one result in one run and a different result in a different run.
- In contrast, *concurrent* program has concurrency as part of its specification. The program must run concurrent threads, each of which can independently perform input/output. The program may be run on many processors, or on one — that is an implementation choice. The behaviour of the program is, necessarily and by design, non-deterministic. Hence, unlike parallelism, concurrency has a substantial semantic impact.

Of these two, my focus in these lectures is exclusively on concurrency, not parallelism. For those who are interested, a good introduction to parallel functional programming is [Trinder et al., 1998], while a recent book gives a comprehensive coverage [Hammond and Michaelson, 1999].

Concurrent Haskell [Peyton Jones et al., 1996] is an extension to Haskell 98 designed to support concurrent programming, and we turn next to its design.

5.1 Threads and `forkIO`

Here is the main loop of the web server:

```
acceptConnections :: Config -> Socket -> IO ()
acceptConnections config socket
  = forever (do { conn <- accept socket ;
                  forkIO (serviceConn config conn) })
```

(We defined `forever` in Section 3.4.) This infinite loop repeatedly calls `accept`, a Haskell version of the Unix procedure of the same name, to accept a new connection. `accept` returns, as part of its result, a `Handle` that can be used to communicate with the client.

```
accept :: Socket -> IO Connection

type Connection = (Handle,    -- Read from here
                   SockAddr)  -- Peer details
```

Having established a connnection, `acceptConnections` then uses `forkIO` to fork off a fresh thread, `(serviceConn config conn)`, to service that connection. The type of `forkIO` is this:

```
forkIO :: IO a -> IO ThreadId
```

It takes an I/O action and arranges to run it concurrently with the "parent" thread. The call to `forkIO` returns immediately, returning as its result an identifier for the forked thread. We will see in Section 6.5 what this `ThreadId` can be used for.

Notice that the forked thread doesn't need to be passed any parameters, as is common in C threads packages. The forked action is a full closure that captures the values of its free variables. In this case, the forked action is `(serviceConn config conn)`, which obviously captures the free variables `config` and `conn`.

5.2 Communication and `MVars`

Suppose we want to add some sort of throttling mechanism, so that when there are more than N threads running the server does something different (e.g. stops accepting new connections or something). To implement this we need to keep track of the total number of (active) forked threads. How can we do this? The obvious solution is to have a counter that the forked thread increments when it begins, and decrements when it is done. But we must of course be careful! If there are lots of threads all hitting on the same counter we must make sure that we don't get race hazards. The increments and decrements must be indivisible.

To this end, Concurrent Haskell supports a synchronised version of an `IORef` called an `MVar`:

```
data MVar a   -- Abstract
newEmptyMVar :: IO (MVar a)
takeMVar     :: MVar a -> IO a
putMVar      :: MVar a -> a -> IO ()
```

Like an `IORef`, an `MVar` is (a reference to) a mutable location that either can contain a value of type `a`, *or can instead be empty.* Like `newIORRef`, `newEmptyMVar` creates an `MVar` but, unlike an `IORef`, the `MVar` is created empty.

`putMVar` fills an empty `MVar` with a value, and `takeMVar` takes the contents of an `MVar` out, leaving it empty. If it was empty in the first place, the call to `takeMVar` blocks until another thread fills it by calling `putMVar`. A call to `putMVar` on an `MVar` that is already full blocks until the `MVar` becomes empty[5].

With the aid of `MVars` it is easy to implement our counter:

```
acceptConnections :: Config -> Socket -> IO ()
acceptConnections config socket
  = do { count <- newEmptyMVar ;
         putMVar count 0 ;
         forever (do { conn <- accept socket ;
                       forkIO (do { inc count ;
                                    serviceConn config conn ;
                                    dec count}) }

inc,dec :: MVar Int -> IO ()
inc count = do { v <- takeMVar count; putMVar count (v+1) }
dec count = do { v <- takeMVar count; putMVar count (v-1) }
```

Presumably there would also be some extra code in `acceptLoop` to inspect the value of the counter, and take some action if it gets too large.

The update of the counter, performed by `inc` and `dec` is indivisible because, during the brief

[5] This represents a change from an earlier version of Concurrent Haskell, in which `putMVar` on a full `MVar` was a program error.

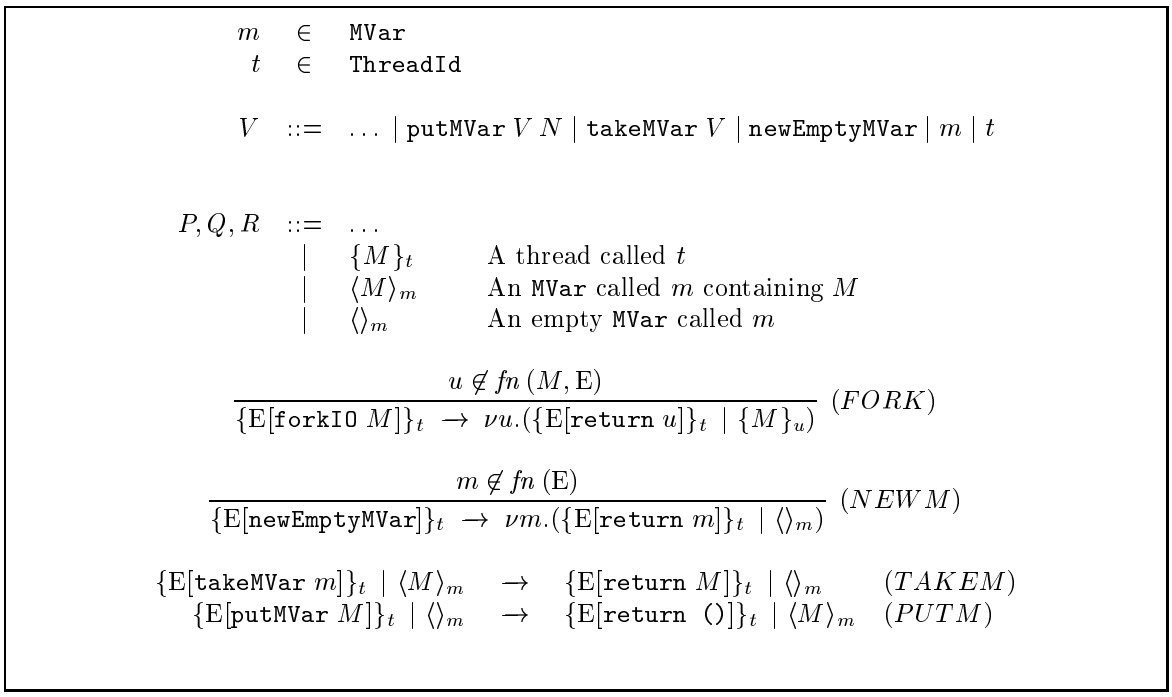

Figure 5: Extensions to support concurrency

moment while `inc` has read the counter but not yet written it back, the counter location is empty. So any threads that try to read it at that moment will simply block.

5.3 Semantics

One of the advantages of the operational semantics we set up in Section 4 is that it can readily be extended to support concurrency and `MVars`. The necessary extensions are given in Figure 5:

- We add new values to represent (a) each new primitive `MVar` operation; (b) the name of an `MVar`; and (c) the name of a thread.
- We extend program states by adding a form for an `MVar`, both in the full state $\langle M \rangle_m$, and in the empty state $\langle\rangle_m$; and a form for a named thread $\{M\}_t$.
- We provide transition rules for the new primitives.

Rules (FORK) and (NEWM) work in a very similar way as the (NEWIO) rule that we described in Section 4.5. In particular, they use ν in an identical fashion to control the new names that are required. Rules (PUTM) and (TAKEM) are similar to (WRITEIO) and (READIO), except that (TAKEM) leaves the `MVar` empty, while (PUTM) fills it.

For the first time, the semantics of the program has become non-deterministic. If there are two threads both of which want to take the contents of an `MVar`, the semantics leaves deliberately unspecified which one "wins". Once it has emptied the `MVar` with rule (TAKEM), however, the other thread can make no progress until some other thread fills it.

Notice that there is no explicit rule for "blocking" a thread when it tries to take the contents of an `MVar` that is empty. All that happens is that there is no valid transition rule involving that thread, so it stays unchanged in the program state until the `MVar` it is trying to take is filled.

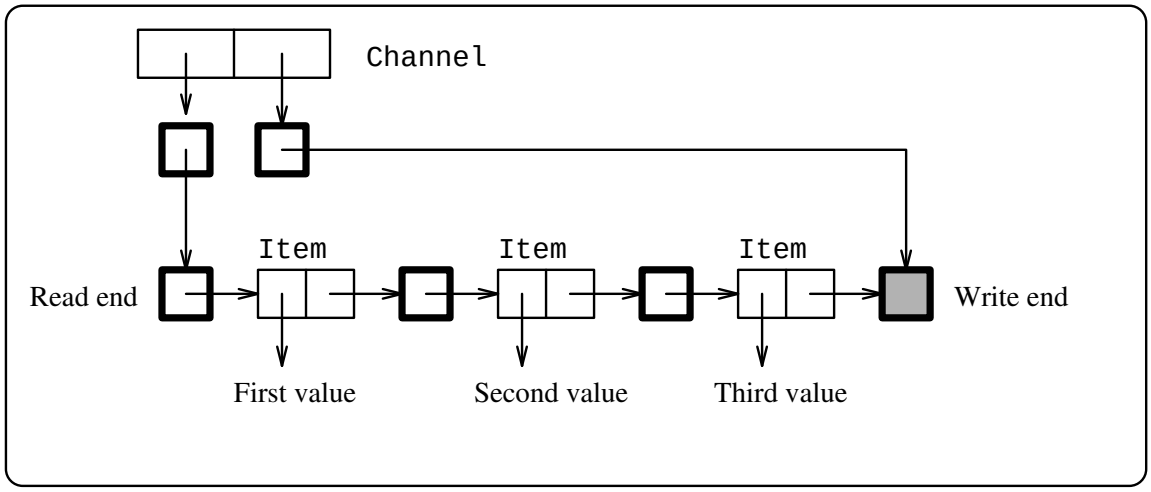

Figure 6: A channel with unbounded buffering

5.4 Channels

The thread created by `forkIO` and its parent thread can each independently perform input and output. We can think of the state of the world as a shared, mutable object, and race conditions can, of course, arise. For example, if two threads are foolish enough to write to the same file, say, bad things are likely to happen.

But what if we *want* to have two threads write to the same file, somehow merging their writes, at some suitable level of granularity? Precisely this behaviour is needed in our web server, because we want to log errors found by the client-service threads to a single error-log file. The simplest thing to do is to create a single thread whose business is to write to the error-log file; to log an error, a client-service thread need only send a message to the error-logging thread. But we have just pushed the problem to a different place: what does it mean to "send a message"?

Using `MVars` we can define a new type of buffered channels, which we will implement in this section:

```
type Channel a
newChan :: IO (Channel a)
putChan :: Channel a -> a -> IO ()
getChan :: Channel a -> IO a
```

A `Channel` permits multiple processes to write to it, and read from it, safely. The error-logging thread can now repeatedly do `getChan`, and write the value it receives into the file; meanwhile a client-service thread wanting to log an error can use `putChan` to send the error message to the error logger.

One possible implementation of `Channel` is illustrated in Figure 6. The channel is represented by a pair of `MVars` (drawn as small boxes with thick borders), that hold the read end and write end of the buffer:

```
type Channel a = (MVar (Stream a),     -- Read end
                  MVar (Stream a))     -- Write end (the hole)
```

The `MVars` in a `Channel` are required so that channel put and get operations can atomically modify the write and read end of the channels respectively. The data in the buffer is held in a `Stream`; that is, an `MVar` which is either empty (in which case there is no data in the `Stream`), or holds an `Item` (a data type we will define shortly):

```
type Stream a = MVar (Item a)
```

An `Item` is just a pair of the first element of the `Stream` together with a `Stream` holding the rest of the data:

```
data Item a = MkItem a (Stream a)
```

A `Stream` can therefore be thought of as a list, consisting of alternating `Item`s and full `MVar`s, terminated with a "hole" consisting of an empty `MVar`. The write end of the channel points to this hole.

Creating a new channel is now just a matter of creating the read and write `MVar`s, plus one (empty) `MVar` for the stream itself:

```
newChan = do { read  <- newEmptyMVar ;
               write <- newEmptyMVar ;
                hole  <- newEmptyMVar ;
               putMVar read hole ;
               putMVar write hole ;
               return (read,write) }
```

Putting into the channel entails creating a new empty `Stream` to become the hole, extracting the old hole and replacing it with the new hole, and then putting an `Item` in the old hole.

```
putChan (read,write) val
  = do { new_hole <- newEmptyMVar ;
         old_hole <- takeMVar write ;
         putMVar write new_hole
         putMVar old_hole (MkItem val new_hole) }
```

Getting an item from the channel is similar. Notice that `getChan` may block at the second `takeMVar` if the channel is empty, until some other process does a `putChan`.

```
getChan (read,write)
  = do { cts <- takeMVar read ;
         MkItem val new_cts <- takeMVar cts ;
         putMVar read new_cts ;
         return val }
```

It is worth noting that any number of processes can safely write into the channel and read from it. The values written will be merged in (non-deterministic, scheduling-dependent) arrival order, and each value read will go to exactly one process.

Other variants are readily programmed. For example, consider a multi-cast channel, in which there are multiple readers, each of which should see all the values written to the channel. All that is required is to add a new operation:

```
dupChan :: Channel a -> IO (Channel a)
```

The idea is that the channel returned by `dupChan` can be read independently of the original, and sees all (and only) the data written to the channel after the `dupChan` call. The implementation is simple, since it amounts to setting up a separate read pointer, initialised to the current write pointer:

```
dupChan (read,write)
  = do { new_read <- newEmptyMVar ;
         hole <- takeMVar write ;
         putMVar write hole ;
         putMVar new_read hole ;
         return (new_read, write) }
```

```
forkIO        ::  IO a -> IO ThreadId
threadDelay   ::  Int -> IO ()            -- Sleep for n microseconds

data MVar a       -- Abstract
newEmptyMVar  ::  IO (MVar a)             -- Created empty
newMVar       ::  a -> IO (MVar a)        -- Initialised

takeMVar      ::  MVar a -> IO a          -- Blocking take
putMVar       ::  MVar a -> a -> IO ()    -- Blocking put

tryTakeMVar   ::  MVar a -> IO (Maybe a)  -- Non-blocking take
tryPutMVar    ::  MVar a -> a -> IO Bool  -- Non-blocking put
isEmptyMVar   ::  MVar a -> IO Bool       -- Test for emptiness
```

Figure 7: The most important concurrent operations

Another easy modification, left as an exercise for the reader, is to add an inverse to `getChan`:

```
unGetChan :: Channel a -> a -> IO ()
```

5.5 Summary

Adding `forkIO` and `MVar`s to Haskell leads to a qualitative change in the sorts of applications one can write. The extensions are simple to describe, and our operational semantics was readily extended to describe them. Figure 7 lists the main operations in Concurrent Haskell, including some that we have not discussed.

You will probably have noticed the close similarity between `IORef`s (Section 3.5) and `MVar`s (Section 5.2). Are they both necessary? Probably not. In practice we find that we seldom use `IORef`s at all:

- Although they have slightly different semantics (an `IORef` cannot be empty) it is easy to simulate an `IORef` with an `MVar` (but not vice versa).
- An `MVar` is not much more expensive to implement than an `IORef`.
- An `IORef` is fundamentally unsafe in a concurrent program, unless you can prove that only one thread can access it at a time.

I introduced `IORef`s in these notes mainly as a presentational device; they allowed me to discuss the idea of updatable locations, and the operational machinery to support them, before getting into concurrency.

While the primitives are simple, they are undoubtedly primitive. `MVar`s are surprisingly often useful "as is", especially for holding shared state, but they are a rather low-level device. However, they provide the raw material from which one can fashion more sophisticated abstractions, and a higher-order language like Haskell is well suited for such a purpose. `Channel`s are an example of such an abstraction, and we give several more in [Peyton Jones et al., 1996]. Einar Karlsen's thesis describes a very substantial application (a programming workbench) implemented in Concurrent Haskell, using numerous concurrency abstractions [Karlsen, 1998].

It is not the purpose of these notes to undertake a proper comparative survey of concurrent programming, but I cannot leave this section without mentioning two other well-developed approaches to concurrency in a declarative setting. Erlang is a (strict) functional language developed at Ericsson for programming telecommunications applications, for which purpose it has been extremely

successful [Armstrong et al., 1996]. Erlang must be the most widely-used concurrent functional language in the world. Concurrent ML (CML) is a concurrent extension of ML, with a notion of first-class events and synchonisation constructs. CML lays particular emphasis on concurrency abstractions, and is well described in Reppy's excellent book [Reppy, 1999].

6 Exceptions and timeouts

The next member of the Awkward Squad is robustness and error recovery. A robust program should not collapse if something unexpected happens. Of course, one tries to write programs in such a way that they will not fail, but this approach alone is insufficient. Firstly, programmers are fallible and, secondly, some failures simply cannot be avoided by careful programming.

Our web server, for example, should not cease to work if

- A file write fails because the disk is full.
- A client requests a seldom-used service, and that code takes the head of an empty list or divides by zero.
- A client vanishes, so the client-service thread should time out and log an error.
- An error in one thread makes it go into an infinite recursion and grow its stack without limit.

All these events are (hopefully) rare, but they are all unpredictable. In each case, though, we would like our web server to recover from the error, and continue to offer service to existing and new clients.

We cannot offer this level of robustness with the facilities we have described so far. We could check for failure on every file operation, though that would be rather tedious. We could try to avoid dividing by zero — but we will never know that we have found every bug. And timeouts and loops are entirely inaccessible.

This is, of course, exactly what exceptions were invented for. An exception handler can enclose an arbitrarily large body of code, and guarantee to give the programmer a chance to recover from errors arising anywhere in that code.

6.1 Exceptions in Haskell 98

Like many languages, Haskell's `IO` monad offers a simple form of exception handling. I/O operations may *raise* an exception if something goes wrong, and that exception can be *caught* by a handler. Here are the primitives that Haskell 98 offers:

```
userError :: String -> IOError
ioError   :: IOError -> IO a
catch     :: IO a -> (IOError -> IO a) -> IO a
```

You can raise an exception by calling `ioError` passing it an argument of type `IOError`. You can construct an `IOError` from a string using `userError`. Finally, you can catch an exception with `catch`. The call `(catch a h)` is an action that, when performed, attempts to perform the action `a` and return its results. However, if performing `a` raises an exception, then `a` is abandoned, and instead `(h e)` is returned, where `e` is the `IOError` in the exception.

Here is an example of how we might extend our main web-server loop to

```
acceptConnections :: Config -> Socket -> IO ()
acceptConnections config socket
```

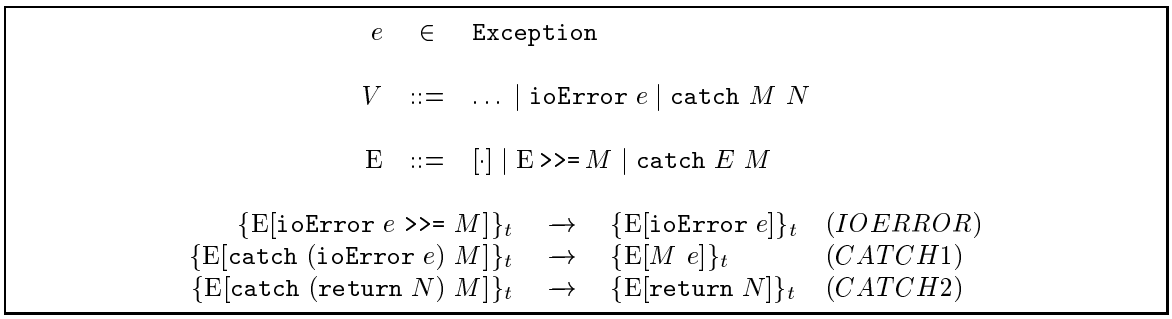

Figure 8: Extensions for exceptions

```
= forever (do { conn <- accept socket ;
                forkIO (service conn) }
where
  service conn = catch (serviceConn config conn)
                       (handler conn)

  handler conn (h,_) = do { logError config e ;
                            hClose h }
```

Now the forked thread `(service conn)` has an exception handler wrapped around it, so that if anything goes wrong, `handler` will be invoked. This handler logs the error (presumably by sending a message to the error-logging thread through a channel held in `config`), and closes the connection handle `h`.

Figure 8 gives the extra semantics required to support Haskell 98 exceptions, in a style that by now will be familiar. The extra evaluation context says that we should evaluate inside a `catch`. Rule (IOERROR) says that a call to `ioError` is propagated by `(>>=)`; this is what correponds to "popping the stack" in a typical implementation. Rules (CATCH1) describes what happens when the exception meets a `catch`: it is passed on to the handler. Lastly, (CATCH2) explains that `catch` does nothing if execution of the protected code terminates normally with `return` N.

6.2 What's wrong with Haskell 98

The Haskell 98 design falls short in two ways:

- It does not handle things that might go wrong in purely-functional code, because an exception can only be raised in the `IO` monad. A pattern-match failure[6], or division by zero, brings the entire program to a halt. We address this problem in Section 6.3
- It does not deal with *asynchronous* exceptions. A synchronous exception arises as a direct result of executing some piece of code — opening a non-existent file, for example. Synchronous exceptions can be raised only at well-defined places. An asynchronous exception, in contrast, is raised by something in the thread's environment: a timeout or user interrupt is an asynchronous exception. It is useful to treat resource exhaustion, such as stack overflow, in the same way. An asynchronous exception can strike at any time, and this makes them much harder to deal with than their synchronous cousins. We tackle asynchronous exceptions in Section 6.5

In the rest of this section we briefly sketch how to handle both of these shortcomings.

[6] A pattern-match failure occurs when a function defined by pattern-matching is applied to a value for which no pattern matches. Example: taking the head of an empty list.

6.3 Synchronous exceptions in pure code

Why does Haskell 98 not allow the program to raise an exception in purely-functional code? The reason is that, as with input/output, Haskell's unconstrained order of evaluation makes it hard to say what the program means. Suppose we invented a new primitive to raise an exception:

```
throw :: Exception -> a
```

(`throw` differs from `ioError` in that it lacks an `IO` on its result type.) There are two difficulties:

- Consider the following expression:

  ```
  length [throw ex1]
  ```

  Does the expression raise exception `ex1`? Since `length` does not evaluate the elements of its argument list, the answer is presumably "no". So *whether an exception is raised depends on how much evaluation takes place.*

- Which exception does the following expression raise, `ex1` or `ex2`?

  ```
  throw ex1 + throw ex2
  ```

  The answer clearly depends on the order in which the arguments to `(+)` are evaluated. So *which exception is raised depends on evaluation order.*

As with input/output (right back in Section 1, one possibility is to fully define evaluation order and, as before, we reject that alternative.

6.3.1 Imprecise exceptions

The best approach is to take the hint from denotational semantics. The purely-functional part of the language should have a straightforward denotational semantics, and that requires us to answer the question: "what value does `throw e` return?". The answer must be "an exceptional value". So we divide the world of values (or denotations) into *ordinary values* (like `'a'` or `True` or `132`) and *exceptional values*. The IEEE Floating Point standard defines certain bit-patterns as "not-a-numbers", or NaNs. A NaN is returned by a floating point operation that fails in some way, such as division by zero. Intel's IA-64 architecture extends this idea to arbitrary data types, using "not-a-thing" (NaT) values to represent the result of speculative operations that have failed. In our terminology, a NaN or NaT is an exceptional value.

So `throw` simply constructs an exceptional value. It is a perfectly well-behaved value provided you never actually evaluate it; only then is the exception raised. The situation is very similar to that for a divergent (non-terminating) expression in a lazy language. Useful programs may contain such values; the program will only diverge if it actually evaluates the divergent term.

That deals with the first bullet, but how about the second? A good solution is to say that the denotation of an expression is

- A single ordinary value, or
- A *set* of exceptions.

By making the denotation into a *set* of exceptions we can finesse the question of which exception is raised if many could be. Let us return to our troublesome example

```
throw ex1 + throw ex2
```

The denotation of this expression is now an exceptional value consisting of a set of two exceptions, `ex1` and `ex2`. In saying this, we do not need to say anything about evaluation order.

I am *not* suggesting that an implementation should actually *construct* the set of exceptions. The idea is that an implementation can use an entirely conventional exception-handling mechanism: when it evaluates an exceptional value, it rolls back the stack looking for a handler. In effect it chooses a single member of the set of exceptions to act as its representative [Hughes and O'Donnell, 1989].

6.4 Catching an imprecise exception

I describe this scheme as using "imprecise" exceptions, because we are deliberately imprecise about which exception is chosen as the representative. How, then, can we catch and handle an exception? At first we might try a non-`IO` version of `catch`:

```
bogusCatch :: a -> (Exception -> a) -> a        -- Bogus
```

`bogusCatch` evaluates its first argument; if it is an ordinary value, `bogusCatch` just returns it; if it is an exceptional value, `bogusCatch` applies the hander to the exception. But `bogusCatch` is problematic if the exceptional value contains a set of exceptions – which member of the set should be chosen? The trouble is that if the compiler decided to change evaluation order (e.g. optimisation is switched on) a different exception might be encountered, and the behaviour of the program would change.

A better approach is to separate the choice of which exception to throw from the exception-catching business:

```
evaluate :: a -> IO a
```

`evaluate x` evaluates its argument `x`; if the resulting value is an ordinary value, it just returns it. If `x` returns an exceptional value, `evaluate` chooses an arbitrary member, say `e`, from the set of exceptions, and then behaves just like `ioError e`; that is, it throws the exception `e`. Since `evaluate` is in the `IO` monad, there is no reason to suppose that it will choose the same member of the set of exceptions each time. It is free to perform input/output, so it can consult some external oracle (whether it is raining, say) to decide which member of the set to choose. Returning to our example, consider

```
catch (evaluate (throw ex1 + throw ex2)) h
```

The handler `h` will be applied to either `ex1` or `ex2`, and there is no way to tell which. It is up to `evaluate` to decide. In practice, the decision is really taken by the evaluation order chosen by the compiler when it compiles the protected expression. Notice what we have done:

- An exception can be *raised* anywhere, including in purely-functional code. This is tremendously useful. For example, pattern-match failure can now raise an exception rather than bringing execution to a halt. Similarly, Haskell 98 provies a function `error`:

  ```
  error :: String -> a
  ```

  When `error` is called, the string in printed, and execution comes to a halt. In our extended version of Haskell, `error` instead raises an exception, which gives the rest of the program a chance to recover from the failure.

- An exception can only be *caught* by `catch`, which is in the `IO` monad. This confines recovery to the monadic-I/O layer of the program, unlike ML (say) where you can catch an exception anywhere. In my view, this restriction is not irksome, and has great semantic benefits. In particular, *by confining the non-deterministic choice to the* `IO` *monad we have prevented non-determinism from infecting the entire language.*

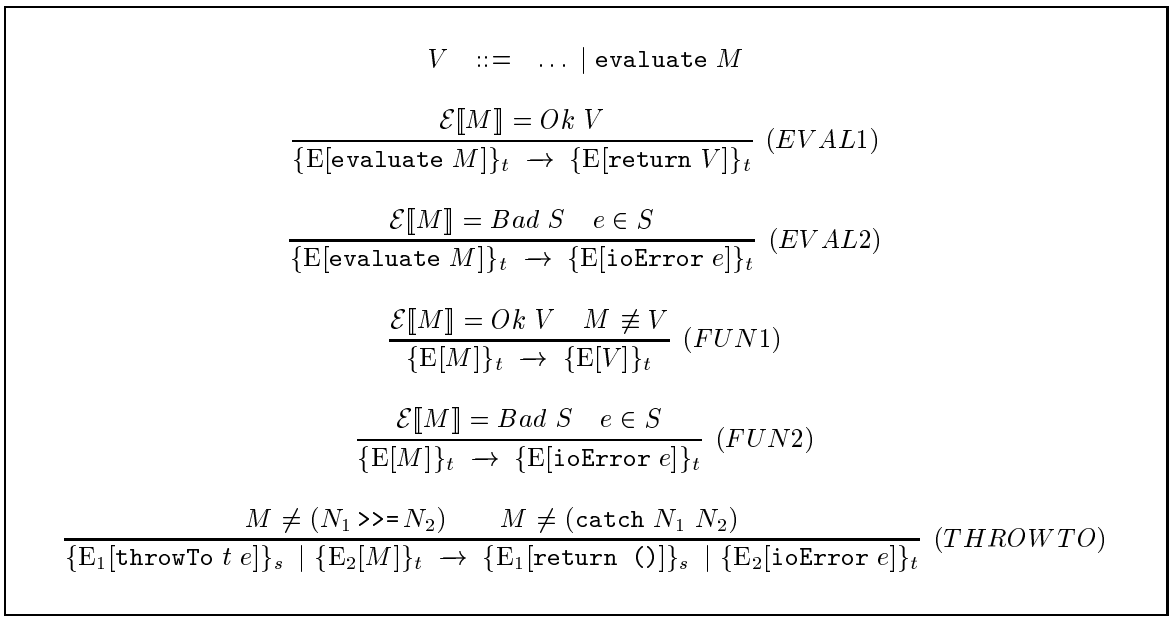

Figure 9: Further extensions for exceptions

6.4.1 Semantics of imprecise exceptions

This approach to synchronous exceptions in Haskell is described in much more detail in [Peyton Jones et al., 1999]. In particular, the paper describes how to extend a standard denotational semantics to include exceptional values, something we have not treated formally here. We will not discuss that here, for lack of space, but will content ourselves with saying that the meaning function $\mathcal{E}[\![M]\!]$ returns either $Ok\ v$ for an ordinary value v, or $Bad\ S$ for an exceptional value, where S is a non-empty set of exceptions. For example, here is the semantics of addition:

$$\mathcal{E}[\![e_1 \texttt{+} e_2]\!] = \mathcal{E}[\![e_1]\!]\ +'\ \mathcal{E}[\![e_2]\!]$$

where $+'$ an addition function defined over the semantic domain of values, thus:

$$\begin{array}{rcl} (Ok\ v_1)\ +'\ (Ok\ v_2) & = & Ok\ (v_1 + v_2) \\ (Ok\ v_1)\ +'\ (Bad\ s_2) & = & Bad\ s_2 \\ (Bad\ s_1)\ +'\ (Ok\ v_2) & = & Bad\ s_1 \\ (Bad\ s_1)\ +'\ (Bad\ s_2) & = & Bad\ (s_1 \cup s_2) \end{array}$$

The first equation deals with the normal case. The second and third deal with the case when one or other of the arguments throws an exception. The last equation handles the case when both arguments throw and exception; in this case $+'$ takes the union of the exceptions that can be thrown by the two arguments. The whole point is that $+'$ is commutative, so that $\mathcal{E}[\![e_1 \texttt{+} e_2]\!] = \mathcal{E}[\![e_2 \texttt{+} e_1]\!]$.

Given this, Figure 9 gives the extra semantics for `evaluate`. If the argument to `evaluate` is an ordinary value, `evaluate` just returns that value (EVAL1); if the value is an exceptional value, `evaluate` choooses an arbitrary member of the set of exceptions, and throws that exception using `ioError`. This deliberately-unconstrained choice is where the non-determinism shows up in the operational semantics.

Since $\mathcal{E}[\![]\!]$ has changed we must do something to rule (FUN). This is a place where our semantics forces us to recognise something we might otherwise have forgotten. Rules (FUN1) and (FUN2) replace (FUN). (FUN2) says that if the next action to perform is itself an exceptional value, then we should just propagate that as an `IO`-monad exception using `ioError`. If it is not, then we behave just like (FUN). Here is an example that shows the importance of this change:

```
catch (if (1/0) then a1 else a2)
      (\ex -> recovery_code x)
```

Before `catch` can perform the action that is its first argument, it must evaluate it; in this case, evaluating it gives divide-by-zero exception, and rule (FUN2) propagates that into an `ioError`.

The `Exception` data type is really the same as `IOError`, except that "`IOError`" does not seem an appropriate name any more. To keep things simple, we just say that `IOError` is a synonym for `Exception`:

```
type IOError = Exception
```

6.5 Asynchronous exceptions

We now turn our attention to asynchronous exceptions. For asynchronous exceptions, we add the following new primitive:

```
throwTo :: ThreadId -> Exception -> IO ()
```

This allows one thread to interrupt another. So far as the interrupted thread is concerned, the situation is just as if it abruptly called `ioError`; an exception is raised and propagated to the innermost enclosing `catch`. This is where the `ThreadId` of a forked thread becomes really useful: we can use it as a handle to send an interrupt to another thread. One thread can raise an exception in another only if it has the latter's `ThreadId`, which is returned by `forkIO`. So a thread is in danger of interruption only from its parent, unless its parent passes on its `ThreadId` to some other thread.

6.5.1 Using asynchronous exceptions

Using `throwTo` we can implement a variety of abstractions that are otherwise inaccessible. For example, we can program the combinator `parIO`, which "races" its two argument actions against each other in parallel. As soon as one terminates, it kills the other, and the overall result is the one returned by the "winner".

```
parIO :: IO a -> IO a -> IO a
```

How can we implement this? We can use an MVar to contain the overall result. We spawn two threads, that race to fill the result MVar; the first will succeed, while the second will block. The parent takes the result from the MVar, and then kills both children:

```
parIO :: IO a -> IO a -> IO a
parIO a1 a2
  = do { m <- newEmptyMVar ;
         c1 <- forkIO (child m a1) ;
         c2 <- forkIO (child m a2) ;
         r <- takeMVar m ;
         throwTo c1 Kill ;
         throwTo c2 Kill ;
    }
  where
    child m a = do { r <- a ; putMVar m r }
```

Using `parIO` we can implement a simple timeout:

```
timeout :: Int -> IO a -> IO (Maybe a)
```

The idea here is that `(timeout n a)` returns `Nothing` if `a` takes longer than `n` microseconds to complete, and `Just r` otherwise, where `r` is the value returned by `a`:

```
timeout :: Int -> IO a -> IO (Maybe a)
timeout n a = parIO (do { r <- a; return (Just r) })
                    (do { threadDelay n; return Nothing })
```

Now we might want to answer questions like this: what happens if a thread is interrupted (via a `throwTo`) while it is executing under a timeout? We can't say for sure until we give a semantics to `throwTo`, which is what we do next.

6.5.2 Semantics of asynchronous exceptions

We can express the behaviour of `throwTo` nicely in our semantics: a `throwTo` in one thread makes the target thread abandon its current action and replace it with `ioError`:

$$\frac{M \neq (N_1 \texttt{>>=} N_2) \qquad M \neq (\texttt{catch}\ N_1\ N_2)}{\{\mathrm{E}_1[\texttt{throwTo}\ t\ e]\}_s \mid \{\mathrm{E}_2[M]\}_t \ \rightarrow\ \{\mathrm{E}_1[\texttt{return}\ ()]\}_s \mid \{\mathrm{E}_2[\texttt{ioError}\ e]\}_t} \ (THROWTO)$$

The conditions above the line are essential to ensure that the context E is maximal; that is, it includes all the active `catch`es.

It should be clear that external interrupts, such as the user pressing Control-C, can also be modelled in this way. Before we can write the semantics we have to answer several questions. Does a Control-C interrupt every thread, or just a designated thread? If the latter, how does a thread get designated? These are good questions to be forced to answer, because they really do make a difference to the programmer.

Having a semantics is very helpful in answering questions like: what happens if a thread is interrupted when it is blocked waiting for an `MVar`? In the semantics, such a thread is simply stuck, with a `takeMVar` at the active site, so (THROWTO) will cause the `takeMVar` to be replaced with `ioError`. So being blocked on an `MVar` doesn't stop a thread receiving an interrupt.

Now we can say what happens to a thread that executes a sub-computation using `timeout`, but is interrupted by `throwTo` while it is waiting for the sub-computation to complete. The parent thread receives the interrupt while it is blocked on the "`takeMVar m`" inside `parIO` (Section 6.5.1); so it abandons the wait and proceeds to the innermost `catch` handler. But that means that the two threads spawned by `parIO` are not killed, and we probably want them to be. So we have to go back to fix up `parIO` somehow. In fact this turns out to be tricky to do: we have to make sure that there is no "window" in which the parent has spawned a child thread but has not set up a handler that will kill the child if the parent is interrupted.

Indeed, programming in the presence of asynchronous exceptions is notoriously difficult, so much so that Modula-3, for example, simply outlaws them. (Instead, well-behaved threads regularly poll an *alert* flag, and commit suicide if it is set [Nelson, 1991].) Haskell differs from Modula in two ways that are relevant here. First, there are fewer side effects, so there are fewer windows of vulnerability to worry about. Second, there are large parts of purely-functional code that we would like to be able to interrupt — and can indeed do so safely — but where any polling mechanism would be very undesirable. These considerations led us to define new primitive combinators to allow a thread to mask and un-mask external interrupts. This further complicates the semantics, but as a result we can write code where we have a chance of *proving* that it has no race hazards. The details are in [Marlow et al., 2000].

6.6 Summary

This section on exceptions is the most experimental of our main themes. Two papers, [Peyton Jones et al., 1999, Marlow et al., 2000], give a great deal more detail on the design, which we have introduced here only in outline.

```
    decl  ::=  foreign impexp [callconv] [string] [unsafe] varid :: primtype

  impexp  ::=  import | export
 callconv ::=  ccall | stdcall

primtype  ::=  IO primres | primres | primarg -> primtype

 primarg  ::=  extty | newarg | ForeignObj

 primres  ::=  extty | newres | ()

    extty ::=  Int | Float | Double | Char | Bool
           |   Addr | StablePtr type
           |   Int8 | Int16 | Int32 | Int64
           |   Word8 | Word16 | Word32 | Word64

  newarg  ::=  A Haskell newtype of a primarg
  newres  ::=  A Haskell newtype of a primres
```

Figure 10: The Haskell FFI proposal syntax

Adding exceptions undoubtedly complicates the language and its semantics, and that is never desirable. But they allow a qualitiative change in the robustness of a program. Now, if there is a pattern match failure almost anywhere in the code of the web server, the system can recover cleanly. Without exceptions, such a failure would be fatal.

7 Interfacing to other programs

In the programming-language world, one rule of survival is simple: dance or die. It is not enough to make a beautiful language. You must also make it easy for programs written in your beautiful language to interact with programs written in other languages. Java, C++, and C all have huge, and hugely useful, libraries available. It is fruitless to reproduce many of these libraries in Haskell; instead, we want to make it easy to call them. Similarly, if we want to plug a small Haskell program into a large project, it is necessary to enable other programs to call Haskell. It is hubristic to expect the Haskell part to always be "on top".

Haskell 98 does not specify any way to call foreign-language procedures, but there has been a lot of progress on this front in the last few years, which we survey in this section. In particular, a proposal has emerged for a Haskell language extension to support foreign-language interfacing. We will call this proposal the *the Haskell Foreign Function Interface (FFI) proposal*; it is documented at `http://haskell.org/definition/ffi`.

7.1 Calling C from Haskell, and Haskell from C

Here is how you can call a C procedure from Haskell, under the FFI proposal:

```
foreign import ccall putChar :: Char -> IO ()
```

The `foreign` declaration brings into scope a Haskell function `putChar` with the specified type. When this function is called, the effect is to call a C procedure, also called `putChar`. Simple, eh?

The complete syntax is given in Figure 10. The following points are worth noting:

- As usual, we use the `IO` monad in the result type of `putChar` to indicate that `putChar` may perform I/O, or have some other side effect. However, some foreign procedures really are foreign *functions*: they have no side effects, and their output really does depend only on their input(s). You can indicate this simply by omitting the `IO` on the result type:

  ```
  foreign import ccall toLower :: Char -> Char
  ```

  The compiler cannot check that you are telling the truth, of course!

- The keyword "`ccall`" indicates the calling convention to use; that is, which arguments are passed in which registers, which on the stack, where the result is returned, and so on. The only other currently-defined calling convention at the moment is "`stdcall`", used on Win32 platforms.

- If the foreign procedure does not have the same name as its Haskell counterpart — for example, it might start with a capital letter, which is illegal for Haskell functions — you can specify the foreign name directly:

  ```
  foreign import ccall "PutChar" putChar :: Char -> IO ()
  ```

- Foreign procedures may take several arguments. Their Haskell type is curried, as is usually the case for multi-argument Haskell functions, but on the C side the arguments are passed all at once, as is usual for C:

  ```
  foreign import ccall drawLine :: Int -> Int -> IO ()
  ```

- There is a strictly limited range of Haskell types that can be used in arguments and results, namely the "atomic" types such as `Int`, `Float`, `Double`, and so on. So how can we pass structured types, such as strings or arrays? We address this question in Section 7.3.

- An implementation of the FFI proposal must provide a collection of new atomic types (Figure 10). In particular, `Addr` is the type of uninterpreted[7] machine addresses; for example, a pointer to a `malloc`'d structure, or to a C procedure.

"`foreign import`" lets you call a C procedure from Haskell. Dually, "`foreign export`" lets you expose a Haskell function as a C procedure. For example:

```
foreign export ccall "Foo" foo :: Int -> Int
foreign export ccall       bar :: Float -> IO Float
```

These declarations are only valid if the same module defines (or imports) Haskell functions `foo` and `bar`, which have the specified types. An exported function may have an `IO` type, but it does not have to — here, `bar` does, and `foo` does not. When the module is compiled, it will expose two procedures, `Foo` and `bar`, which can be called from C.

7.2 Dynamic calls

It is quite common to make an *indirect* call to an external procedure; that is, one is supplied with the address of the procedure and one wants to call it. An example is the dynamic dispatch of a method call in an object-oriented system, indirecting through the method table of the object.

To make such an indirect call from Haskell, use the `dynamic` keyword:

[7] "Uninterpreted" in the sense that they are treated simply as bit patterns. The Haskell garbage collector does not follow the pointer.

```
foreign import dynamic ccall foo :: Addr -> Int -> IO Int
```

The first argument must be of type `Addr`, and is taken to be the address of the external procedure to be called.

There is also a dynamic version of `foreign export`:

```
foreign export dynamic mkCB :: (Int -> IO Int) -> IO Addr
```

This declaration defines a Haskell function `mkCB`. When `mkCB` is given an arbitrary Haskell function of type `(Int->IO Int)`, it returns a C function pointer (of type `Addr`) that can be called by C. Typically, this `Addr` is then somehow passed to the C program, which subsequently uses it to call the Haskell function using a C indirect call.

7.3 Marshalling

Transferring control is, in some ways, the easy bit. Transferring data "across the border" is much harder. For "atomic" types, such as `Int` and `Float`, it is clear what to do, but for structured types, matters are much murkier.

For example, suppose we wanted to import a function that operates on strings:

```
foreign import ccall uppercase :: String -> String
```

- First there is the question of data representation. One has to decide either to alter the Haskell language implementation, so that is string representation is identical to that of C, or to translate the string from one representation to another at run time. This translation is conventionally called *marshalling*.

  Since Haskell is lazy, the second approach is required. In any case, it is tremendously constraining to try to keep common representations between two languages. For example, C terminates strings with a null character, but other languages may keep a length field. Marshalling, while expensive, serves to separate the implementation concerns of the two different languages.

- Next come questions of allocation and lifetime. Where should we put the translated string? In a static piece of storage? (But how large a block should we allocate? Is it safe to re-use the same block on the next call?) Or in Haskell's heap? (But what if the called procedure does something that triggers garbage collection, and the transformed string is moved? Can the called procedure hold on to the string after it returns?) Or in C's `malloc`'d heap? (But how will it get deallocated? And `malloc` is expensive.)

- C procedures often accept pointer parameters (such as strings) that can be `NULL`. How is that to be reflected on the Haskelll side of the interface? For example, if `uppercase` did something sensible when called with a `NULL` string (e.g. returns a `NULL` string) we might like the Haskell type for `uppercase` to be

  ```
  foreign import ccall uppercase :: Maybe String -> Maybe String
  ```

  so that we can model `NULL` by `Nothing`.

The bottom line is this: there are many somewhat-arbitrary choices to make when marshalling parameters from Haskell to C and vice versa. And that's only C! There are even more choices when we consider arbitrary other languages.

What are we to do? The consensus in the Haskell community is this:

> We define a *language extension* that is as small as possible, and build *separate tools* to generate marshalling code.

The `foreign import` and `foreign export` declarations constitute the language extension. They embody just the part of foreign-language calls that cannot be done in Haskell itself, *and no more*. For example, suppose you want to import a procedure that draws a line, whose C prototype might look like this:

```
void DrawLine( float x1, float y1, float x2, float y2 )
```

One might ideally like to import this procedure with the following Haskell signature.

```
type Point = (Float,Float)
drawLine :: Point -> Point -> IO ()
```

The FFI proposal does not let you do this directly. Instead you have to do some marshalling yourself (in this case, unpacking the pairs):

```
type Point = (Float,Float)

drawLine :: Point -> Point -> IO ()
drawLine (x1,y1) (x2,y2) = dl_help x1 y1 x2 y2

foreign import ccall "DrawLine"
      dl_help :: Float -> Float -> Float -> Float -> IO ()
```

Writing all this marshalling code can get tedious, especially when one adds arrays, enumerations, in-out parameters passed by reference, NULL pointers, and so on. There are now several tools available that take some specification of the interface as input, and spit out Haskell code as output. Notably:

Green Card [Nordin et al., 1997] is a pre-processor for Haskell that reads directives embedded in a Haskell module and replaces these directives with marshalling code. Using Green Card one could write

```
type Point = (Float,Float)
drawLine :: Point -> Point -> IO ()
%call (float x1, float y1) (float x2, float y2)
%code DrawLine( x1, y1, x2, y2 )
```

Green Card is C-specific, and doesn't handle the foreign-export side of things at all.

`C->Haskell` [Chakravarty, 1999] reads both a Haskell module with special directives (or "binding hooks") and a standard C header file, and emits new Haskell module with all the marshalling code added. The advantage compared to Green Card is that less information need be specified in the binding hooks than in Green Card directives.

H/Direct [Finne et al., 1998] instead reads a description of the interface written in *Interface Definition Language* (IDL), and emits a Haskell module containing the marshalling code. IDL is a huge and hairy language, but it is neither Haskell-specific nor C-specific. H/Direct deals with both import and export, can read Java class files as well as IDL files, and can generate code to interface to C, COM, and Java.

It is well beyond the scope of these notes to give a detailed introduction to any of these tools here. However, in all cases the key point is the same: *any of these tools can be used with any Haskell compiler that implements the* `foreign` *declaration.* The very fact that there are three tools stresses the range of possible design choices, and hence the benefit of a clear separation.

7.4 Foreign objects and stable pointers

This section is not yet written.

7.5 Other languages and systems

So far we have concentrated exclusively on interfacing to programs written in C. Good progress has also been made for other languages and software architectures:

COM is Microsoft's Component Object Model, a language-independent, binary interface for composing software components. Because of its language independence COM is a very attractive target for Haskell. H/Direct directly supports both calling COM objects from Haskell, and implementing COM objects in Haskell [Peyton Jones et al., 1998, Finne et al., 1998, Finne et al., 1999, Leijen and Hook, 1999].

CORBA addresses similar goals to COM, but with a very different balance of design choices. H/Direct can read CORBA's flavour of IDL, but cannot yet generate the relevant marshalling and glue code. There is a good CORBA interface for the functional/logic language Mercury, well described in [Jeffery et al., 1999].

Lambada [Meijer and Finne, 2000] offers a collection of Haskell libraries that makes it easy to write marshalling code for calling, and being called by, Java programs. Lambada also offers a tool that reads Java class files and emits IDL that can then be fed into H/Direct to generate the marshalling code. There is ongoing work on extending the `foreign` declaration construct to suppport Java calling conventions.

7.6 Implementation notes

It is relatively easy to implement the `foreign import` declaration. The code generator needs to be taught how to generate code for a call, using appropriate calling conventions, marshalling parameters from the small, fixed range of types required by the FFI. The `dynamic` variant of `foreign import` is no harder.

A major implementation benefit is that all the I/O libraries can be built on top of such `foreign import`s; there is no need for the code generator to treat `getChar`, say, as a primitive.

Matters are not much harder for `foreign export`; here, the code generator must produce a procedure that can be called by the foreign language, again marshalling parameters appropriately. `foreign export dyanmic` is tricker, though, because we have to generate a single, static address that encapsulates a full Haskell closure. The details are given in [Finne et al., 1999]

8 Have we lost the plot?

Now that we have discussed the monadic approach in some detail, you may well be asking the following question: once we have added imperative-looking input/output, concurrency, shared variables, and exceptions, have we not simply re-invented good old procedural programming? Have we "lost the plot" — that is, forgotten what the original goals of functional programming were?

I believe not. The differences between conventional procedural programming and the monadic funtional style remain substantial:

- There is a clear distinction, enforced by the type system, between *actions* which may have side effects, and *functions* which may not. The distinction is well worth making from software engineering point of view. A function can be understood as an independent entity. It can only affect its caller through the result it returns. Whenever it is called with the same arguments it will deliver the same result. And so on.

In contrast, the interaction of an action with its caller is complex. It may read or write `MVars`, block, raise exceptions, fork new threads... and none of these things are explicit in its type.

- No reasoning laws are lost when monadic I/O is added. For example, it remains unconditionally true that

  $$\texttt{let } x \texttt{ = } e \texttt{ in } b \quad = \quad b[e/x]$$

  There are no caveats of side conditions, such as "x must not have side effects".

- In our admittedly-limited experience, most Haskell programs consist almost entirely of functions, not actions: a small monadic-I/O "skin" surrounds a large body of purely-functional code. While it is certainly possible to write a Haskell program that consists almost entirely of I/O, it is unusual do to so.

- Actions are first class values. They can be passed as arguments to functions, returned as results, stored in data structures, and so on. This gives unusual flexibility to the programmer.

Another good question is this: is the `IO` monad a sort of "sin-bin", used whenever we want to do something the breaks the purely-functional paradigm? Could we be a bit more refined about it? In particular, if we argue that it is good to know from the type of an expression that it has no side effects, would it not also be useful to express in the type some limits on the side effects it may cause? Could we have a variant of `IO` that allowed exceptions but not I/O? Or I/O but not concurrency? The answer is technically, yes of course. There is a long history of research into so-called *effect systems*, that track what kind of effects an expression can have [Jouvelot and Gifford, 1991]. Such effect systems can be expressed in a monadic way, or married with a monadic type system [Wadler, 1998]. However, the overhead on the programmer becomes greater, and I do not know of any language that uses such a system[8]. An interesting challenge remains, to devise a more refined system that is still practical; there is some promising work in this direction [Benton et al., 1998, Wadler, 1998, Tolmach, 1998, Benton and Kennedy, 1999]. Meanwhile I argue that a simple pure-or-impure distinction offers an excellent cost/benefit tradeoff.

9 Summary

We have surveyed Haskell's monadic I/O system, along with three signficant language extensions[9]. It is easy to extend a language, though! Are these extensions any good? Are they just an *ad hoc* set of responses to an *ad hoc* set of demands? Will every new demand lead to a new extension? Could the same effect be achieved with something simpler and more elegant?

I shall have to leave these judgements to you, gentle reader. These notes constitute a status report on developments in the Haskell community at the time of writing. The extensions I have described cover the needs of a huge class of applications, so I believe we have reached at least a plateau in the landscape. Everything I describe is implemented in the Glasgow Haskell compiler.

Nevertheless the resulting language is undeniably complicated, and the monadic parts have a very imperative feel. I would be delighted to find a way to make it simpler and more declarative.

[8] Some smart compilers use type-based effect systems to guide their optimisers, but that is different from the programmer-visible type system.

[9] I describe them all as "language extensions" because, while none has a significant impact on Haskell's syntax or type system, all have an impact on its semantics and implementation

Acknowledgements

These notes have been improved immeasurably by many conversations with Tony Hoare. Thank you Tony! I also want to thank Conal Elliott, Sigbjorn Finne, and Julian Seward for their constructive feedback.

References

[Ariola and Sabry, 1998] Ariola, Z. and Sabry, A. (1998). Correctness of monadic state: An imperative call-by-need calculus. In *25th ACM Symposium on Principles of Programming Languages (POPL'98)*, San Diego. ACM.

[Armstrong et al., 1996] Armstrong, J., Virding, R., , Wikstrom, C., and Williams, M. (1996). *Concurrent programming in Erlang (2nd edition)*. Prentice Hall.

[Barendsen and Smetsers, 1996] Barendsen, E. and Smetsers, S. (1996). Uniqueness typing for functional languages with graph rewriting semantics. *Mathematical Structures in Computer Science*, 6:579–612.

[Benton and Kennedy, 1999] Benton, N. and Kennedy, A. (1999). Monads, effects, and transformations. In *Higher Order Operational Techniques in Semantics: Third International Workshop*, number 26 in Electronic Notes in Theoretical Computer Science, pages 1–18. Elsevier.

[Benton et al., 1998] Benton, N., Kennedy, A., and Russell, G. (1998). Compiling Standard ML to Java bytecodes. In [ICFP98, 1998], pages 129–140.

[Bird and Wadler, 1988] Bird, R. and Wadler, P. (1988). *Introduction to Functional Programming*. Prentice Hall.

[Chakravarty, 1999] Chakravarty, M. (1999). C -> Haskell: yet another interfacing tool. In *International Workshop on Implementing Functional Languages (IFL'99)*, Lecture Notes in Computer Science, Lochem, The Netherlands. Springer Verlag.

[Felleisen and Hieb, 1992] Felleisen, M. and Hieb, R. (1992). The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 103:235–271.

[Finne et al., 1998] Finne, S., Leijen, D., Meijer, E., and Peyton Jones, S. (1998). H/Direct: a binary foreign language interface for Haskell. In [ICFP98, 1998], pages 153–162.

[Finne et al., 1999] Finne, S., Leijen, D., Meijer, E., and Peyton Jones, S. (1999). Calling hell from heaven and heaven from hell. In *ACM SIGPLAN International Conference on Functional Programming (ICFP'99)*, pages 114–125, Paris. ACM.

[Hammond and Michaelson, 1999] Hammond, K. and Michaelson, G., editors (1999). *Research Directions in Parallel Functional Programming*. Springer-Verlag.

[Hudak, 2000] Hudak, P. (2000). *The Haskell school of expression*. Cambridge University Press.

[Hughes and O'Donnell, 1989] Hughes, R. and O'Donnell, J. (1989). Expressing and reasoning about non-deterministic functional programs. In Davis, K. and Hughes, R., editors, *Functional Programming, Glasgow 1989*, pages 308–328. Workshops in Computing, Springer Verlag.

[ICFP98, 1998] ICFP98 (1998). *ACM SIGPLAN International Conference on Functional Programming (ICFP'98)*, Baltimore. ACM.

[Jeffery et al., 1999] Jeffery, D., Dowd, T., and Somogyi, Z. (1999). MCORBA: a CORBA binding for Mercury. In Gupta, editor, *Practical Applications of Declarative Languages*, number 1551 in Lecture Notes in Computer Science, pages 211–227, San Antonio, Texas. Springer Verlag.

[Jouvelot and Gifford, 1991] Jouvelot, P. and Gifford, D. (1991). Algebraic reconstruction of types and effects. In *18'th ACM Symposium on Principles of Programming Languages (POPL), Orlando*. ACM.

[Karlsen, 1998] Karlsen, E. (1998). *Tool integration in a functional programming language*. PhD thesis, University of Bremen.

[Leijen and Hook, 1999] Leijen, D. and Hook, E. M. J. (1999). Haskell as an automation controller. In *Third International School on Advanced Functional Programming (AFP'98)*, number 1608 in Lecture Notes in Computer Science, Braga, Portugal. Springer Verlag.

[Marlow et al., 2000] Marlow, S., Peyton Jones, S., and Moran, A. (2000). Asynchronous exceptions in Haskell. In *4th International Workshop on High-Level Concurrent Languages (HLCL'00)*, Montreal, Canada.

[Meijer and Finne, 2000] Meijer, E. and Finne, S. (2000). Lambada, Haskell as a better Java. Technical report, Computer Science Dept, University of Utrecht.

[Moggi, 1989] Moggi, E. (1989). Computational lambda calculus and monads. In *Logic in Computer Science, California*. IEEE.

[Nelson, 1991] Nelson, G., editor (1991). *Systems Programming with Modula-3*. Prentice Hall, Englewood Cliffs, NJ.

[Nordin et al., 1997] Nordin, T., Peyton Jones, S., and Reid, A. (1997). Green Card: a foreign-language interface for Haskell. In Launchbury, J., editor, *Haskell workshop*, Amsterdam.

[Peyton Jones et al., 1996] Peyton Jones, S., Gordon, A., and Finne, S. (1996). Concurrent Haskell. In *23rd ACM Symposium on Principles of Programming Languages (POPL'96)*, pages 295–308, St Petersburg Beach, Florida. ACM.

[Peyton Jones et al., 1998] Peyton Jones, S., Meijer, E., and Leijen, D. (1998). Scripting COM components in Haskell. In *Fifth International Conference on Software Reuse*, pages 224–233, Los Alamitos, CA. IEEE Computer Society.

[Peyton Jones et al., 1999] Peyton Jones, S., Reid, A., Hoare, C., Marlow, S., and Henderson, F. (1999). A semantics for imprecise exceptions. In *ACM Conference on Programming Languages Design and Implementation (PLDI'99)*, pages 25–36, Atlanta. ACM.

[Peyton Jones and Wadler, 1993] Peyton Jones, S. and Wadler, P. (1993). Imperative functional programming. In *20th ACM Symposium on Principles of Programming Languages (POPL'93)*, pages 71–84. ACM.

[Reppy, 1999] Reppy, J. (1999). *Concurrent programming in ML*. Cambridge University Press.

[Roscoe, 1998] Roscoe, B. (1998). *The Theory and Practice of Concurrency*. Prentice Hall.

[Thompson, 1999] Thompson, S. (1999). *Haskell: the craft of functional programming*. Addison Wesley.

[Tolmach, 1998] Tolmach, A. (1998). Optimizing ML using a hierarchy of monadic types. In *Workshop on Types in Compilation '98, Kyoto*, Lecture Notes in Computer Science, pages 97–113. Springer Verlag.

[Trinder et al., 1998] Trinder, P., Hammond, K., Loidl, H.-W., and Peyton Jones, S. (1998). Algorithm + strategy = parallelism. *Journal of Functional Programming*, 8:23–60.

[Wadler, 1992a] Wadler, P. (1992a). Comprehending monads. *Mathematical Structures in Computer Science*, 2:461–493.

[Wadler, 1992b] Wadler, P. (1992b). The essence of functional programming. In *20th ACM Symposium on Principles of Programming Languages (POPL'92)*, pages 1–14. ACM, Albuquerque.

[Wadler, 1995] Wadler, P. (1995). Monads for functional programming. In Jeuring, J. and Meijer, E., editors, *Advanced Functional Programming*, volume 925 of *Lecture Notes in Computer Science*. Springer Verlag.

[Wadler, 1997] Wadler, P. (1997). How to declare an imperative. *ACM Computing Surveys*, 29(3).

[Wadler, 1998] Wadler, P. (1998). The marriage of effects and monads. In [ICFP98, 1998], pages 63–74.

[Wright and Felleisen, 1994] Wright, A. and Felleisen, M. (1994). A syntactic approach to type soundness. *Information and Computation*, 115:38–94.