

A Recursive do for Haskell: Design and Implementation

Levent Erkök

John Launchbury

Oregon Graduate Institute of Science and Technology

Abstract

Certain programs making use of monads need to perform recursion over the values of monadic actions. Although the do-notation of Haskell provides a convenient framework for monadic programming, it lacks the generality to support such recursive bindings. To remedy this problem, we propose an extension to Haskell's do-notation and describe its translation into the core language.

1 Introduction

Suppose you are designing an interpreter for a language that has `let` expressions for introducing local bindings. Operationally, the expression `let v = e in b` denotes the same value as `b` where `e` is substituted for all free occurrences of the variable `v`. The abstract syntax of your language might include:

```
data Exp = ...
  | Let Var Exp Exp
```

Assuming the language is applicative, the natural choice for the implementation would be the environment monad. In this setting, the section of the interpreter that handles `let` expressions might look like:

```
eval (Let v e b) =
  do ev <- eval e
  inExtendedEnv (v, ev) (eval b)
```

where `inExtendedEnv` is a non-proper morphism of the environment monad extending the environment with the binding $\{v \mapsto ev\}$ before passing it on. This approach yields a very satisfactory implementation.

Notice that our `let` bindings are not recursive: The variable `v` is not known in the expression `e`. Consider what happens if we lift this restriction. All we need is a way to extend the environment with the value of `v` when we evaluate the expression `e`. That is, we want to write:

```
eval (Let v e b) =
  do ev <- inExtendedEnv (v, ev) (eval e)
  inExtendedEnv (v, ev) (eval b)
```

Unfortunately this is not valid Haskell: The variable `ev` is undefined when it is referenced in the first generator. The problem is that do-expressions of Haskell suffer from a similar problem that we are trying to solve in our toy language: Just as our tiny language can not express recursive bindings in `let`, the do-notation of Haskell can not express recursion over the values of monadic actions. Here, the function

`inExtendedEnv` performs a monadic action (that of extending the environment and passing it on to its second argument), but in doing so it depends on a value that it defines, i.e. the value of the variable `ev`.

Having failed in our naive approach, how can we implement the recursive `let`? Assuming `Val` represents the type of values that our language can produce and the following declaration for the environments:

```
data Env a = Env [(Var, Val)] -> a)
```

We can write the `Let` case of our interpreter as:

```
eval (Let v e b) =
  Env (\env ->
    let Env f = eval e
        ev    = f ((v, ev):env)
        Env g = eval b
    in g ((v, ev):env))
```

Although it works, this solution is quite annoying. First of all, we had to reveal how environments are actually implemented. Worse, our code will only work with that particular implementation: any change in the representation of environments will require change(s) in the interpreter. The code is no longer easy to understand or maintain: Almost all benefits of using a monad based implementation is lost. The failed approach that used recursive bindings in the do-notation had none of these problems.

Fortunately there is a way out. Recently, we have shown that such recursive bindings make sense in a variety of monads satisfying certain requirements [1]. In particular, a certain fixed point operator, called `mfix`, must be available for the monad in which we want to express recursion over the results of monadic actions. The research reported in our earlier work mainly concentrated on theoretical issues, such as axiomatization of the required recursive behavior, and the demonstration of satisfactory definitions for various monads. It also described a naive translation for a recursive do-notation.

The current paper, on the other hand, concentrates on the issues from a language design point of view. In this sense, we describe the translation for the new do-notation that might be employed in a real Haskell compiler. It turns out that there are several design choices involved and we explain why we prefer the proposed solutions in detail. We hope that this paper will fulfill two purposes: First of all, it will serve as a Haskell proposal for a new semantics for do-expressions. Secondly, it will provide a guideline for Haskell implementors (or thrill seekers), in case the new translation gets adopted in Haskell.

$$\begin{aligned}
\text{mfix } (\text{return} \cdot h) &= \text{return } (\text{fix } h) & (1) \\
\text{mfix } (\lambda x. a \gg= f \ x) &= a \gg= \lambda y. \text{mfix } (\lambda x. f \ x \ y) & (2) \\
\text{mfix } (\lambda^{\sim}(x, _). \text{mfix } (\lambda^{\sim}(_, y). f \ (x, y))) &= \text{mfix } f & (3)
\end{aligned}$$

Figure 1: Axioms for mfix. In axiom 2, x is not free in a .

2 Recursive monads

We first recall the definition of recursive monads:

Definition 2.1 *Recursive Monads.* A monad m is recursive if there is a function $\text{mfix} :: \forall a. (a \rightarrow m \ a) \rightarrow m \ a$ satisfying the axioms given in Figure 1.

For clarity, we will call the recursive version of the do-notation the μdo -notation, and we will write recursive do-expressions using the keyword μdo . (In actual program text, we will use `mdo`.) Whenever we refer to the do-notation, we will mean the currently available notation in Haskell that does not allow recursive bindings.

As outlined in [1], a μdo -expression can be naively translated into a do-expression as follows:

$$\begin{array}{ccc}
\mu\text{do } p_1 \leftarrow e_1 & & \text{mfix } (\lambda^{\sim} BV. \text{do } p_1 \leftarrow e_1 \\
\vdots & \Rightarrow & \vdots \\
p_n \leftarrow e_n & & p_n \leftarrow e_n \\
e & & v \leftarrow e \\
& & \text{return } BV) \\
& & \gg= \lambda \ BV. \text{return } v
\end{array}$$

where BV stands for the k -tuple consisting of all the variables occurring in all the binding patterns $p_1 \dots p_n$ plus the brand new variable v . The variables occurring in these patterns may not be multiply bound: neither in the same pattern, nor in different patterns.

In this translation, we say that a μdo -expression is well-typed if the underlying monad is recursive and the resulting code is well-typed. We will refer to this translation as the naive translation.

As an example, here is the translation for the interpreter code given in the previous section:

```

eval (Let v e b) =
  mfix (\(ev, f) ->
    do ev <- inExtendedEnv (v, ev) (eval e)
    f <- inExtendedEnv (v, ev) (eval b)
    return (ev, f))
  >>= \(ev, f) -> return f

```

This translation is well-typed as long as the underlying environment monad is declared to be recursive with a suitable definition of `mfix`. (Further details for the environment monad and the definition of `mfix` can be found in [2]).

3 A tour of μdo

In this section, we consider a number of issues and ideas that will refine our naive translation. We will also consider the case of `let`-generators in detail.

3.1 Recursive variables

The basic idea behind the naive translation is to create a function of all the variables that are bound in the μdo -expression, which will then be handed to the function `mfix`.

If there are k variables that are bound in the μdo -expression, then the function that is passed to `mfix` will have a $k + 1$ -tuple as its argument. For instance, consider the following μdo -expression, and the function created by the naive translation for it:

$$\begin{array}{ccc}
\mu\text{do } a \leftarrow f \ a & & \lambda^{\sim}(a, b, c, v). \\
b \leftarrow g \ a & \Rightarrow & \text{do } a \leftarrow f \ a \\
c \leftarrow h \ a \ b & & b \leftarrow g \ a \\
e \ c & & c \leftarrow h \ a \ b \\
& & v \leftarrow e \ c \\
& & \text{return } (a, b, c, v)
\end{array}$$

Intuitively, there is no need to recurse over the variables b and c : They are not used before being bound within the body of the do-expression. That is, the following function will work just fine:

$$\begin{array}{l}
\lambda^{\sim}(a, v). \text{do } a \leftarrow f \ a \\
\quad b \leftarrow g \ a \\
\quad c \leftarrow h \ a \ b \\
\quad v \leftarrow e \ c \\
\quad \text{return } (a, v)
\end{array}$$

(Notice that we can't leave the variable v out: its value will be projected out after the application of `mfix`.)

This observation yields the first refinement: The k -tuple BV should only contain those variables that are referenced before defined in a μdo -expression.

3.2 Segmentation

Consider the following μdo -expression:

```

μdo putStr "forming a list of 1s"
    ones ← return (1:ones)
    putStr "forming a list of 2s"
    twos ← return (2:twos)
    return (ones, twos)

```

Our translation would produce:

$$\begin{array}{l}
\text{mfix } (\lambda^{\sim}(\text{ones}, \text{twos}, v). \\
\quad \text{do putStr "forming a list of 1s"} \\
\quad \quad \text{ones} \leftarrow \text{return } (1:\text{ones}) \\
\quad \quad \text{putStr "forming a list of 2s"} \\
\quad \quad \text{twos} \leftarrow \text{return } (2:\text{twos}) \\
\quad \quad v \leftarrow \text{return } (\text{ones}, \text{twos}) \\
\quad \quad \text{return } (\text{ones}, \text{twos}, v)) \\
\gg= \lambda \ (\text{ones}, \text{twos}, v). \text{return } v
\end{array}$$

But this translation is quite unsatisfactory: The only recursion we need is in independently computing the lists `ones` and `twos`. From an intuitive point of view, recursion needs only be performed over sections of the code that actually need it. This suggests the following translation:

```

do putStr "forming a list of 1s"
  ones ← μdo ones ← return (1:ones)
        return ones
  putStr "forming a list of 2s"
  ones ← μdo twos ← return (2:twos)
        return twos
  return (ones, twos)

```

where inner μ do-expressions will further be translated accordingly. This is analogous to an optimization performed by Haskell compilers for compiling ordinary `let` expressions, where the bindings that are mutually dependent are grouped together. In the case of the `let`, this brings efficiency (no unnecessary knots need to be tied) and it enhances the polymorphic types of bound variables [4]. In our case, we increase the number of calls to `mfix`, but each `mfix` has a smaller piece of code to work on, hence, we might expect a gain in efficiency. However, there is a deeper reason why we favor this translation. There are cases when the segmentation based translation might produce values while the naive version fails to terminate. As we explained in [1], the segmentation idea corresponds to the right shrinking law, which tells us when we are allowed to shrink the scope of an μ do-expression from the bottom:

$$\begin{aligned} \text{mfix } (\lambda \tilde{(x,y)}.f \ x \gg \lambda z.g \ z \gg \lambda w.\text{return } (z,w)) \\ \sqsubseteq \text{mfix } f \gg \lambda z.g \ z \gg \lambda w.\text{return } (z,w) \end{aligned}$$

While some monads satisfy right shrinking as an equality, (identity, state, reader, output, etc.), some monads don't (maybe, lists, trees). Hence, performing segmentation will limit the scope of `mfix` calls to minimal segments, possibly improving the termination behavior. As a concrete example, consider the function:

```

g :: [Int] -> Maybe Int
g [x] = Nothing
g _   = Nothing

```

This function will return `Nothing` if its input can successfully be pattern matched against `[x]`. In particular, it will produce \perp for the input `1 : \perp` . Now consider:

```

mdo xs <- Just (1:xs)
  g xs

```

We would expect the value of this expression to be `Nothing`. Unfortunately, if we do not perform segmentation, the translation will be:

```

mfix (\~(xs, v) -> do xs <- Just (1:xs)
  v <- g xs
  return (xs, v))
>>= \~(xs, v) -> return v

```

which will evaluate to \perp . As we have discussed in [1], when the list `xs` is computed we expect to get the chain $\{\perp, 1 : \perp, 1 : 1 : \perp, \dots\}$, but applying `g` to these values before feeding them back to the list producer will produce \perp for the second element, hence short-circuiting the evaluation to \perp . If, on the other hand, we perform segmentation, we will get:

```

do xs <- mfix (\~(xs, v) -> do xs <- Just (1:xs)
  v <- return xs
  return (xs, v))
>>= \~(xs, v) -> return xs
g xs

```

which will be evaluated to `Nothing` as expected.

Hence, the segmentation idea serves two purposes. First, if we have a huge μ do-expression where only small parts of it have recursive dependencies, the recursive computation will take place only over those parts, rather than over the entire body. Secondly, and somewhat unexpectedly, monadic actions might interfere with values of bindings in unexpected ways, and segmentation will prevent such problems when the interference is not intended.

3.3 Let bindings

The `do`-notation of Haskell allows polymorphic `let` bindings. How can we accomodate them within the μ do-notation? An obvious extension suggests the following translation:

$$\begin{aligned} \mu\text{do } \dots 1 \dots \quad \text{let } p_1 = e_1 \quad \dots \quad p_n = e_n \quad \dots 2 \dots \quad e \\ \Rightarrow \quad \text{mfix } (\lambda \tilde{BV}. \text{do } \dots 1 \dots \quad \text{let } p_1 = e_1 \quad \dots \quad p_n = e_n \quad \dots 2 \dots \quad v \leftarrow e \quad \text{return } BV) \\ \gg \lambda \tilde{BV}. \text{return } v \end{aligned}$$

The translation is similar to what we had before, except now the variables bound in $p_1 \dots p_n$ appear in BV as well. The intuition is that a `let` bound variable will be visible anywhere in a μ do-expression.

Unfortunately, this idea leads to unfortunate translations. Consider the following example:

```

mdo z <- return (f 2 z)
  y <- return (f 'a' y)
  let f x y = x
  return ()

```

which will be translated to:

```

mfix (\~(z, y, f, v) ->
  do z <- return (f 2 z)
  y <- return (f 'a' y)
  let f x y = x
  v <- return ()
  return (z, y, f, v))
>>= \~(z, y, f, v) -> return v

```

The translated code will *not* type-check: The function `f`, which was used polymorphically, has become monomorphic. In fact, the situation is even worse: Referring to the schematic translation above, the `let` bound variables in patterns $p_1 \dots p_n$ will be monomorphic in the code section represented by $\dots 1 \dots$, but they will be polymorphic in the code section $\dots 2 \dots$ and in expressions $e_1 \dots e_n$. This is quite bizarre.

One solution to this problem would be to restrict `let`-bound variables to be visible only in the textually following section of the μ do-expression. This restriction, however, is not in the spirit of μ do in the first place: any defined variable should be visible in the entire body.

A second solution would be to lift `let`-generators to the top in a μ do-expression. As a concrete example, the above code can be translated to:

```

mfix (\~(z, y, v) ->
  do let f x y = x
  z <- return (f 2 z)
  y <- return (f 'a' y)
  v <- return ()
  return (z, y, v))
>>= \~(z, y, v) -> return v

```

This translation will type-check just fine. But, are we justified in moving let-generators to the top? Unfortunately not. Moving let expression to the top might change the termination behavior. Consider the expression:

```
t :: Maybe Int
t = mdo x <- f x
    let f x = Just 4
    return x
```

If we lift let-generators to top, we get:

```
t' :: Maybe Int
t' = mdo let f x = Just 4
        x <- f x
        return x
```

When run, t yields \perp , while t' computes to `Just 4`. The reason is that the introduction of f before the recursive binding provides additional information that is used in the fixpoint computation. Abstractly, moving let-generators around within a μ do-expression corresponds to the parametricity law from [1], namely:

$$\begin{aligned} \text{mfix } (\lambda x. f x \gg \text{return} \cdot h) \\ = \text{mfix } (\lambda x. \text{return } (h x) \gg f) \gg \text{return} \cdot h \end{aligned}$$

This law requires a strict h for equality. (When h is not strict, we only get an inequation.) Notice that f is not strict in our example.

Although moving let-generators to the top can be viewed as an optimization increasing termination, we refrain from doing so since we can not guarantee that we can do it all the time: Consider a situation where we use recursive bindings and moving let-generators to the top improves termination. If we ever rearrange the expression so that we cannot move the let-generators any more, (by creating a nested μ do or by just manually converting a let-generator to its equivalent `return` form), the expression may no longer terminate as often as it did before. This is not a particularly desirable situation: A perfectly valid rearrangement of the code should not fail to work just because an optimization no longer applies. For instance, consider t'' , defined as:

```
t'' :: Maybe Int
t'' = mdo x <- f x
    f <- return (\x -> Just 4)
    return x
```

Intuitively, both t and t'' should compute the same value. If our translation optimizes t to t' , then t would produce `Just 4`, but the computation of t'' will not terminate. Hence, an otherwise correct transformation might cause non-termination.

Given that we can not move let-generators around, one can think of a more complicated scheme to solve the typing problem by specifying rank-2 polymorphic types. That is, we want our very first example to translate to:

```
mfix (\(z, y, f :: forall a b. a -> b -> a, v) ->
do z <- return (f 2 z)
  y <- return (f 'a' y)
  let f x y = x
  v <- return ()
  return (z, y, f, v))
>>= \(z, y, f, v) -> return v
```

Unfortunately rank-2 polymorphism is not part of Haskell yet (although many implementations support them). Even

if Haskell had rank-2 polymorphism, we doubt that the complexity of this translation is worth it in practice.

The solution we adopt is to require let bindings to be monomorphic in a μ do-expression. That is, let becomes just a syntactic sugar within μ do, translated as:

```
let p1 = e1      =>  p1 <- return e1
...              =>  ...
pn = en         =>  pn <- return en
```

A function binding is translated similarly:

```
let f x y = x    =>  f <- return (\x y. x)
```

This gives us a uniform and simple design. If a polymorphic value definition is required, one should use the standard let expressions of Haskell, rather than the let-generator. The standard let will create its own scope with polymorphic names, as intended. For instance, our very first example should be written as:

```
mdo let f x y = x
    in mdo z <- return (f 2 z)
        y <- return (f 'a' y)
        return ()
```

which makes the intended use of f much more clear.

We expect that this restriction will be negligible in practice. Such let-generators in do-expressions are generally used for giving a name to a common pure expression in the code to follow, and such expressions are rarely polymorphic.¹ Given that there is a way to create polymorphic pure values (by using a usual let-expression), we consider that the simplicity of this design far outweighs the generality we might obtain by a much more complicated translation scheme.

3.4 Exported variables

The final refinement to the translation is about an optimization based on the observation that not all variables need to be known in an enclosing environment when a segment is formed. For instance, consider:

```
mdo x <- f x
    return 1
```

When segments are formed, we get:

```
do x <- mdo x <- f x
    return x
    return 1
```

Obviously, the variable x need not be returned from the inner μ do:

```
do mdo x <- f x
    return ()
    return 1
```

Although this is not a significant optimization, the produced code will at least be easier to compile and when printed out it will reveal more of the dependencies in the input code.

¹To see how important polymorphic let-bindings within the do-notation, we have recently polled the Haskell mailing list, the primary discussion medium for discussing Haskell-related issues on the Internet. The consensus was that such polymorphic let-generators are hardly ever used in practice and even if needed, there is always an obvious way to rewrite the expression without using them. We consider this as an indication that the monomorphism restriction is hardly an issue for let-generators. Also, a quick look at the Nofib benchmark suite reveals that polymorphism in let-generators is not an essential tool in practice.

4 Type checking μ do-expressions

In this section, we discuss how we can type check μ do-expressions. First, we fix the syntax of μ do-expressions for the purposes of this paper. We assume the non-terminal e ranges over all Haskell expressions and we extend it with the μ do expressions: (The meta-variable p ranges over patterns and `mdo` is a new keyword)

$$\begin{aligned} e &::= \mu do \\ &\quad | \dots \text{ other expressions of Haskell} \\ \mu do &::= mdo \text{ } bs \\ bs &::= \{s\} e \\ s &::= p \leftarrow e \\ &\quad | e \\ &\quad | \text{let } p = e \{ p = e \} \end{aligned}$$

where $\{x\}$ denotes one or more repetitions of x . The non-terminal s denotes statements, the building blocks of the μ do-notation. Furthermore, the variables bound in patterns within a μ do-expression may not be repeated: neither in different patterns, nor in the same pattern.

As usual, we will assume that a single expression e in a generator position in μ do is interpreted as the generator $_ \leftarrow e$. To further simplify the matters, we will also assume let-generators are turned into their `return` equivalents as described in Section 3.3. That is, the non-terminal s simply becomes:

$$s ::= p \leftarrow e$$

Once the simplifications described above are made, type checking μ do-expressions is quite straightforward. Let $V(p)$ be the set of variables that are bound by a pattern p . To type “ $\mu do \{p_i \leftarrow e_i\} e$ ” in a typing environment Γ , we first compute the set $BV = \cup_i V(p_i)$. Each $V(p_i)$ must be distinct, corresponding to the fact that variables can not be repeated in bindings. (We will discuss this in more detail in Section 6.1.) Let

$$\Gamma' = \Gamma + \{v_k : \tau_k \mid v_k \in BV\}$$

where we introduce a fresh τ_k for each bound variable v_k . Now, the typing rule simply becomes:

$$\frac{\Gamma' \vdash e_i : m \tau_i \quad \Gamma' \vdash p_i : \tau_i \quad \Gamma' \vdash e : m \tau}{\Gamma \vdash \mu do \{p_i \leftarrow e_i\} e : m \tau}$$

with the side condition that m must belong to the `MonadRec` class defined as:

```
class Monad m => MonadRec m where
  mfix :: (a -> m a) -> m a
```

5 The translation for μ do-expressions

In this section, we try to formalize the ideas and describe the translation of μ do-expressions into the core language in detail. We start by giving some definitions regarding variables and statements in μ do-expressions.

5.1 Some definitions

In these definitions, we assume that the `let`-generators are desugared into their `return` forms.

Definition 5.1 *Defined variables.* The variables defined by a generator $p \leftarrow e$ are those variables that appear in the pattern p . A sequence of bindings define variables that are defined by any statement in the sequence.

Definition 5.2 *Used variables.* The variables used by a generator $p \leftarrow e$, are those that appear free in e . A sequence of statements use those variables used by any statement in the sequence.

Definition 5.3 *Free variables.* The free variables of a sequence of statements are those that are used but not defined in the sequence.

Definition 5.4 *Recursive variables.* The recursive variables of a sequence of statements are those that are used *before* defined in that segment. Notice that being defined is an essential prerequisite: if the segment does not define that variable, it’s not recursive.

Definition 5.5 *Connected statements.* A statement s is connected to a textually following statement s' , if any of the following hold:

- s' defines a variable that is used by s ,
- s' textually appears in between s and s'' , where s is connected to s'' .

The second condition can be considered as interval closure. Notice that, unlike a usual `let`-expression, we can not reorder the bindings: Order does matter in performing side effects. Hence, if two statements are connected, then so are all the others that appear textually in between.

Definition 5.6 *Segments.* A sequence of statements within a μ do-expression is maximally connected if no statement following the sequence is connected to any statement in the sequence. Maximally connected sequences of statements are called segments. The segments of a μ do-expression can be computed as follows:

- The first statement in the body of an μ do-expression belongs to a segment.
- The furthest statement in the μ do-expression that the first statement is connected to, and all the other statements in between, form the segment. If there is no such statement, then this first statement forms a segment by itself.
- Once a segment is found, the next statement in the body of the μ do-expression (if any) starts a new segment.

Notice that the number of segments are bound by the number of statements in a μ do-expression.

Definition 5.7 *Exported variables of a segment.* The exported variables from a segment are those variables that are defined in the segment and used in any of the following segments.

5.2 Translation algorithm

We describe the algorithm step by step using the following schematic running example:

```

μdo <a b> ← <c d>      s0
    <e> ← <f>          s1
    <g> ← <h>          s2
    <f> ← <a>          s3
    <i j> ← <i e>       s4
    <j g>              s5

```

where $\langle v_1 \dots v_n \rangle$ stands for any pattern binding variables $v_1 \dots v_n$ on the left hand side of a generator and for any expression freely referencing variables $v_1 \dots v_n$ on the right hand side. Notice that the actual patterns/expressions are not important for our purposes.

Segmentation Step: Starting with the first statement, form the segments as described in Definition 5.6.

To perform this step, we will need the defined and used variables of each statement. Luckily, for our running example, these sets are obvious:

$D_0 = \{a, b\}$	$U_0 = \{c, d\}$
$D_1 = \{e\}$	$U_1 = \{f\}$
$D_2 = \{g\}$	$U_2 = \{h\}$
$D_3 = \{f\}$	$U_3 = \{a\}$
$D_4 = \{i, j\}$	$U_4 = \{i, e\}$
$D_5 = \emptyset$	$U_5 = \{j, g\}$

Applying the computation given in Definition 5.6, we get four segments: $S_0 = \{s_0\}$, $S_1 = \{s_1, s_2, s_3\}$, $S_2 = \{s_4\}$ and $S_4 = \{s_5\}$.

Analysis step: For each segment do the following: Compute recursive variables of the segment, call this set R . If R is empty, then this segment does not need fixed-point computation, leave it untouched. If R is not empty, then we will replace this segment with a single μ do expression as follows: First determine the exported variables of this segment. Let this set be E . Then create the expression

`return (v_1, \dots, v_k)`

where $v_1 \dots v_k$'s are the elements of E . (If E is empty, we'll have the expression `return ()`.) Attach this expression to the end of the segment. Mark this segment as **RECURSIVE** for future processing.

Returning to our example, here are the sets R and E for each segment, notice that we need E only when R is non-empty:

$R_0 = \emptyset$	
$R_1 = \{f\}$	$E_1 = \{e, g\}$
$R_2 = \{i\}$	$E_2 = \{j\}$
$R_3 = \emptyset$	

Since only R_1 and R_2 are non-empty, we need to add a return statement to them for their exported variables, and mark them as **RECURSIVE**. Furthermore, we add the statement `return (e, g)` to S_1 and `return j` to S_2 .

Translation step: At this step, we are left with a number of segments, some of which are marked **RECURSIVE** by the previous step. For each marked segment, do the following:

- Create a brand new variable v ,
- Modify the final expression (created by the previous step), so that its value will be bound to v ,

- Create a tuple corresponding to the R set for this segment, add v to this tuple as well. Call this tuple RT . Also create the tuple corresponding to the set E , call it ET .

- Form the expression:

```

ET ← mfix (λ~RT. do ...
           ...
           v ← return ET
           return RT)
≫ λ RT. return v

```

Notice that every segment that was marked **REC** becomes a single generator. Returning to our example, we create the following generator for segment S_1 :

```

(e, g) ← mfix (λ~(f, v). do <e> ← <f>
                          <g> ← <h>
                          <f> ← <a>
                          v ← return (e, g)
                          return (f, v))
≫ λ(f, v). return v

```

And for S_2 , we create:

```

j ← mfix (λ~(i, v). do <i, j> ← <i, e>
                     v ← return j
                     return (i, v))
≫ λ(i, v). return v

```

Finalization step: Now, concatenate all segments and form a single do-expression out of them. For our example, we obtain:

```

do <a b> ← <c d>
  (e, g) ← mfix (λ~(f, v). do <e> ← <f>
                          <g> ← <h>
                          <f> ← <a>
                          v ← return (e, g)
                          return (f, v))
                          ≫ λ(f, v). return v
  j ← mfix (λ~(i, v). do <i, j> ← <i, e>
                     v ← return j
                     return (i, v))
                     ≫ λ(i, v). return v
<j g>

```

De-sugaring step: Now, we are left with nothing but non-recursive do-expressions. We apply the well known translation algorithm (described in [4]) to get rid of the `do`. The resulting expression is now in the core language.

6 Other issues

In this section, we briefly discuss some issues related to the integration of μ do-notation in Haskell.

6.1 Unified do-notation

Although we have used a separate keyword `mdo`, we do not believe that there is any need for retaining both `do` and `mdo` as distinct constructs. In the same way that Haskell's `let` plays the role of `letrec`, we believe that the `do`-notation should be changed to capture our translation. The compiler, upon analyzing the body of the `do`, should perform

the appropriate translation depending on whether recursion is indeed used.

Of course, this brings along questions of compatibility: Will old `do`-expressions retain their meaning? Axiom 2 indeed guarantees us that this is the case [1]. There are two minor incompatibilities, however. The first is mentioned above – `let` generators become monomorphic. The second is that the current syntax for `do`-expressions in Haskell allow repeated variables in binding patterns. A new binding simply shadows the earlier one. If we allow repetitions in the μ do-notation, however, the translation would not treat them as independent. Furthermore, a repeated variable might change its type in the `do`-notation, and this will fail to type check for the μ do-notation. More importantly, one might expect that repeated variables will provide a way of constraining the values that they might take in the μ do-notation, which is not what the translation implies. Hence, even if the translation goes through, this might lead to misunderstandings.² Therefore, the syntax of the new μ do-notation explicitly prohibits a variable from being repeated in different patterns (repetition within the same pattern is disallowed following the usual Haskell convention).

Luckily, it's not a common practice in the Haskell community to repeat variable names in bindings, and in most cases the type-checker should be helpful in locating any problems. In case the translation goes through, the value of the `do`-expression will likely be “wrong enough” to alert the programmer.

6.2 A new type class

To accomodate for the recursive-monads a `MonadRec` subclass of should be added to the standard prelude:

```
class Monad m => MonadRec m where
  mfix :: (a -> m a) -> m a
```

with the understanding that the `mfix` axioms are proof obligations on the user for each `MonadRec` instance declaration. In this way, whenever a `do`-expression is used recursively, the translation will place a call to the function `mfix`, which places a `MonadRec` instance requirement on the underlying monad. The type system will then be able to warn the users appropriately in case an error is detected. The prelude can also supply the instance declarations for the `Maybe` and `List` monads [1]. The current IO extension libraries contain the `mfix` functions for the internal `ST` and `IO` monads (the functions `fixST` and `fixIO` respectively). These definitions should also be moved into `MonadRec` instance declarations for the `ST` and `IO` monads.

6.3 Generators as final expressions

Recall that the `do`-notation of Haskell always requires a final expression. We keep the same restriction for μ do-notation. However, it might be argued that we relax this requirement and allow naming the final expressions value, in case it is needed in earlier expressions. For instance, the expression:

```
mdo { x <- Just (take 5 (1:x)) }
```

²In a similar vein, it can be argued that repetitions should not have been allowed in the old `do`-notation either. List comprehensions become especially horrible: `f x = [x | x <- [x..4], x <- [x..8]]` is a confusing (yet legal) Haskell function.

would produce `Just [1,1,1,1,1]`. Although attractive, we think requiring a final expression is a better choice for two reasons. First of all, the syntax should be similar to the non-monadic version, i.e. the usual `let` expressions of Haskell. (Notice that `let` always requires a final expression.) It's also not clear what `mdo {}` means, i.e. we lack a base case. The other alternative, requiring that there should be at least one generator in a μ do, is less attractive. The second reason comes from an observation about commutative monads: In a commutative monad, it does not matter which order the bindings are performed. In a μ do, however, moving the last generator around will change the meaning of the whole expression, which is not a desirable situation.

On the other hand, if such μ do expressions are found to be very practical in the Haskell community, the syntax can be changed to allow for this possibility. Ignoring other details, our translation for this case will look like:

$$\begin{array}{ccc} \mu\text{do } p_1 \leftarrow e_1 & & \text{mfix } (\lambda BV. \text{do } p_1 \leftarrow e_1 \\ \dots & \Rightarrow & \dots \\ p_n \leftarrow e_n & & \text{v}\text{do } p_n \leftarrow e_n \\ & & \text{return } BV) \\ & & \gg= \lambda BV. \text{return } v \end{array}$$

7 Current Status

We have implemented the naive version of the translation by modifying the source code of the Hugs system.³

In this simple implementation, no type checking is performed on μ do-expressions, hence occurrences of `let` expressions are not required to be monomorphic, and the associated typing problems are ignored. Furthermore, one uses the keyword `mdo` to use the μ do-notation: The use of keyword `do` does not trigger any translation. This implementation also allows naming the final expressions value, as described in the previous section.

We have a prototype implementation of the full translation described in this paper, working on a simple subset Haskell. We plan to integrate the translation into Hugs by the end of the year.

8 An Example: Doubly linked circular lists

In this section, we give a hands-on example of using monadic recursive bindings. This example is aimed at illustrating what our current implementation is capable of expressing. (The code for this section is a literate Haskell script, extended with the μ do-notation.) We will create a doubly linked circular list which can be traversed forwards or backwards. At each node, we will store a flag indicating whether the node has been visited before. Since this flag needs to be mutable, we will use the internal `IO` monad to get access to the required reference cells. Obviously, this example can be generalized to any problem where we have interlinked stateful objects with possibly cyclic dependencies.

We first import the `MonadRec` and `IOExts` libraries. The first one declares the `MonadRec` class. The `IOExts` library provides `fixIO`, as discussed above, along with functions to create and manipulate mutable variables:

```
> import MonadRec
> import IOExts
```

³More information and downloading instructions are available online at URL: <http://www.cse.ogi.edu/PacSoft/projects/muhugs>.

The `MonadRec` declaration for IO is trivial:

```
> instance MonadRec IO where
>   mfix = fixIO
```


By this declaration, the μ do-notation becomes available for the IO monad. Each node in our list will have a mutable boolean value indicating whether it has been visited, left and right nodes and a single integer value for the data:

```
> data N = N (IORef Bool, N, Int, N)
```

To create a new node with value *i* in between the nodes *b* and *f*, we use the function `newNode`:

```
> newNode :: N -> Int -> N -> IO N
> newNode b i f = do v <- newIORef False
>   return (N (v, b, i, f))
```

Notice that the visited flag is set to `False`. We will use this function to create the following structure:



Here's the code for it:

```
> ll = mdo n0 <- newNode n3 0 n1
>   n1 <- newNode n0 1 n2
>   n2 <- newNode n1 2 n3
>   n3 <- newNode n2 3 n0
>   return n0
```

The use of μ do is essential: the cyclic nature of the construction is not expressible using an ordinary do-expression. We can test our implementation with a traversal function:

```
> data Dir = F | B deriving Eq
>
> traverse :: Dir -> N -> IO [Int]
> traverse d (N (v, b, i, f)) =
>   do visited <- readIORef v
>   if visited
>   then return []
>   else do writeIORef v True
>         let next = if d == F then f else b
>         is <- traverse d next
>         return (i:is)
```

Here's a sample run:

```
Main> ll >>= traverse F >>= print
[0,1,2,3]
Main> ll >>= traverse B >>= print
[0,3,2,1]
```

The inverse function that takes a non-empty list and constructs a doubly linked circular list out of its elements illustrates the use of μ do further:

```
> l2dll :: [Int] -> IO N
> l2dll (x:xs) = mdo c <- newNode 1 x f
>   (f, l) <- l2dll' c xs
>   return c
>
> l2dll' :: N -> [Int] -> IO (N, N)
> l2dll' p [] = return (p, p)
> l2dll' p (x:xs) = mdo c <- newNode p x f
>   (f, l) <- l2dll' c xs
>   return (c, l)
```

Note in particular the essential use of μ do in the construction of the linked list.

9 Related Work

As far as the implementation is concerned, the only directly related work we know of took place within the context of the O'Haskell programming language. O'Haskell is a concurrent, object oriented extension of Haskell designed for addressing issues in reactive functional programming [3]. One application of O'Haskell is in programming layered network protocols. Each layer interacts with its predecessor and successor by receiving and passing information in both directions. In order to connect two protocols that have mutual dependencies, one needs a recursive knot-tying operation. Since O'Haskell objects are monadic, recursive monads are employed in establishing connections between objects. To facilitate for this operation, O'Haskell extends the do-notation with a keyword `fix`, whose translation is a simplified version of ours. This extension arose from a practical need in the O'Haskell work and it was not particularly designed to meet a general need.

10 Conclusions

In this paper, we have described how to extend the do-notation of Haskell to allow for recursive bindings using the ideas given in [1]. We have started with a naive translation and refined it using various ideas to obtain a final translation strategy. It is our hope that the μ do-notation will replace the do-notation of Haskell in the future and this work will serve as a guide for Haskell implementors in integrating the new translation into their compilers and interpreters.

11 Acknowledgements

We are thankful to the members of the OGI PacSoft Research Group for valuable discussions.

The research reported in this paper is supported by the National Science Foundation (CCR-9970980).

References

- [1] ERKÖK, L., AND LAUNCHBURY, J. Recursive monadic bindings. In *Proceedings of the ACM SIGPLAN International Conference on Functional Programming (ICFP '00)* (2000, to appear.).
- [2] ERKÖK, L., AND LAUNCHBURY, J. Recursive monadic bindings: Proofs and details. Tech. rep., Department of Computer Science and Engineering, Oregon Graduate Institute, 2000 (To appear).
- [3] NORDLANDER, J. *Reactive Objects and Functional Programming*. PhD thesis, Chalmers University of Technology, Göteborg, Sweden, 1999.
- [4] PEYTON JONES, S. L., AND HUGHES, J. (Editors.) Report on the programming language Haskell 98, a non-strict purely-functional programming language. Available at: <http://www.haskell.org/onlinereport>, Feb. 1999.