# Deriving Backtracking Monad Transformers

## Functional Pearl

Ralf Hinze
Institut für Informatik III
Universität Bonn
Römerstraße 164, 53117 Bonn, Germany
ralf@informatik.uni-bonn.de

## ABSTRACT

In a paper about pretty printing J. Hughes introduced two fundamental techniques for deriving programs from their specification, where a specification consists of a signature and properties that the operations of the signature are required to satisfy. Briefly, the first technique, the term implementation, represents the operations by terms and works by defining a mapping from operations to observations — this mapping can be seen as defining a simple interpreter. The second, the context-passing implementation, represents operations as functions from their calling context to observations. We apply both techniques to derive a backtracking monad transformer that adds backtracking to an arbitrary monad. In addition to the usual backtracking operations — failure and nondeterministic choice — the prolog cut and an operation for delimiting the effect of a cut are supported.

## Categories and Subject Descriptors

D.1.1 [**Programming Techniques**]: Applicative (Functional) Programming; D.3.2 [**Programming Languages**]: Language Classifications—*applicative (functional) languages*; D.3.3 [**Programming Languages**]: Language Constructs and Features—*control structures; polymorphism*; F.3.1 [**Logics and Meanings of Programs**]: Specifying and Verifying and Reasoning about Programs—*specification techniques*; F.3.2 [**Logics and Meanings of Programs**]: Semantics of Programming Languages—*algebraic approaches to semantics*; F.3.3 [**Logics and Meanings of Programs**]: Studies of Program Constructs—*control primitives*

## General Terms

Design, languages, theory, verification

## Keywords

Program derivation, monads, monad transformers, backtracking, cut, continuations, Haskell, Prolog

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
*ICFP '00,* Montreal, Canada.
Copyright 2000 ACM 1-58113-202-6/00/0009 ..$5.00

## 1. INTRODUCTION

Why should one derive a program from its specification? Ideally, a derivation explains and motivates the various design choices taken in a particular implementation. At best a derivation eliminates the need for so-called eureka steps, which are usually inevitable if a program is explained, say, by means of example.

In a paper about pretty printing J. Hughes [6] introduced two fundamental techniques for deriving programs from their specification. Both techniques provide the programmer with considerable guidance in the process of program derivation. To illustrate their utility and versatility we apply the framework to derive several monad transformers, which among other things add backtracking to an arbitrary monad.

Briefly, a monad transformer is a mapping on monads that augments a given monad by a certain computational feature such as state, exceptions, or nondeterminism. Traditionally, monad transformers are introduced in a single big eureka step. Even the recent introductory textbook on functional programming [2] fails to explain the particular definitions of monad transformers. After defining an exception monad transformer R. Bird remarks: "Why have we chosen to write [ ... ]? The answer is: because it works.". Building upon Hughes' techniques we will try to provide a more satisfying answer. The reader should be prepared, however, that the results are somewhat different from the standard textbook examples.

The paper is organized as follows. Sec. 2 reviews monads and monad transformers. Sec. 3 introduces Hughes' techniques by means of a simple example. Sec. 4 applies the framework to derive a backtracking monad transformer that adds backtracking to an arbitrary monad. Finally, Sec. 5 extends the design of Sec. 4 to include additional control constructs: Prolog's cut and an operation for delimiting the effect of cut. Finally, Sec. 6 concludes and points out directions for future work.

## 2. PRELIMINARIES

Monads have been proposed by Moggi as a means to structure denotational semantics [11, 12]. Wadler popularized Moggi's idea in the functional programming community by using monads to structure functional programs [15, 16, 17]. In Haskell 98 [13] monads are captured by the class definition in Fig. 1. The essential idea of monads is to distinguish between *computations* and *values*. This distinction is reflected on the type level: an element of $m\ a$ represents a computation that yields a value of type $a$. The trivial computation

```
class Monad m where
  return  ::  a → m a
  (>>=)   ::  m a → (a → m b) → m b
  (>>)    ::  m a → m b → m b
  fail    ::  String → m a
  m >> n  =   m >>= const n
  fail s  =   error s
```

**Figure 1: The *Monad* class.**

that immediately returns the value $a$ is denoted *return* $a$. The operator $(\gg\!=)$, commonly called 'bind', combines two computations: $m \gg\!= k$ applies $k$ to the result of the computation $m$. The derived operation $(\gg)$ provides a handy shortcut if one is not interested in the result of the first computation. The operation *fail* is useful for signaling error conditions and will be used to this effect. Note that *fail* does not stem from the mathematical concept of a monad, but has been added to the monad class for pragmatic reasons, see [13, Sec. 3.14].

The operations are required to satisfy the following so-called *monad laws*.

$$return\ a \gg\!= k \quad = \quad k\ a \tag{M1}$$

$$m \gg\!= return \quad = \quad m \tag{M2}$$

$$(m \gg\!= k_1) \gg\!= k_2 \quad = \quad m \gg\!= (\lambda a \to k_1\ a \gg\!= k_2) \tag{M3}$$

For an explanation of the laws we refer the reader to [2, Sec. 10.3]. Note that *fail* is intentionally left unspecified.

Different monads are distinguished by the computational features they support. Each computational feature is typically accessed through a number of additional operations. For instance, a backtracking monad additionally supports the operations *false* and $(|)$ denoting failure and nondeterministic choice. It is relatively easy to construct a monad that supports only a single computational feature. Unfortunately, there is no uniform way of combining two monads, which support different computational features. The reason is simply that two features may interact in different ways. There is, however, a uniform method for augmenting a given monad by a certain computational feature. This method is captured by the following class definition which introduces *monad transformers* [9].

```
class Transformer τ where
  promote  ::  (Monad m) ⇒ m a → τ m a
  observe  ::  (Monad m) ⇒ τ m a → m a
```

A monad transformer is basically a type constructor $\tau$ that takes a monad $m$ to a monad $\tau\ m$. It must additionally provide two operations: an operation for embedding computations from the underlying monad into the transformed monad and an inverse operation, which allows us to observe 'augmented' computations in the underlying monad. Since *observe* forgets structure, it will in general be a partial function. In what follows we will abbreviate *observe* by $\omega$ and *promote* by $\pi$. Turning to the laws we require promotion to respect the monad operations.

$$\pi\ (return\ a) \quad = \quad return\ a \tag{P1}$$

$$\pi\ (m \gg\!= k) \quad = \quad \pi\ m \gg\!= (\pi \cdot k) \tag{P2}$$

These laws determine $\pi$ as a *monad morphism*. In general, $\pi$ should respect every operation the underlying monad provides in order to guarantee that a program that does not use new features behaves the same in the underlying and in the transformed monad. The counterpart of $\pi$ is not quite a monad morphism.

$$\omega\ (return\ a) \quad = \quad return\ a \tag{O1}$$

$$\omega\ (\pi\ m \gg\!= k) \quad = \quad m \gg\!= (\omega \cdot k) \tag{O2}$$

The second law is weaker than the corresponding law for $\pi$. It is unreasonable to expect more since computations in $\tau\ m$ can, in general, not be mimicked in $m$.

# 3. ADDING ABNORMAL TERMINATION

This section reviews Hughes' technique by means of a simple example. We show how to augment a given monad by an operation that allows one to terminate a computation abnormally. Monads with additional features are introduced as subclasses of *Monad*.

```
type Exception  =  String
class (Monad m) ⇒ Raise m where
  raise         ::  Exception → m a
```

The call *raise* $e$ terminates the current computation. This property is captured by the law:

$$raise\ e \gg\!= k \quad = \quad raise\ e, \tag{R1}$$

which formalizes that *raise* $e$ is a left zero of $(\gg\!=)$. Now, let us try to derive a monad transformer for this feature. Beforehand, we must determine how *raise* $e$ is observed in the base monad. We specify:

$$\omega\ (raise\ e) \quad = \quad fail\ e, \tag{O3}$$

which appears to be the only reasonable choice since we know nothing of the underlying monad.

*Remark.* We do not consider an operation for trapping exceptions (such as *handle*) in order to keep the introductory example short and simple. It is worth noting, however, that the derivation of a fully-fledged exception monad transformer proceeds similar to the derivation given in Sec. 5.

## 3.1 A term implementation

The term implementation represents operations simply by terms of the algebra and works by defining an interpreter for the language. Since we have four operations — *return*, $(\gg\!=)$, *raise*, and $\pi$ — the datatype that implements the term algebra consequently comprises four constructors. We adopt the convention that monad transformers are given names that are all in upper case. For the constructor names we re-use the names of the operations with the first letter in upper case; operators like $(\gg\!=)$ are prefixed by a colon.

```
data RAISE m a
  =  Return a
  |  ∀b. (RAISE m b) :>>= (b → RAISE m a)
  |  Raise Exception
  |  Promote (m a)
```

Note that the definition involves an existentially quantified type variable [8] in the type of $(:\gg\!=)$. We use GHC/Hugs syntax for existential quantification: the existentially quantified variable is bound by an explicit *universal* quantifier written *before* the constructor.

```
data RAISE m a                 =  Return a
                               |  ∀b.(RAISE m b) :>>= (b → RAISE m a)
                               |  Raise Exception
                               |  Promote (m a)
instance Monad (RAISE m) where
    return                     =  Return
    (>>=)                      =  (:>>=)
instance Raise (RAISE m) where
    raise                      =  Raise
instance Transformer RAISE where
    promote                    =  Promote
    observe (Return a)         =  return a
    observe (Return a :>>= k)  =  observe (k a)
    observe ((m :>>= k1) :>>= k2) = observe (m :>>= (λa → k1 a :>>= k2))
    observe (Raise e :>>= k)   =  fail e
    observe (Promote m :>>= k) =  m >>= (observe · k)
    observe (Raise e)          =  fail e
    observe (Promote m)        =  m
```

**Figure 2: A term implementation of *RAISE*.**

Now, each of the operations *return*, $(\gg\!=)$, *raise*, and $\pi$ is implemented by the corresponding constructor. In other words, the operations do nothing. All the work is performed by $\omega$ which can be seen as defining a tiny interpreter for the monadic language. Except for one case the definition of $\omega$ is straightforward.

$$\begin{array}{lcl}\omega\ (Return\ a) & = & return\ a\\ \omega\ (m :\gg\!= k) & = & \\ \omega\ (Raise\ e) & = & fail\ e\\ \omega\ (Promote\ m) & = & m\end{array}$$

Can we fill in the blank on the right-hand side? It appears impossible to define $\omega\ (m :\gg\!= k)$ in terms of its constituents. The only way out of this dilemma is to make a further case distinction on $m$:

$$\begin{array}{lcl}\omega\ (Return\ a :\gg\!= k) & = & \omega\ (k\ a)\\ \omega\ ((m :\gg\!= k_1) :\gg\!= k_2) & = & \omega\ (m :\gg\!= (\lambda a \to k_1\ a :\gg\!= k_2))\\ \omega\ (Raise\ e :\gg\!= k) & = & fail\ e\\ \omega\ (Promote\ m :\gg\!= k) & = & m \gg\!= (\omega \cdot k).\end{array}$$

Voilà. Each equation is a simple consequence of the monad laws and the laws for $\omega$. In particular, the second equation employs (M3), the associative law for $(\gg\!=)$, to reduce the size of $(:\gg\!=)$'s first argument. This rewrite step is analogous to rotating a binary tree to the right. Fig. 2 summarizes the term implementation. Note that in the sequel we will omit trivial instance declarations like *Monad* (*RAISE* $m$) and *Raise* (*RAISE* $m$).

What about correctness? First of all, the definition of $\omega$ is exhaustive. It is furthermore terminating since the size of $(:\gg\!=)$'s left argument is steadily decreasing. We can establish termination using a so-called *polynomial interpretation* of the operations [4]:

$$\begin{array}{lcl@{\qquad\qquad}lcl}Return_\tau\ a & = & 1 & Raise_\tau\ e & = & 1\\ m :\gg\!>_\tau n & = & 2 \times m + n & Promote_\tau\ m & = & 1.\end{array}$$

A multivariate polynomial $op_\tau$ of $n$ variables is associated with each $n$-ary operation $op$. For each equation $\omega\ \ell = \ldots\ \omega\ r\ \ldots$ we must show that $\tau\ \ell > \tau\ r$ for all variables (ranging over positive integers) where $\tau$ is given by $\tau(op\ e_1 \ldots e_n) = op_\tau(\tau\ e_1) \ldots (\tau\ e_n)$. Note that we consider bind only for the special case that the result of the first argument is ignored. The inclusion of $m :\gg\!= k$ in its full generality is feasible but technically more involving since the interpretation of $k$ depends on the value $m$ computes.

Does the implementation satisfy its specification? Since we are working in the free algebra, the laws do not hold: the expressions *Return* $a$ and *Return* $a :\gg\!=$ *Return*, for example, are distinct, unrelated terms. The laws of the specification only hold *under observation*. The monad laws become:

$$\begin{array}{rcl}\omega\ (return\ a \gg\!= k) & = & \omega\ (k\ a)\\ \omega\ (m \gg\!= return) & = & \omega\ m\\ \omega\ ((m \gg\!= k_1) \gg\!= k_2) & = & \omega\ (m \gg\!= (\lambda a \to k_1\ a \gg\!= k_2)).\end{array}$$

The first and the third are direct consequences of $\omega$'s definition. The second can be shown by induction on $m$. Fortunately, we can live with the weakened laws, since the only way to run computations of type *RAISE* $m$ is to use $\omega$.

## 3.2 A simplified term implementation

Can we do better than the naive term implementation? A major criticism of the first attempt is that the operations do not exploit the algebraic laws. It is conceivable that we can work with a subset of the term algebra. For instance, we need not represent both *Raise* $e$ and *Raise* $e :\gg\!=$ *Return*. A rather systematic way to determine the required subset of terms is to program a simplifier for the datatype *RAISE*, which exploits the algebraic laws as far as possible. It turns out that we only need to modify $\omega$ slightly.

$$\begin{array}{lcl}\sigma & :: & RAISE\ m\ a \to RAISE\ m\ a\\ \sigma\ (Return\ a) & = & Return\ a\\ \sigma\ (Return\ a :\gg\!= k) & = & \sigma\ (k\ a)\\ \sigma\ ((m :\gg\!= k_1) :\gg\!= k_2) & = & \sigma\ (m :\gg\!= (\lambda a \to k_1\ a :\gg\!= k_2))\\ \sigma\ (Raise\ e :\gg\!= k) & = & Raise\ e\\ \sigma\ (Promote\ m :\gg\!= k) & = & Promote\ m :\gg\!= (\sigma \cdot k)\\ \sigma\ (Raise\ e) & = & Raise\ e\end{array}$$

$$\sigma\ (Promote\ m) \quad = \quad Promote\ m$$

Inspecting the right hand sides we see that we require $(:\gg=)$ only in conjunction with *Promote*. Since $\pi\ m$ is furthermore equivalent to $\pi\ m \gg= return$ we can, in fact, restrict ourselves to the following subset of the term algebra.

$$\begin{array}{l}\textbf{data}\ RAISE\ m\ a\\ \quad = \quad Return\ a\\ \quad \mid \quad \forall b.\ PromoteBind\ (m\ b)\ (b \to RAISE\ m\ a)\\ \quad \mid \quad Raise\ Exception\end{array}$$

Following Hughes [6] we call elements of the new datatype *simplified terms*. We avoid the term normal form or canonical form since distinct terms may not necessarily be semantically different. For instance, *return a* can be represented both by *Return a* and *PromoteBind* (*return a*) *Return*. Nonetheless, using this representation the definition of $\omega$ is much simpler. It is, in fact, directly based on the laws (O1)–(O3). The complete implementation appears in Fig. 3. If we are only interested in defining a monad (not a monad transformer), then we can omit the constructor *PromoteBind*. The resulting datatype corresponds exactly to the standard definition of the exception monad.

What about efficiency? The naive implementation — or rather, the first definition of $\omega$ has a running time that is proportional to the size of the computation. Unfortunately, the 'improved' term implementation has a quadratic worst-case behaviour. Consider the expression

$$\omega\ (\cdots\ ((\pi\ (return\ 0) \gg= inc) \gg= inc)\ \cdots \gg= inc).$$

where *inc* is given by $inc\ n = \pi\ (return\ (n+1))$. Since the amortized running time of bind is proportional to the size of its first argument, it takes $O(n^2)$ steps to evaluate the expression above. The situation is analogous to flattening a binary tree. Bad luck.

## 3.3 A context-passing implementation

Since we cannot improve the implementation of the operations without sacrificing the runtime efficiency, let us try to improve the definition of $\omega$. While rewriting $\omega$ we will work out a specification for the final *context-passing implementation*. For a start, we can avoid some pattern matching if we specialize $\omega$ for $op \gg= k$. To this end we replace the equations concerning $(\gg=)$ by the single equation

$$\omega\ (op :\gg= c) \quad = \quad \omega_1\ op\ c$$

and define $\omega_1$ by

$$\begin{array}{lcl}\omega_1\ (Return\ a)\ c & = & \omega\ (c\ a)\\ \omega_1\ (m :\gg= k)\ c & = & \omega_1\ m\ (\lambda a \to k\ a :\gg= c)\\ \omega_1\ (Raise\ e)\ c & = & fail\ e\\ \omega_1\ (Promote\ m)\ c & = & m \gg= \lambda a \to \omega\ (c\ a).\end{array}$$

Interestingly, the parameter $c$ is used twice in conjunction with $\omega$. In an attempt to eliminate the mutual recursive dependence on $\omega$ we could try to pass $\omega \cdot c$ as a parameter instead of $c$. This variation of $\omega_1$, which we call $\underline{\omega}$, can be specified as follows.

$$\begin{array}{l}\underline{\omega}\ op\ \underline{c} = \omega\ (op :\gg= c)\\ \quad \Longleftarrow \quad \forall a.\ \underline{c}\ a = \omega\ (c\ a)\end{array} \qquad (1)$$

Let us derive the definition of $\underline{\omega}$ for $op = Return\ a$. We assume that precondition (1) holds — note that the equation number refers to the precondition only — and reason:

$$\begin{array}{ll}& \underline{\omega}\ (Return\ a)\ \underline{c}\\ = & \quad \{\ \text{specification and assumption (1)}\ \}\\ & \omega\ (Return\ a :\gg= c)\\ = & \quad \{\ \text{definition } \omega\ \}\\ & \omega\ (c\ a)\\ = & \quad \{\ \text{assumption (1)}\ \}\\ & \underline{c}\ a.\end{array}$$

The calculations for $Promote\ m$ and $Raise\ e$ are similar. It remains to infer the definition for $op = (m :\gg= k)$:

$$\begin{array}{ll}& \underline{\omega}\ (m :\gg= k)\ \underline{c}\\ = & \quad \{\ \text{specification and assumption (1)}\ \}\\ & \omega\ ((m :\gg= k) :\gg= c)\\ = & \quad \{\ \text{definition } \omega\ \}\\ & \omega\ (m :\gg= (\lambda a \to k\ a :\gg= c))\\ = & \quad \{\ \text{specification}\ \}\\ & \underline{\omega}\ m\ (\lambda a \to \omega\ (k\ a :\gg= c))\\ = & \quad \{\ \text{specification and assumption (1)}\ \}\\ & \underline{\omega}\ m\ (\lambda a \to \underline{\omega}\ (k\ a)\ \underline{c}).\end{array}$$

Voilà. The dependence on $\omega$ has vanished. To summarize, $\underline{\omega}$ is given by

$$\begin{array}{lcl}\underline{\omega}\ (Return\ a) & = & \lambda \underline{c} \to \underline{c}\ a\\ \underline{\omega}\ (m :\gg= k) & = & \lambda \underline{c} \to \underline{\omega}\ m\ (\lambda a \to \underline{\omega}\ (k\ a)\ \underline{c})\\ \underline{\omega}\ (Raise\ e) & = & \lambda \underline{c} \to fail\ e\\ \underline{\omega}\ (Promote\ m) & = & \lambda \underline{c} \to m \gg= \underline{c}.\end{array}$$

Note that the constructors appear only on the left-hand sides. This means that we are even able to remove the interpretative layer, ie *return a* can be implemented directly by $\lambda \underline{c} \to \underline{c}\ a$ instead of *Return*. In general, we consistently replace $\underline{\omega}\ op$ by $op$. Silently, we have converted the term implementation into a *context-passing implementation*. To see why the term 'context-passing' is appropriate, consider the final specification of the context-passing implementation.

$$\begin{array}{l}op\ \underline{c} = \omega\ (op \gg= c)\\ \quad \Longleftarrow \quad \forall a.\ \underline{c}\ a = \omega\ (c\ a)\end{array} \qquad (2)$$

The parameter $\underline{c}$ of $op$ can be seen as a representation of $op$'s calling context $\omega\ (\bullet \gg= c)$ — we represent a context by an expression that has a hole in it. This is the nub of the story: every operation knows the context in which it is called and it is furthermore able to access and to rearrange the context. This gives the implementor a much greater freedom of manoeuvre as compared to the simplified term algebra. For instance, $(\gg=)$ can use the associative law to improve efficiency. By contrast, $(\gg=)$ of the simplified term variety does not know of any outer binds and consequently falls into the efficiency trap.

It is quite instructive to infer the operations of the context-passing implementation from scratch using the specification above. Fig. 4 summarizes the calculations. Interestingly, each monad law, the law for *raise*, and each law for $\omega$ is invoked exactly once. In other words, the laws of the specification are necessary and sufficient for deriving an implementation.

It remains to determine the type of the new monad transformer. This is most easily accomplished by inspecting the

```
data RAISE m a                 =  Return a
                               |  ∀b.PromoteBind (m b) (b → RAISE m a)
                               |  Raise Exception
instance Monad (RAISE m) where
    return a                   =  Return a
    Return a >>= k             =  k a
    (PromoteBind m k1) >>= k2  =  PromoteBind m (λa → k1 a >>= k2)
    Raise e >>= k              =  Raise e
instance Raise (RAISE m) where
    raise e                    =  Raise e
instance Transformer RAISE where
    promote m                  =  PromoteBind m Return
    observe (Return a)         =  return a
    observe (PromoteBind m k)  =  m >>= (observe · k)
    observe (Raise e)          =  fail e
```

**Figure 3: A simplified term implementation of *RAISE*.**

definition of $\pi$. Note that $\pi\ m$ equals $(\gg\!=)\ m$ and recall that $(\gg\!=)$ possesses the type $\forall a.\forall b.m\ a \to (a \to m\ b) \to m\ b$ which is equivalent to $\forall a.m\ a \to (\forall b.(a \to m\ b) \to m\ b)$. Consequently, the new transformer has type $\forall b.(a \to m\ b) \to m\ b$. So, while the term implementation requires existential quantification, the context-passing implementation makes use of universal quantification. The final implementation appears in Fig. 5.[1] The cognoscenti would certainly recognize that the implementation is identical with the definition of the *continuation monad transformer* [9]. Only the types are different: *RAISE* involves rank-2 types while the continuation monad transformer is additionally parameterized with the answer type: $CONT\ ans\ m\ a = (a \to m\ ans) \to m\ ans$. The transformer $RAISE\ m$ constitutes the smallest extension of $m$ that allows one to add *raise*. Note, for instance, that *callcc* is definable in $CONT\ ans\ m$ but not in $RAISE\ m$. We will see in Sec. 4.3 that rank-2 types have advantages over parameterized types.

# 4. ADDING BACKTRACKING

By definition, a *backtracking monad* is a monad with two additional operations: the constant *false*, which denotes failure, and the binary operation $(|)$, which denotes nondeterministic choice. The class definition contains a third operation, termed *cons*, which provides a handy shortcut for $return\ a \mid m$.

```
class (Monad m) ⇒ Backtr m where
    false   ::  m a
    (|)     ::  m a → m a → m a
    cons    ::  a → m a → m a
    cons a m  =  return a | m
```

The operations are required to satisfy the following laws.

$$false \mid m = m \quad \text{(B1)}$$
$$m \mid false = m \quad \text{(B2)}$$
$$(m \mid n) \mid o = m \mid (n \mid o) \quad \text{(B3)}$$
$$false \gg\!= k = false \quad \text{(B4)}$$
$$(m \mid n) \gg\!= k = (m \gg\!= k) \mid (n \gg\!= k) \quad \text{(B5)}$$

That is, *false* and $(|)$ form a monoid; *false* is a left zero of $(\gg\!=)$, and $(\gg\!=)$ distributes leftward through $(|)$. Now, since we aim at defining a backtracking monad transformer, we must also specify the interaction of promoted operations with $(|)$:

$$(\pi\ m \gg\!= k) \mid n = \pi\ m \gg\!= \lambda a \to k\ a \mid n. \quad \text{(B6)}$$

Consider $\pi\ m$ as a deterministic computation, ie a computation that succeeds exactly once. Then (B6) formalizes our intuition that a deterministic computation can be pushed out of a disjunction's left branch. Finally, we must specify how the backtracking operations are observed in the base monad.

$$\omega\ false = fail\ \texttt{"false"} \quad \text{(O4)}$$
$$\omega\ (return\ a \mid m) = return\ a \quad \text{(O5)}$$

So we can observe the first answer of a nondeterministic computation.

## 4.1 A term implementation

The free term algebra of the backtracking monad is given by the following type definition.

```
data BACKTR m a
  =  Return a
  |  ∀b. (BACKTR m b) :>>= (b → BACKTR m a)
  |  False
  |  BACKTR m a :| BACKTR m a
  |  Promote (m a)
```

Let us try to derive an interpreter for this language. The definition of the base cases follows immediately from the specification. For $m :\gg\!= k$ we obtain:

$$\begin{aligned}
&\omega\ (Return\ a :\gg\!= k) &&= \omega\ (k\ a)\\
&\omega\ ((m :\gg\!= k_1) :\gg\!= k_2) &&= \omega\ (m :\gg\!= (\lambda a \to k_1\ a :\gg\!= k_2))\\
&\omega\ (False :\gg\!= k) &&= fail\ \texttt{"false"}\\
&\omega\ ((m :| n) :\gg\!= k) &&= \omega\ ((m :\gg\!= k) :| (n :\gg\!= k))\\
&\omega\ (Promote\ m :\gg\!= k) &&= m \gg\!= (\omega \cdot k).
\end{aligned}$$

[1] Note that *RAISE* must actually be defined using **newtype** instead of **type**. This, however, introduces an additional data constructor that affects the readability of the code. Instead we employ **type** declarations as if they worked as **newtype** declarations.

$$
\begin{array}{ll}
 & (return\ a)\ \underline{c} \\
= & \quad \{\ \text{specification and assumption (2)}\ \} \\
 & observe\ (return\ a \gg\!= c) \\
= & \quad \{\ \text{(M1)}\ \} \\
 & observe\ (c\ a) \\
= & \quad \{\ \text{assumption (2)}\ \} \\
 & \underline{c}\ a
\end{array}
$$

$$
\begin{array}{ll}
 & (m \gg\!= k)\ \underline{c} \\
= & \quad \{\ \text{specification and assumption (2)}\ \} \\
 & observe\ ((m \gg\!= k) \gg\!= c) \\
= & \quad \{\ \text{(M3)}\ \} \\
 & observe\ (m \gg\!= (\lambda a \to k\ a \gg\!= c)) \\
= & \quad \{\ \text{specification}\ \} \\
 & m\ (\lambda a \to observe\ (k\ a \gg\!= c)) \\
= & \quad \{\ \text{specification and assumption (2)}\ \} \\
 & m\ (\lambda a \to k\ a\ \underline{c})
\end{array}
$$

$$
\begin{array}{ll}
 & (raise\ e)\ \underline{c} \\
= & \quad \{\ \text{specification and assumption (2)}\ \} \\
 & observe\ (raise\ e \gg\!= c) \\
= & \quad \{\ \text{(R1)}\ \} \\
 & observe\ (raise\ e) \\
= & \quad \{\ \text{(O3)}\ \} \\
 & fail\ e
\end{array}
$$

$$
\begin{array}{ll}
 & (promote\ m)\ \underline{c} \\
= & \quad \{\ \text{specification and assumption (2)}\ \} \\
 & observe\ (promote\ m \gg\!= c) \\
= & \quad \{\ \text{(O2)}\ \} \\
 & m \gg\!= \lambda a \to observe\ (c\ a) \\
= & \quad \{\ \text{assumption (2)}\ \} \\
 & m \gg\!= \underline{c}
\end{array}
$$

$$
\begin{array}{ll}
 & observe\ m \\
= & \quad \{\ \text{(M2)}\ \} \\
 & observe\ (m \gg\!= return) \\
= & \quad \{\ \text{specification}\ \} \\
 & m\ (\lambda a \to observe\ (return\ a)) \\
= & \quad \{\ \text{(O1)}\ \} \\
 & m\ return
\end{array}
$$

**Figure 4: Deriving a context-passing implementation of *RAISE*.**

$$
\begin{array}{l}
\textbf{type}\ RAISE\ m\ a \quad = \quad \forall b.(a \to m\ b) \to m\ b \\
\textbf{instance}\ (Monad\ m) \Rightarrow Monad\ (RAISE\ m)\ \textbf{where} \\
\quad return\ a \quad = \quad \lambda \underline{c} \to \underline{c}\ a \\
\quad m \gg\!= k \quad = \quad \lambda \underline{c} \to m\ (\lambda a \to k\ a\ \underline{c}) \\
\textbf{instance}\ (Monad\ m) \Rightarrow Raise\ (RAISE\ m)\ \textbf{where} \\
\quad raise\ e \quad = \quad \lambda \underline{c} \to fail\ e \\
\textbf{instance}\ Transformer\ RAISE\ \textbf{where} \\
\quad promote\ m \quad = \quad \lambda \underline{c} \to m \gg\!= \underline{c} \\
\quad observe\ m \quad = \quad m\ return
\end{array}
$$

**Figure 5: A context-passing implementation of *RAISE*.**

Similarly, for $m :| n$ we make a case distinction on $m$:

$$
\begin{array}{lcl}
\omega\ (Return\ a :| \mathrm{f}) & = & return\ a \\
\omega\ (m :\gg\!= k :| \mathrm{f}) & = & \\
\omega\ (False :| \mathrm{f}) & = & \omega\ \mathrm{f} \\
\omega\ ((m :| n) :| \mathrm{f}) & = & \omega\ (m :| (n :| \mathrm{f})) \\
\omega\ (Promote\ m :| \mathrm{f}) & = & m.
\end{array}
$$

Unfortunately, one case remains. There is no obvious way to simplify $\omega\ (m :\gg\!= k :| \mathrm{f})$. As usual, we help ourselves by making a further case distinction on $m$.

$$
\begin{array}{lcl}
\omega\ ((Return\ a :\gg\!= k) :| \mathrm{f}) & = & \omega\ (k\ a :| \mathrm{f}) \\
\omega\ (((m :\gg\!= k_1) :\gg\!= k_2) :| \mathrm{f}) & = & \omega\ ((m :\gg\!= (\lambda a \to k_1\ a \\
 & & \qquad :\gg\!= k_2)) :| \mathrm{f}) \\
\omega\ ((False :\gg\!= k) :| \mathrm{f}) & = & \omega\ \mathrm{f} \\
\omega\ (((m :| n) :\gg\!= k) :| \mathrm{f}) & = & \omega\ ((m :\gg\!= k) \\
 & & \qquad :| ((n :\gg\!= k) :| \mathrm{f})) \\
\omega\ ((Promote\ m :\gg\!= k) :| \mathrm{f}) & = & m \gg\!= \lambda a \to \omega\ (k\ a :| \mathrm{f})
\end{array}
$$

Voilà. We have succeeded in building an interpreter for backtracking. Fig. 6 lists the complete implementation.

Now, what about correctness? Clearly, the case distinction is exhaustive. To establish termination we can use the following polynomial interpretation.

$$
\begin{array}{lcllcl}
Return_\tau\ a & = & 2 & m :|_\tau n & = & 2 \times m + n \\
m :\gg\!=_\tau n & = & m^2 \times n & Promote_\tau\ m & = & 2 \\
False_\tau & = & 2 & &
\end{array}
$$

As before, the laws of the specification only hold under observation.

## 4.2 A simplified term implementation

Let us take a brief look at the simplified term implementation. Inspecting the definition of $\omega$ — recall that a simplifier is likely to make the same case distinction as $\omega$ — we see that we need at most six terms: *False*, *Return a*, *Return a* $:|$ f, *Promote m*, *Promote m* $:\gg\!=$ *k*, and *Promote m* $:|$ f. We can eliminate three of them using $return\ a = cons\ a\ false$, $\pi\ m = \pi\ m \gg\!= return$, and $\pi\ m :| \mathrm{f} = \pi\ m \gg\!= \lambda a \to cons\ a\ \mathrm{f}$. This explains the following definition of simplified terms.

$$
\begin{array}{l}
\textbf{data}\ BACKTR\ m\ a \\
\quad = \quad False \\
\quad \mid \quad Cons\ a\ (BACKTR\ m\ a) \\
\quad \mid \quad \forall b.\ PromoteBind\ (m\ b)\ (b \to BACKTR\ m\ a)
\end{array}
$$

In essence, the simplified term algebra is an extension of the datatype of parametric lists with *False* corresponding to [] and *Cons* corresponding to (:). The additional constructor *PromoteBind* makes the difference between a monad and a monad transformer. Note that the standard list monad transformer, $LIST\ m\ a = m\ [a]$, can only be applied to so-called *commutative monads* [7]. By contrast, *BACKTR* works for arbitrary monads.

## 4.3 A context-passing implementation

In Sec. 3.3 we have seen that the context-passing implementation essentially removes the interpretative layer from the 'naive' term implementation. If we apply the same steps, we can derive very systematically a context-passing implementation of backtracking. We leave the details to the reader and sketch only the main points. First, from the case analysis $\omega$ performs we may conclude that the most complex context has the form $\omega\ (\bullet :\gg\!= c :| \mathrm{f})$. All other contexts can be rewritten into this form. Second, if we inspect the

```
data BACKTR m a          =  Return a
                         |  ∀b.(BACKTR m b) :>>= (b → BACKTR m a)
                         |  False
                         |  BACKTR m a :|: BACKTR m a
                         |  Promote (m a)
instance Transformer BACKTR where
  promote                =  Promote
  observe (Return a)     =  return a
  observe (Return a :>>= k)        =  observe (k a)
  observe ((m :>>= k1) :>>= k2)    =  observe (m :>>= (λa → k1 a :>>= k2))
  observe (False :>>= k)           =  fail "false"
  observe ((m :|: n) :>>= k)       =  observe ((m :>>= k) :|: (n :>>= k))
  observe (Promote m :>>= k)       =  m >>= (observe · k)
  observe False          =  fail "false"
  observe (Return a :|: f)         =  return a
  observe ((Return a :>>= k) :|: f)        =  observe (k a :|: f)
  observe (((m :>>= k1) :>>= k2) :|: f)    =  observe ((m :>>= (λa → k1 a :>>= k2)) :|: f)
  observe ((False :>>= k) :|: f)           =  observe f
  observe (((m :|: n) :>>= k) :|: f)       =  observe ((m :>>= k) :|: ((n :>>= k) :|: f))
  observe ((Promote m :>>= k) :|: f)       =  m >>= λa → observe (k a :|: f)
  observe (False :|: f)                    =  observe f
  observe ((m :|: n) :|: f)                =  observe (m :|: (n :|: f))
  observe (Promote m :|: f)                =  m
  observe (Promote m)                      =  m
```

**Figure 6: A term implementation of *BACKTR*.**

equations that are concerned with $\omega\ (\bullet \gg\!= c \mid f)$ we see that f appears once in the context $\omega\ \bullet$. Likewise, $c$ is used twice in the context $\omega\ (\bullet\ a \mid f)$. These observations motivate the following specification.

$$
\begin{aligned}
&op\ \underline{c}\ \underline{f} = \omega\ (op \gg\!= c \mid f) \\
&\quad\Longleftarrow\quad \underline{f} = \omega\ f \qquad (3)\\
&\quad\wedge\quad \forall f'\ \underline{f'}.\ (\forall a.\ \underline{c}\ a\ \underline{f'} = \omega\ (c\ a \mid f')) \Longleftarrow \underline{f'} = \omega\ f' \qquad (4)
\end{aligned}
$$

The nice thing about Hughes' technique is that mistakes made at this point will be discovered later when the operations are derived. For instance, it may seem unnecessary that $\underline{c}$ is parameterized with $\underline{f'}$. However, if we simply postulate $\forall a.\ \underline{c}\ a = \omega\ (c\ a \mid f)$, then we will not be able to derive a definition for $(\mid)$. Better still, one can develop the specification above while making the calculations. The derivation of *false*, for instance, motivates assumption (3); the derivation of *return* suggests either $\forall a.\ \underline{c}\ a = \omega\ (c\ a \mid f)$ or assumption (4) and the derivation of $(\mid)$ confirms that (4) is the right choice. The complete derivation appears in Fig. 7. Interestingly, each equation of the specification is invoked exactly once.

It remains to determine the type of the backtracking monad transformer. If we assume that the second parameter, the so-called *failure continuation*, has type $m\ b$, then the first parameter, the so-called *success continuation*, is of type $a \to m\ b \to m\ b$. It follows that the type of the new transformer is $\forall a.(a \to m\ b \to m\ b) \to m\ b \to m\ b$. Again, the answer type is universally quantified. We will see shortly why this is a reasonable choice. Fig. 8 summarizes the implementation.

Reconsider Fig. 7 and note that the derivation of *return*, $(\gg\!=)$, *false*, and $(\mid)$ is completely independent of $\omega$'s specification. The laws (O4) and (O5) are only required in the derivation of $\omega$. Only $\pi$ relies on (O3) which, however, appears to be the only sensible way to observe promoted operations. This suggests that we can define different observations without changing the definitions of the other operations. In other words, we may generalize the specification as follows (here $\varphi$ is an arbitrary observer function).

$$
\begin{aligned}
&op\ \underline{c}\ \underline{f} = \varphi\ (op \gg\!= c \mid f) \qquad (5)\\
&\quad\Longleftarrow\quad \underline{f} = \varphi\ f\\
&\quad\wedge\quad \forall f'\ \underline{f'}.\ (\forall a.\ \underline{c}\ a\ \underline{f'} = \varphi\ (c\ a \mid f')) \Longleftarrow \underline{f'} = \varphi\ f'\\
&\quad\wedge\quad \forall m\ k.\ \varphi\ (\pi\ m \gg\!= k) = m \gg\!= (\varphi \cdot k)
\end{aligned}
$$

To illustrate the use of the generalized specification assume that we want to collect all solutions of a nondeterministic computation. To this end we specify an observation *solve* of type $(Monad\ m) \Rightarrow BACKTR\ m\ a \to m\ [a]$:

$$
\begin{aligned}
solve\ false &= return\ [\,] && \text{(S1)}\\
solve\ (return\ a \mid m) &= a \lhd solve\ m && \text{(S2)}\\
solve\ (\pi\ m \gg\!= k) &= m \gg\!= (solve \cdot k), && \text{(S3)}
\end{aligned}
$$

where $(\lhd)$ is given by

$$
\begin{aligned}
(\lhd) &:: (Monad\ m) \Rightarrow a \to m\ [a] \to m\ [a]\\
a \lhd ms &= ms \gg\!= \lambda as \to return\ (a : as).
\end{aligned}
$$

An implementation for *solve* can be readily derived if we specialize (5) for $c = return$ and $f = false$. We obtain:

$$
\begin{aligned}
&\varphi\ op = op\ (\oplus)\ e\\
&\quad\Longleftarrow\quad \varphi\ false = e\\
&\quad\wedge\quad \forall a\ f'.\ \varphi\ (return\ a \mid f') = a \oplus \varphi\ f'\\
&\quad\wedge\quad \forall m\ k.\ \varphi\ (\pi\ m \gg\!= k) = m \gg\!= (\varphi \cdot k).
\end{aligned}
$$

$$
\begin{array}{ll}
 & (\mathit{return}\ a)\ \underline{c}\ \underline{\mathrm{f}} \\
= & \quad \{\ \text{specification and assumptions (3) \& (4)}\ \} \\
 & \mathit{observe}\ (\mathit{return}\ a \gg\!= c \mid \mathrm{f}) \\
= & \quad \{\ \text{(M1)}\ \} \\
 & \mathit{observe}\ (c\ a \mid \mathrm{f}) \\
= & \quad \{\ \text{assumptions (3) \& (4)}\ \} \\
 & \underline{c}\ a\ \underline{\mathrm{f}}
\end{array}
$$

$$
\begin{array}{ll}
 & (m \gg\!= k)\ \underline{c}\ \underline{\mathrm{f}} \\
= & \quad \{\ \text{specification and assumptions (3) \& (4)}\ \} \\
 & \mathit{observe}\ ((m \gg\!= k) \gg\!= c \mid \mathrm{f}) \\
= & \quad \{\ \text{(M3)}\ \} \\
 & \mathit{observe}\ (m \gg\!= (\lambda a \to k\ a \gg\!= c) \mid \mathrm{f}) \\
= & \quad \{\ \text{specification and assumption (3)}\ \} \\
 & m\ (\lambda a\ \underline{\mathrm{f}'} \to \mathit{observe}\ (k\ a \gg\!= c \mid \mathrm{f}'))\ \underline{\mathrm{f}} \\
= & \quad \{\ \text{specification and assumption (4)}\ \} \\
 & m\ (\lambda a\ \underline{\mathrm{f}'} \to k\ a\ \underline{c}\ \underline{\mathrm{f}'})\ \underline{\mathrm{f}}
\end{array}
$$

$$
\begin{array}{ll}
 & \mathit{false}\ \underline{c}\ \underline{\mathrm{f}} \\
= & \quad \{\ \text{specification and assumptions (3) \& (4)}\ \} \\
 & \mathit{observe}\ (\mathit{false} \gg\!= c \mid \mathrm{f}) \\
= & \quad \{\ \text{(B4)}\ \} \\
 & \mathit{observe}\ (\mathit{false} \mid \mathrm{f}) \\
= & \quad \{\ \text{(B1)}\ \} \\
 & \mathit{observe}\ \mathrm{f} \\
= & \quad \{\ \text{assumption (3)}\ \} \\
 & \underline{\mathrm{f}}
\end{array}
$$

$$
\begin{array}{ll}
 & (m \mid n)\ \underline{c}\ \underline{\mathrm{f}} \\
= & \quad \{\ \text{specification and assumptions (3) \& (4)}\ \} \\
 & \mathit{observe}\ ((m \mid n) \gg\!= c \mid \mathrm{f}) \\
= & \quad \{\ \text{(B5)}\ \} \\
 & \mathit{observe}\ ((m \gg\!= c \mid n \gg\!= c) \mid \mathrm{f}) \\
= & \quad \{\ \text{(B3)}\ \} \\
 & \mathit{observe}\ (m \gg\!= c \mid (n \gg\!= c \mid \mathrm{f})) \\
= & \quad \{\ \text{specification and assumption (4)}\ \} \\
 & m\ \underline{c}\ (\mathit{observe}\ (n \gg\!= c \mid \mathrm{f})) \\
= & \quad \{\ \text{specification and assumptions (3) \& (4)}\ \} \\
 & m\ \underline{c}\ (n\ \underline{c}\ \underline{\mathrm{f}})
\end{array}
$$

$$
\begin{array}{ll}
 & (\mathit{promote}\ m)\ \underline{c}\ \underline{\mathrm{f}} \\
= & \quad \{\ \text{specification and assumptions (3) \& (4)}\ \} \\
 & \mathit{observe}\ (\mathit{promote}\ m \gg\!= c \mid \mathrm{f}) \\
= & \quad \{\ \text{(B6)}\ \} \\
 & \mathit{observe}\ (\mathit{promote}\ m \gg\!= (\lambda a \to c\ a \mid \mathrm{f})) \\
= & \quad \{\ \text{(O3)}\ \} \\
 & m \gg\!= \lambda a \to \mathit{observe}\ (c\ a \mid \mathrm{f}) \\
= & \quad \{\ \text{assumptions (3) \& (4)}\ \} \\
 & m \gg\!= \lambda a \to \underline{c}\ a\ \underline{\mathrm{f}}
\end{array}
$$

$$
\begin{array}{ll}
 & \mathit{observe}\ m \\
= & \quad \{\ \text{(M2) and (B2)}\ \} \\
 & \mathit{observe}\ (m \gg\!= \mathit{return} \mid \mathit{false}) \\
= & \quad \{\ \text{specification}\ \} \\
 & m\ (\lambda a\ \underline{\mathrm{f}'} \to \mathit{observe}\ (\mathit{return}\ a \mid \mathrm{f}'))\ (\mathit{observe}\ \mathit{false}) \\
= & \quad \{\ \text{(O4)}\ \} \\
 & m\ (\lambda a\ \underline{\mathrm{f}'} \to \mathit{observe}\ (\mathit{return}\ a \mid \mathrm{f}'))\ (\mathit{fail}\ \texttt{"false"}) \\
= & \quad \{\ \text{(O5)}\ \} \\
 & m\ (\lambda a\ \underline{\mathrm{f}'} \to \mathit{return}\ a)\ (\mathit{fail}\ \texttt{"false"})
\end{array}
$$

**Figure 7: Deriving a context-passing implementation of *BACKTR*.**

$$
\begin{array}{lcl}
\multicolumn{3}{l}{\textbf{type}\ \mathit{BACKTR}\ m\ a} \\
 & = & \forall b.(a \to m\ b \to m\ b) \to m\ b \to m\ b \\
\multicolumn{3}{l}{\textbf{instance}\ (\mathit{Monad}\ m) \Rightarrow \mathit{Monad}\ (\mathit{BACKTR}\ m)\ \textbf{where}} \\
\quad \mathit{return}\ a & = & \lambda \underline{c} \to \underline{c}\ a \\
\quad m \gg\!= k & = & \lambda \underline{c} \to m\ (\lambda a \to k\ a\ \underline{c}) \\
\multicolumn{3}{l}{\textbf{instance}\ (\mathit{Monad}\ m) \Rightarrow \mathit{Backtr}\ (\mathit{BACKTR}\ m)\ \textbf{where}} \\
\quad \mathit{false} & = & \lambda \underline{c} \to \mathit{id} \\
\quad m \mid n & = & \lambda \underline{c} \to m\ \underline{c} \cdot n\ \underline{c} \\
\multicolumn{3}{l}{\textbf{instance}\ \mathit{Transformer}\ \mathit{BACKTR}\ \textbf{where}} \\
\quad \mathit{promote}\ m & = & \lambda \underline{c}\ \underline{\mathrm{f}} \to m \gg\!= \lambda a \to \underline{c}\ a\ \underline{\mathrm{f}} \\
\quad \mathit{observe}\ m & = & m\ (\lambda a\ \underline{\mathrm{f}} \to \mathit{return}\ a)\ (\mathit{fail}\ \texttt{"false"})
\end{array}
$$

**Figure 8: A context-passing implementation of *BACKTR*.**

Consequently, $\mathit{solve}\ \mathit{op} = \mathit{op}\ (\lhd)\ (\mathit{return}\ [\,])$. Now, instead of providing *solve* as an additional observer function we promote it into the backtracking monad.

$$
\begin{array}{l}
\mathit{sols} :: (\mathit{Monad}\ m) \Rightarrow \mathit{BACKTR}\ m\ a \to \mathit{BACKTR}\ m\ [a] \\
\mathit{sols}\ m \quad = \quad \pi\ (m\ (\lhd)\ (\mathit{return}\ [\,]))
\end{array}
$$

This way we can use the all solution collecting function as if it were a new computational primitive. Since $\pi$ is a monad morphism, we furthermore know that *sols* satisfies suitable variants of (S1)–(S3). Note that the implementation of *sols* makes non-trivial use of rank-2 types. If we used a variant of *BACKTR* that is parameterized with the answer type, then *sols* cannot be assigned a type $t\ a \to t\ [a]$ for some $t$.

## 5. ADDING CONTROL

Let us extend our language by two additional Prolog-like control constructs. The first, called cut and denoted '!', allows us to reduce the search space by dynamically pruning unwanted computation paths. The second, termed *call*, is provided for controlling the effect of cut. Both constructs are introduced as a subclass of *Backtr*.

$$
\begin{array}{lcl}
\multicolumn{3}{l}{\textbf{class}\ (\mathit{Backtr}\ m) \Rightarrow \mathit{Cut}\ m\ \textbf{where}} \\
\quad ! & :: & m\ () \\
\quad \mathit{cutfalse} & :: & m\ a \\
\quad \mathit{call} & :: & m\ a \to m\ a \\
\quad ! & = & \mathit{return}\ () \mid \mathit{cutfalse} \\
\quad \mathit{cutfalse} & = & ! \gg \mathit{false}
\end{array}
$$

The operational reading of '!' and *call* is as follows. The cut succeeds exactly once and returns (). As a side-effect it discards *all* previous alternatives. The operation *call* delimits the effect of cut: $\mathit{call}\ m$ executes $m$; if the cut is invoked in $m$, it discards only the choices made since $m$ was called. The class definition contains a third operation, called *cutfalse*, which captures a common programming idiom in Prolog, the so-called cut-fail combination [14].

Note that instances of the class *Cut* must define either '!' or *cutfalse*. The default definitions already employ our knowledge about the properties of the operations, which we shall consider next. We sketch the axiomatization only briefly, for a more in-depth treatment the interested reader is referred to [5]. The cut is characterized by the following

three equations.

$$(! \gg m) \mid n = ! \gg m \quad (!1)$$
$$! \gg (m \mid n) = m \mid ! \gg n \quad (!2)$$
$$! \gg return\ () = ! \quad (!3)$$

The first equation formalizes our intuition that a cut discards *past* choice points, ie alternatives which appear 'above' or to its left. On the other hand, the cut does not affect *future* choice points, ie alternatives which appear to its right. This fact is captured by (!2). Axiom (!3) simply records that cut returns (). An immediate consequence of the axioms is $! = return\ () \mid ! \gg false$, which explains the default definition of cut. To see why this relation holds replace $m$ by *return* () and $n$ by *false* in (!2).

The operation *cutfalse* enjoys algebraic properties which are somewhat easier to remember: *cutfalse* is a left zero of both ($\gg\!=$) and ($\mid$).

$$cutfalse \gg\!= k = cutfalse \quad (CF1)$$
$$cutfalse \mid m = cutfalse \quad (CF2)$$

The default definitions use the fact that '!' and *cutfalse* are interdefinable. Likewise, the two sets of axioms are interchangeable. We may either define $cutfalse = ! \gg false$ and take the equations for '!' as axioms — the laws for *cutfalse* are then simple logical consequences — or vice versa.

Finally, *call* is required to satisfy:

$$call\ false = false \quad (C1)$$
$$call\ (return\ a \mid m) = return\ a \mid call\ m \quad (C2)$$
$$call\ (! \gg m) = call\ m \quad (C3)$$
$$call\ (m \mid cutfalse) = call\ m \quad (C4)$$
$$call\ (\pi\ m \gg\!= k) = \pi\ m \gg\!= (call \cdot k). \quad (C5)$$

Thus, *call* $m$ behaves essentially like $m$ except that any cut inside $m$ has only local effect. It remains to lay down how the new operations are observed in the underlying monad.

$$\omega\ (call\ m) = \omega\ m \quad (O6)$$

Note that we need not specify the observation of '!' and *cutfalse* since (C3), (C4), and (O6) imply $\omega\ (! \gg m) = \omega\ m$ and $\omega\ (m \mid cutfalse) = \omega\ m$.

## 5.1 A term implementation

The free term implementation faces two problems, one technical and one fundamental. Let us consider the technical problem first. Inspecting the type signature of cut, we find that cut cannot be turned into a constructor, because it does not have the right type. If we define a type, say, *CUT m a*, then '!' must have exactly this type. Alas, its type signature only allows for a substitution instance, ie *CUT m* (). Here, we stumble over the general problem that Haskell's **data** construct is not capable of expressing arbitrary polymorphic term algebras. Fortunately, the axioms save the day. Since '!' can be expressed in terms of *cutfalse* and this operation has a polymorphic type, we turn *cutfalse* into a constructor.

```
data CUT m a = Return a
             | ∀b. (CUT m b) :>>= (b → CUT m a)
             | False
             | CutFalse
             | CUT m a :| CUT m a
             | Call (CUT m a)
             | Promote (m a)
```

Turning to the definition of $\omega$ we encounter a problem of a more fundamental nature. For a start, we discover that the term $\omega\ (call\ m \gg\!= k)$ cannot be simplified. If we make a further case distinction on $m$, we end up with $\omega\ (call\ (call\ m \gg\!= k_1) \gg\!= k_2)$ which is not reducible either. The crux is that we have no axiom that specifies the interaction of *call* with ($\gg\!=$). And rightly so. Each *call* opens a new scope for cut. Hence, we cannot reasonably expect that nested *call*s can be collapsed. This suggests to define two interpreters, one for $\omega$ and one for *call*, which means, of course, that the implementation is no longer based on the free term algebra. The resulting code, which is mostly straightforward, appears in Fig. 9. The equations involving *cutfalse* use the fact that *cutfalse* is a left zero of both ($\gg\!=$) and ($\mid$), and that *call* maps *cutfalse* to *false*. Note that $\omega$ falls back on *call* to avoid duplication of code.

## 5.2 A simplified term implementation

For the sake of completeness, here is the simplified term algebra, which augments the type *BACKTR* of Sec. 4.2 with an additional constructor for *cutfalse*.

```
data CUT m a
  = False
  | CutFalse
  | Cons a (CUT m a)
  | ∀b. PromoteBind (m b) (b → CUT m a)
```

In essence, we have lists with two different terminators, *False* and *CutFalse*. Interestingly, exactly this structure (without *PromoteBind*) has been used to give a denotational semantics for Prolog with cut [1], where *cutfalse* and *call* are termed *esc* and *unesc*.

## 5.3 A context-passing implementation

We have seen that the realization of cut and *call* is more demanding since there is no way to simplify nested invocations of *call*. With regard to the context-passing implementation this means that we must consider an infinite number of possible contexts. Using a grammar-like notation we can characterize the set of all possible contexts as follows.

$$\mathcal{C} ::= \omega\ (\bullet \gg\!= k \mid \mathrm{f}) \mid \mathcal{C}[call\ (\bullet \gg\!= k \mid \mathrm{f})]$$

A context is either simple or of the form $C[call\ (\bullet \gg\!= k \mid \mathrm{f})]$ where $C$ is the enclosing context. Thus, contexts are organized in a list- or stack-like fashion. As usual we will represent operations as functions from contexts to observations. The main difference to Sec. 4.3 is that each operation must now consider two different contexts and that the contexts are recursively defined. Note, however, the duality between the term and the context-passing implementation: In Sec. 5.1 we had two interpreters, *call* and $\omega$, and each interpreter had to consider each operation. Here we have two contexts and each operation must consider each context.

Turning to the implementation details we will see that the greatest difficulty is to get the types right. The contexts are represented by a recursive datatype with two constructors: *OBCC* (which is an acronym for observe-bind-choice context) and *CBCC* (call-bind-choice context). The first takes two arguments, the success and the failure continuation, while the second expects three arguments, the two continuations and the representation of the enclosing context. In order to infer their types it is useful to consider the

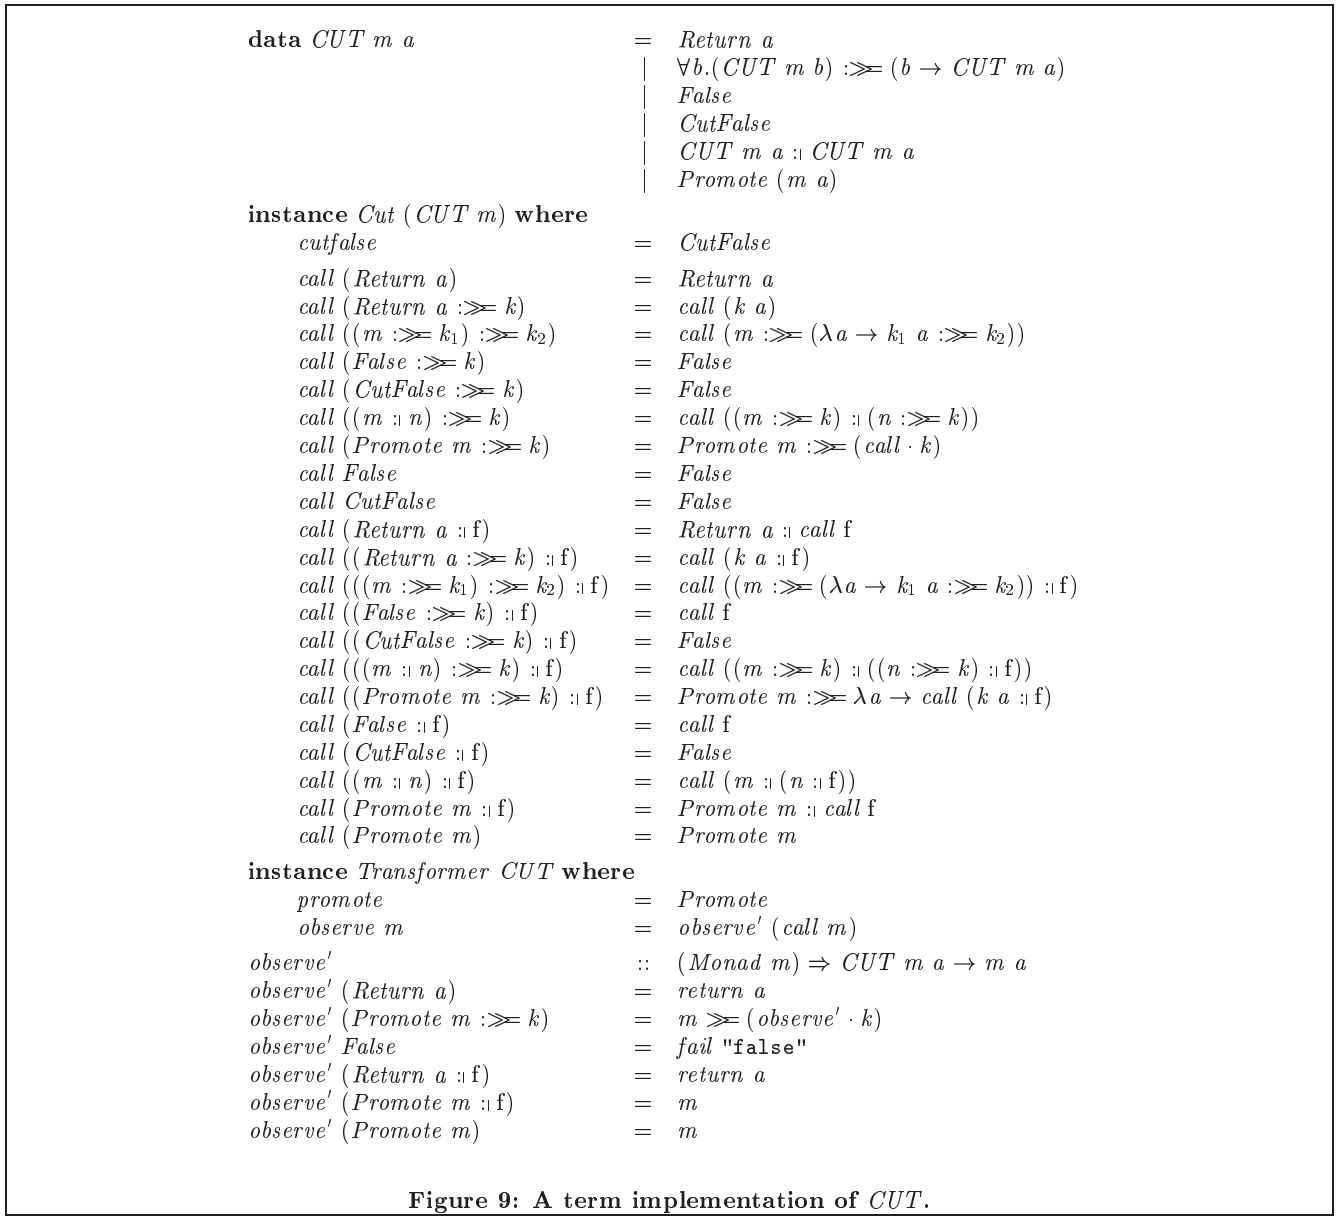


**Figure 9: A term implementation of $CUT$.**

specification of the context-passing implementation beforehand. The specification is similar to the one given in Sec. 4.3 except that we have two clauses, one for each context.

$$
\begin{aligned}
& op\ (OBCC\ \underline{c}\ \underline{f}) = \omega\ (op \gg\!= c \mid f) \\
& \quad \Longleftarrow \quad \underline{f} = \omega\ f && (4) \\
& \quad \wedge \quad \forall f'\ \underline{f'}.\ (\forall a.\ \underline{c}\ a\ \underline{f'} = \omega\ (c\ a \mid f')) \Longleftarrow \underline{f'} = \omega\ f' && (5) \\
& op \cdot CBCC\ \underline{c}\ \underline{f} = call\ (op \gg\!= c \mid f) \\
& \quad \Longleftarrow \quad \underline{f} = call\ f && (6) \\
& \quad \wedge \forall f'\ \underline{f'}.\ (\forall a.\ \underline{c}\ a\ \underline{f'} = call\ (c\ a \mid f')) \Longleftarrow \underline{f'} = call\ f' && (7)
\end{aligned}
$$

The first clause closely corresponds to the specification of Sec. 4.3. For that reason we may assign the components of $OBCC\ \underline{c}\ \underline{f}$ the same types: $\underline{f}$ has type $m\ b$ and $\underline{c}$ has type $a \to m\ b \to m\ b$ where $b$ is the answer type. This implies that the type of contexts must be parameterized with $m$, $a$, and $b$.

$$\mathbf{data}\ \mathcal{C}\ m\ a\ b \quad = \quad OBCC\ (a \to m\ b \to m\ b)\ (m\ b) \mid \ldots$$

The second clause of the specification has essentially the same structure as the first one. The main difference is that the components dwell in the transformed monad rather than in the underlying monad. Furthermore, *CBCC* additionally contains the enclosing context which may have a different type. To illustrate, consider the context $C[\,call\ (\bullet \gg\!= c \mid f)\,]$ of type $\mathcal{C}\ m\ a\ b$. If we assume that the enclosing context $C$ has type $\mathcal{C}\ m\ i\ b$ — there is no reason to require that $C$ has the same argument type as the entire context, but it must have the same answer type — then f has type $CUT\ m\ i$ and $c$ has type $a \to CUT\ m\ i \to CUT\ m\ i$. This motivates the following definition.

$$
\begin{aligned}
\mathbf{data}\ \mathcal{C}\ m\ a\ b \quad = \quad & OBCC\ (a \to m\ b \to m\ b)\ (m\ b) \\
\mid \quad & \forall i.\, CBCC\ (a \to CUT\ m\ i \to CUT\ m\ i) \\
& \qquad (CUT\ m\ i)\ (\mathcal{C}\ m\ i\ b)
\end{aligned}
$$

$$\mathbf{type}\ CUT\ m\ a = \forall b.\mathcal{C}\ m\ a\ b \to m\ b$$

Note that the intermediate type is represented by an existentially quantified variable. The mutually recursive types $\mathcal{C}$ and $CUT$ are somewhat mind-boggling as they involve both universal and existential quantification, a combination of features the author has not seen before.

Now that we have the types right, we can address the derivation of the various operations. Except for $\pi$ the calculations are analogous to those of Sec. 4.3. For $\pi\ m$ we must conduct an inductive proof to show that $m$ propagates through the stack of contexts, ie $(\pi\ m \gg\!= k)\ c = m \gg\!= \lambda a \to k\ a\ c$. The proof is left as an exercise to the reader. To derive cut we reason:

$$
\begin{aligned}
& ! \cdot CBCC\ \underline{c}\ \underline{f} \\
= \quad & \{ \text{ specification and assumptions (6) \& (7) } \} \\
& call\ (! \gg\!= c \mid f) \\
= \quad & \{ \text{ (!3), (M3), and (M1) } \} \\
& call\ (! \gg c\ () \mid f) \\
= \quad & \{ \text{ (!1) and (!2) } \} \\
& call\ (c\ () \mid ! \gg false) \\
= \quad & \{ \text{ assumption (7) } \} \\
& \underline{c}\ ()\ (call\ (! \gg false)) \\
= \quad & \{ \text{ (C3) and (C1) } \} \\
& \underline{c}\ ()\ false.
\end{aligned}
$$

The derivation for the context *OBCC* proceeds in an analogous fashion. For *call* we obtain:

$$
\begin{aligned}
& call\ m \\
= \quad & \{ \text{ (M2) and (B2) } \} \\
& call\ (m \gg\!= return \mid false) \\
= \quad & \{ \text{ specification } \} \\
& m \cdot CBCC\ (\lambda a\ \underline{f'} \to call\ (return\ a \mid f'))\ (call\ false) \\
= \quad & \{ \text{ (C1) and (C2) } \} \\
& m \cdot CBCC\ (\lambda a\ \underline{f'} \to return\ a \mid call\ f')\ false \\
= \quad & \{ \underline{f'} = call\ f' \} \\
& m \cdot CBCC\ (\lambda a\ \underline{f'} \to return\ a \mid \underline{f'})\ false \\
= \quad & \{ \text{ definition } cons \} \\
& m \cdot CBCC\ cons\ false.
\end{aligned}
$$

Thus, *call* installs a new context with *cons* and *fail* as the initial failure continuations. The complete implementation appears in Fig. 10. Note that most of the monad operations *pattern match on the context*. This fact sets the implementation apart from *continuation passing style* (CPS), where the context is an anonymous function that cannot be inspected. By contrast, CPS-based implementations [3, 10] use three continuations (a success, a failure, and a cut continuation).

# 6. CONCLUSION

Naturally, most of the credit goes to J. Hughes for introducing two wonderful techniques for deriving programs from their specification. Many of the calculations given in this paper already appear in [6], albeit specialized to monads. However, the step from monads to monad transformers is not a big one and this is one of the pleasant findings. To be able to derive an implementation of Prolog's control core from a given axiomatization is quite remarkable. We have furthermore applied the techniques to derive state monad transformers, *STATE*, and exception monad transformers, *EXC*. In both cases the techniques worked well.

Some work remains to be done though. We did not address the problem of promotion in general. It is well known that different combinations of transformers generally lead to different semantics of the operations involved. For instance, composing *STATE* with *BACKTR* yields a backtracking monad with a backtrackable state, which is characterized as follows.

$$
\begin{aligned}
store\ s \gg false \quad &= \quad false \\
store\ s \gg (m \mid n) \quad &= \quad store\ s \gg m \mid store\ s \gg n
\end{aligned}
$$

Reversing the order of the two transformers results in a global state, which enjoys a different axiomatization.

$$store\ s \gg (m \mid n) \quad = \quad store\ s \gg m \mid n$$

For both variants it is straightforward to derive an implementation from the corresponding specification — in the first case $(\mid)$ is promoted through *STATE*, in the second case *store* is promoted through *BACKTR*. Unfortunately, some harder cases remain, where the author has not been able to *derive* a promotion in a satisfying way. The problematic operations are, in general, those where the interaction with $(\gg\!=)$ is not explicitly specified. For instance, it is not clear how to derive the promotion of *call* through the state monad transformer.

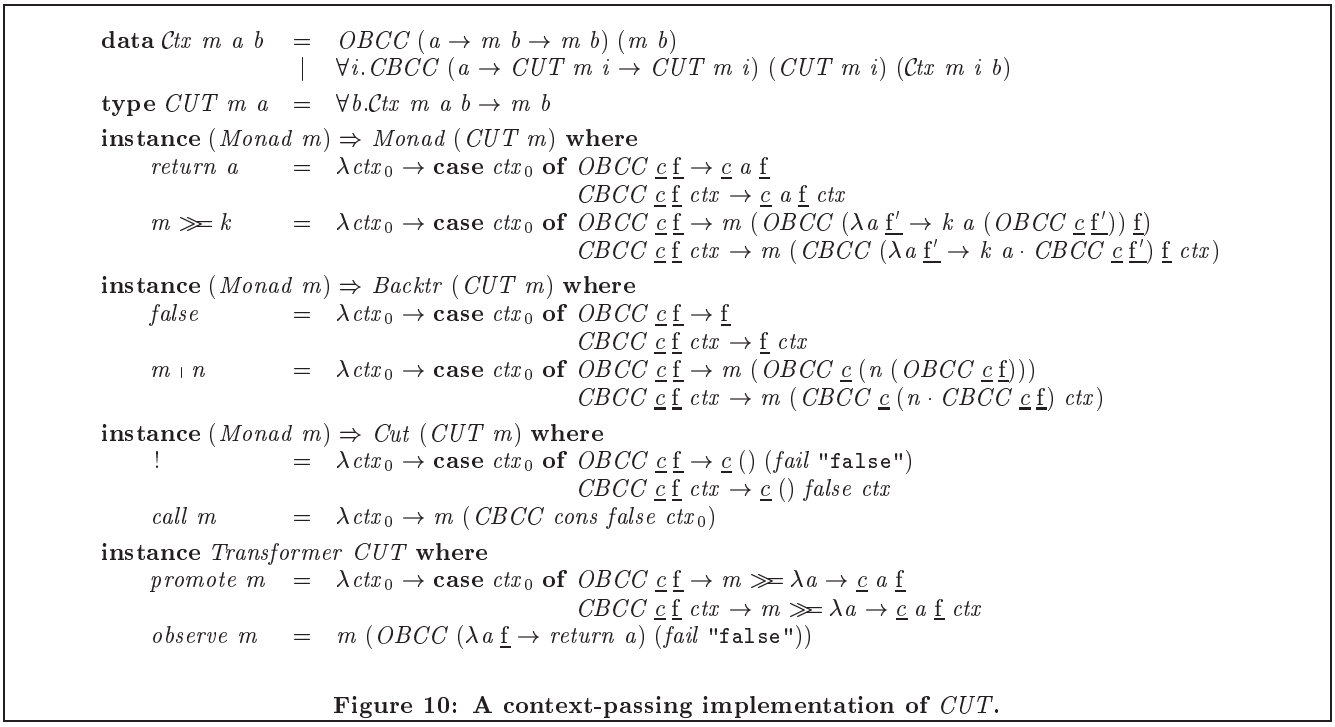

```
data Ctx m a b   =  OBCC (a → m b → m b) (m b)
                 |  ∀i. CBCC (a → CUT m i → CUT m i) (CUT m i) (Ctx m i b)
type CUT m a     =  ∀b. Ctx m a b → m b
instance (Monad m) ⇒ Monad (CUT m) where
  return a       =  λctx0 → case ctx0 of OBCC c f → c a f
                                         CBCC c f ctx → c a f ctx
  m ≫= k         =  λctx0 → case ctx0 of OBCC c f → m (OBCC (λa f' → k a (OBCC c f')) f)
                                         CBCC c f ctx → m (CBCC (λa f' → k a · CBCC c f') f ctx)
instance (Monad m) ⇒ Backtr (CUT m) where
  false          =  λctx0 → case ctx0 of OBCC c f → f
                                         CBCC c f ctx → f ctx
  m ∣ n          =  λctx0 → case ctx0 of OBCC c f → m (OBCC c (n (OBCC c f)))
                                         CBCC c f ctx → m (CBCC c (n · CBCC c f) ctx)
instance (Monad m) ⇒ Cut (CUT m) where
  !              =  λctx0 → case ctx0 of OBCC c f → c () (fail "false")
                                         CBCC c f ctx → c () false ctx
  call m         =  λctx0 → m (CBCC cons false ctx0)
instance Transformer CUT where
  promote m      =  λctx0 → case ctx0 of OBCC c f → m ≫= λa → c a f
                                         CBCC c f ctx → m ≫= λa → c a f ctx
  observe m      =  m (OBCC (λa f → return a) (fail "false"))
```

**Figure 10: A context-passing implementation of *CUT*.**

## 7. ACKNOWLEDGMENTS

I would like to thank four anonymous referees for their valuable comments.

## 8. REFERENCES

[1] M. Billaud. Simple operational and denotational semantics for Prolog with cut. *Theoretical Computer Science*, 71(2):193–208, March 1990.

[2] R. Bird. *Introduction to Functional Programming using Haskell*. Prentice Hall Europe, London, 2nd edition, 1998.

[3] A. de Bruin and E. de Vink. Continuation semantics for prolog with cut. In J. Díaz and F. Orejas, editors, *Proceedings of the International Joint Conference on Theory and Practice of Software Development : Vol. 1*, LNCS 351, pages 178–192. Springer-Verlag, 1989.

[4] N. Dershowitz and J.-P. Jouannaud. Rewrite systems. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*, chapter 6, pages 243–320. Elsevier Science Publishers B.V. (North Holland), 1990.

[5] R. Hinze. Prolog's control constructs in a functional setting — Axioms and implementation. *International Journal of Foundations of Computer Science*, 2000. To appear.

[6] J. Hughes. The design of a pretty-printing library. In J. Jeuring and E. Meijer, editors, *Advanced Functional Programming*, LNCS 925, pages 53–96. Springer-Verlag, 1995.

[7] M. P. Jones and L. Duponcheel. Composing monads. Technical Report YALEU/DCS/RR-1004, Department of Computer Science, Yale University, December 1993.

[8] K. Läufer and M. Odersky. An extension of ML with first-class abstract types. In *Proceedings of the 1992 ACM Workshop on ML and its Applications, San Francisco, California*, pages 78–91. ACM-Press, 1992.

[9] S. Liang, P. Hudak, and M. Jones. Monad transformers and modular interpreters. In *Proceedings of the 21st ACM Symposium on Principles of Programming Languages, San Francisco, California*, pages 333–343. ACM-Press, 1995.

[10] E. Meijer. *Calculating Compilers*. PhD thesis, Nijmegen University, 1992.

[11] E. Moggi. An abstract view of programming languages. Technical Report ECS-LFCS-90-113, Department of Computer Science, Edinburgh University, 1990.

[12] E. Moggi. Notions of computation and monads. *Information and Computation*, 93(1):55–92, 1991.

[13] S. Peyton Jones and J. Hughes, editors. *Haskell 98 — A Non-strict, Purely Functional Language*, February 1999. Available from `http://www.haskell.org/definition/`.

[14] L. Sterling and E. Shapiro. *The Art of Prolog: Advanced Programming Techniques*. The MIT Press, 1986.

[15] P. Wadler. Comprehending monads. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming, Nice*, pages 61–78. ACM-Press, 1990.

[16] P. Wadler. The essence of functional programming. In *Proceedings of the 19th ACM Symposium on Principles of Programming Languages, Sante Fe, New Mexico*, pages 1–14. ACM-Press, 1992.

[17] P. Wadler. Monads for functional programming. In J. Jeuring and E. Meijer, editors, *Advanced Functional Programming*, LNCS 925, pages 24–52. Springer-Verlag, 1995.