Tracing monadic computations and representing effects

Maciej Piróg and Jeremy Gibbons

Department of Computer Science
University of Oxford

Maciej.Pirog@cs.ox.ac.uk   Jeremy.Gibbons@cs.ox.ac.uk

In functional programming, monads are supposed to encapsulate computations, effectfully producing the final result, but keeping to themselves the means of acquiring it. For various reasons, we sometimes want to reveal the internals of a computation. To make that possible, in this paper we introduce monad transformers that add the ability to automatically accumulate observations about the course of execution as an effect. We discover that if we treat the resulting trace as the actual result of the computation, we can find new functionality in existing monads, notably when working with non-terminating computations.

1   Introduction

Consider a simple database in which different users store data under keys. This can be represented in Haskell as a data structure of type $[(User,[(Key,Data)])]$. We can define a function *getData* which looks up a value for a specified user and a key.

$$getData :: [(User,[(Key,Data)])] \rightarrow User \rightarrow Key \rightarrow Maybe\ Data$$
$$getData\ db\ u\ k = lookup\ u\ db \ggg lookup\ k$$

If a user *u* has an entry *d* associated with a key *k* in the database, *getData u k* returns *Just d*; otherwise, it returns *Nothing*. In the latter case we might want to inform the user 'why' the program is not able to deliver data: they might have misspelled their username, which means that *lookup u* will fail, or they might have tried to read from a missing key, which means that *lookup k* will fail. Unfortunately, the *Maybe* monad does not allow one to observe at what point a failing program actually fails. We need to structure our function using a more sophisticated monad.

What are the desired properties of such a monad? For sure, we want it to employ the same kind of effects as *Maybe*, so that we do not have to alter the logic of our program. We would like to have a *lift* operation, which allows us to automatically translate some operations from *Maybe* into the new monad, so that it is not necessary to rewrite standard functions like *lookup*. But most importantly, it must also reveal some observations about the course of execution (a *trace*), such as the number of successful subcomputations, from which we can extract the desired information about the point of failure. One possibility is to explicitly accumulate such observations inside the computation, using monads like *Writer* or *Exception*. In this article, however, we take a different approach: we automate the accumulation inside the monadic structure. The accumulation is transparent inside the computation; that is, it cannot affect the course of execution, and is revealed only on the outside.

We aim for maximum genericity: we construct transformers that add traces to arbitrary monads. Our main tools are free monads (Section 2), which have the ability to represent the very structure of monadic computations. They provide, in a sense, the most informative traces possible. Then, we introduce a transformer, called *Nest*, which allows one to mix the free and effectful computations provided by a monad (Section 3). The genericity pays off, and we find a number of applications for tracing monads:

- Traces of computations can be interpreted in different ways. In Section 4.3 we show an example in which, by revealing the inner structure of a computation in the *List* monad, we can define the Prolog *cut* operator.

- Computations encapsulated in monadic expressions are often monolithic. They are supposed to produce final values only, so there is little space for non-termination. But, assuming non-strict semantics, we can see computations as entities that lazily unfold a trace. So, even if a computation is non-terminating, we can benefit from its infinite trace. If we interpret free parts of a *Nest* value as terms, we discover that *Nest* is a generalisation of Capretta's partiality monad (see Section 4.1).

- For the same reasons, traces are a useful tool for modelling impure interaction with the environment. In Section 6, we sketch out some future work on a novel, coinductive approach to the functional semantics of effects.

2   Simple tracing with free monads

Moggi [19] called monads *notions of computation*, because they describe computational effects in a way that abstracts from the type of values produced by a computation. In a sense, from the categorical point of view, monads abstract even from the exact values produced by the computation, since *return* and *join* are natural transformations. It is the *bind* operator (defined as $m \ggeq f = join\ (fmap\ f\ m)$) that mixes the functorial (looking only at the values) *fmap* and parametric (looking only at the structure) *join*. The laws for monads and functors entail the following equality, called *naturality* of join.

$$fmap\ f \circ join \equiv join \circ fmap\ (fmap\ f)$$

Read as a transformation from left to right, it allows one to move occurrences of *join* to the left of a composition, and occurrences of *fmap* to the right. This way, we can split computations with multiple steps into a "mapping" part and a "joining" part. For example, consider a computation $m \ggeq f \ggeq g$. It can be split as follows.

$$m \ggeq f \ggeq g \equiv (join \circ fmap\ g \circ join \circ fmap\ f)\ m \equiv (join \circ join \circ fmap\ (fmap\ g) \circ fmap\ f)\ m$$

Intuitively, we can see the "mapping" part (that is, $fmap\ (fmap\ g) \circ fmap\ f$) as the construction of the computation and the "joining" part ($join \circ join$) as the execution of effects.

Our first approach to tracing is to capture the mapping part of a computation. We suspend execution of the *join* operators of the monad, so that we can isolate and examine the elements from which the computation is composed.

2.1   Free monads

The type of the mapping part is different for different numbers of nested occurrences of the *fmap* function, and so for computations consisting of different numbers of steps. For example,

Just `'a'`	:: *Maybe Char*
fmap Just (*Just* `'a'`)	:: *Maybe* (*Maybe Char*)
fmap (*fmap Just*) (*fmap Just* (*Just* `'a'`))	:: *Maybe* (*Maybe* (*Maybe Char*))

A reasonable class of datatypes in which to store such expressions are the *free monads*, also known as *f-generated trees*. Allowing ourselves to define monads in terms of *fmap*, *return*, and *join*, rather than the *return* and $\ggeq$ required by Haskell, we can write:

```
data Free f a = Wrap (f (Free f a)) | Return a
instance Functor f ⇒ Functor (Free f) where
   fmap g (Return a) = Return (g a)
   fmap g (Wrap f)   = Wrap (fmap (fmap g) f)
instance Functor f ⇒ Monad (Free f) where
   return            = Return
   join (Return f)   = f
   join (Wrap f)     = Wrap (fmap join f)
```

Each level of the type constructor of a monad corresponds to a level in the *Free* data structure. We store *Just* ʼaʼ :: *Maybe Char* as *Wrap* (*Just* (*Return* ʼaʼ)), and *Just* (*Just* (ʼaʼ)) :: *Maybe* (*Maybe Char*) as *Wrap* (*Just* (*Wrap* (*Just* (*Return* ʼaʼ)))), and so on.

2.2 Monad transformers with *drop*

The type *Free m a* can be seen as a datatype of terms generated by the signature *m* (a functor) and a set of variables *a*. Even if *m* is a monad, *Free m* cannot depend on any effects provided by *m*; the join operation for *Free m* performs only substitution, and is independent of the join for *m*.

In order to couple a monad *m* and the *m*-generated free monad, we need the notion of *monad transformers* [16]. A monad transformer with *drop* is a relation between two monads, *m* and *t*, characterised by two functions, *lift* and *drop*, which translate computations in *m* into computations in *t*, and *vice versa*. The relationship can be defined as the following two-parameter Haskell type class. Though in functional programming *drop* is rarely considered to be a member of *MonadTrans*, it plays an important role here. For the moment, we forget that we usually insist that monad transformers are subject to a certain set of algebraic laws.

```
class (Monad m, Monad t) ⇒ MonadTrans m t | t → m where
   lift :: m a → t a
   drop :: t a → m a
```

For any monad *m*, we define an instance of *MonadTrans m* (*Free m*) as follows. The *lift* operation is straightforward, as it only wraps the value and maps the *Return* constructor. The *drop* operation traverses the structure and flattens each level, thus performing suspended *bind*s of *m*.

```
instance (Functor m, Monad m) ⇒ MonadTrans m (Free m) where
   lift              = Wrap ∘ fmap Return
   drop (Return a) = return a
   drop (Wrap m)   = m ⟫= drop
```

2.3 Examples

Now, we can test our tracing free monad transformer on the *Maybe* monad. The *lift* function allows us to translate any computation in *m* into a computation in *Free m*. To get the original, non-traced computation back, we use the *drop* function.

We can see every composition *Wrap* ∘ *Just* as a "tick", given for each lifted successful subcomputation. The trace forms a unary counter, storing the number of ticks. The final value of the computation is stored in the last *Wrap*. Consider the following conversation with the Haskell interactive shell.

▷ **do** { *lift* (*Just* 2); *lift* (*Just* 4); *lift Nothing* }
Wrap (*Just* (*Wrap* (*Just* (*Wrap Nothing*))))
▷ *drop* (**do** { *lift* (*Just* 2); *lift* (*Just* 4); *lift Nothing* })
Nothing

Similarly, for the *Writer* monad, we can get a list of all the values appended to the monoid followed by the final return value. (For brevity, we show *Writer* (*a*, *Sum s*) as (*s*, *a*).)

▷ **do** { *lift* (*tell* (*Sum* 2)); *lift* (*tell* (*Sum* 3)); *lift* (*tell* (*Sum* 7)); *return* 'a' }
Wrap (2, *Wrap* (3, *Wrap* (7, *Return* 'a')))
▷ *drop* (**do** { *lift* (*tell* (*Sum* 2)); *lift* (*tell* (*Sum* 3)); *lift* (*tell* (*Sum* 7)); *return* 'a' })
(12, 'a')

An important thing to notice is that non-terminating computations in *Writer* do not make much sense. For example, the following computation just diverges.

▷ **let** *w n* = **do** { *tell* (*Sum n*); *w* (*n* + 1) } **in** *w* 0
(*** Exception : stack overflow

In contrast, with the tracing version of *Writer*, we can enjoy an infinite stream of actions that happen during the execution:

▷ **let** *w n* = **do** { *lift* (*tell* (*Sum n*)); *w* (*n* + 1) } **in** *w* 0
Wrap (0, *Wrap* (1, *Wrap* (2, *Wrap* (3, *Wrap* (4, *Wrap* (5, *Wrap* (6, *Wrap* (7, *Wrap* (8, *Wrap*...

3   More advanced tracing with free structures

Tracing computations with free structures is not very flexible: every bind and every join is suspended, creating a new level of structure every time a monadic action is called. In some circumstances we would like to trace only selected parts of the computation—perhaps we are confident that the other parts always succeed, or we want to treat selected pieces of the computation monolithically, and we are not interested in a fine-grained report about their behaviour.

Another issue is the algebraic properties of monad transformers with *drop*. Intuitively, a pair of monads *m* and *t* are related as a monad transformer if *t* incorporates at least the same effects as *m*. This is usually formalised with the following equalities [12].

$$
\begin{aligned}
lift\ (return\ a) &= return\ a \\
lift\ (m \ggg f) &= lift\ m \ggg (lift \circ f) \\
drop\ (return\ a) &= return\ a \\
drop\ (lift\ m \ggg f) &= m \ggg drop \circ f
\end{aligned}
$$

The first two equalities state that *lift* is a monad morphism. The instance *MonadTrans m* (*Free m*) from Section 2.2 does not have this property. For example, *lift* (*Just* 1) ≫ *lift* (*Just* 2) is equal to *Wrap* (*Just* (*Wrap* (*Just* (*Return* 2)))), while *lift* (*Just* 1 ≫ *Just* 2) is equal to *Wrap* (*Just* (*Return* 2)).

For these two reasons, we abandon the idea of using free monads directly to trace computations. Instead, in this section we introduce a general class, *MonadTrace*, which allows one to specify the points at which to make observations about the execution, and a corresponding version of the *Free* monad that is a proper monad transformer.

3.1 The *MonadTrace* class

The *MonadTrace* class introduces a single monadic value, *mark*. Intuitively, this is an operation that stores the current state of execution in the trace.

> **class** *Monad t* $\Rightarrow$ *MonadTrace t* **where**
>   *mark* :: t ()

We call a monad v a *tracer* of a monad m, if m and v form a monad transformer (*MonadTrans m v*), and v is an instance of *MonadTrace*. Additionally, *lift*, *drop* and *mark* should satisfy the following equalities.

> $lift \quad (return\ a) = return\ a$
> $lift \quad (m \ggg f) = lift\ m \ggg (lift \circ f)$
> $drop\ (return\ a) = return\ a$
> $drop\ (v \ggg g) \ = drop\ v \ggg (drop \circ g)$
> $drop\ mark \qquad = return\ ()$

The tracer v cannot perform more effects than the monad m, except for tracing with the *mark* operation. Nonetheless, tracing does not affect the course of computation in any way observable from inside of the v-computation, hence both *lift* and *drop* are monad morphisms. Note that the laws entail $drop \circ lift = id$.

   We use the *mark* gadget wherever we want to make an observation. In the following example, the intuitive semantics of a tracer for the *Maybe* monad is that we get a tick whenever the computation is still successful when placing a mark.

> **do** $\{x \leftarrow lift\ m_0;\ y \leftarrow lift\ m_1;\ mark;\ z \leftarrow lift\ m_2;\ mark;\ return\ (x+y+z)\}$

That means that if m_0 is successful, but m_1 fails, no ticks are stored in the trace (intuitively, it is equivalent to *Nothing*). Only if both m_0 and m_1 are successful, the *mark* operation stores this success (the trace is of the form *Just a*, where a is a result of the rest of the computation). If m_2 is also successful, the second call to *mark* stores this information in the trace (and the trace is of the form *Just (Just a)*, where $a = return\ (x+y+z) = Just\ (x+y+z)$).

   We also define a convenient function, *mind*, to perform a traced lifting.

> $mind :: (MonadTrans\ m\ v,\ MonadTrace\ v) \Rightarrow m\ a \rightarrow v\ a$
> $mind\ m = $ **do** $\{x \leftarrow lift\ m;\ mark;\ return\ x\}$

Then the previous example can be written more concisely:

> **do** $\{x \leftarrow lift\ m_0;\ y \leftarrow mind\ m_1;\ z \leftarrow mind\ m_2;\ return\ (x+y+z)\}$

3.2 The *Nest* monad

Free monads allow one to capture the structure of an m-computation as data; but for tracing, we need also to be able to perform some parts of the computations (the lifted ones) immediately. This suggests considering the composition of the two monads m and *Free m*, in one order or the other. In fact, because we want the lifted parts to be performed immediately, the appropriate order of composition is to have m on the outside and *Free m* inside. We therefore define:

> **newtype** $Nest\ m\ a = Nest\ \{\,unNest :: m\ (Free\ m\ a)\,\}$
>
> **instance** $Functor\ m \Rightarrow Functor\ (Nest\ m)$ **where**
> $fmap\ f\ (Nest\ m) = Nest\ (fmap\ (fmap\ f)\ m)$

It remains to show that *Nest* can be given the structure of a monad. We do this using Jones and Duponcheel's *prod* construction [14]: given monad M with unit $return_M$ and multiplication $join_M$, and similarly monad F with $return_F$ and $join_F$, the composition $M\ F$ forms a monad with unit and multiplication given by

Diagrammatically:

$$return = return_M \circ return_F$$
$$join \ \ = join_M \circ fmap_M\ prod$$

provided that natural transformation $prod :: F\ M\ F \rightarrow M\ F$ satisfies the three properties

$$prod \circ return_F \qquad\quad = id \qquad\qquad (1)$$
$$prod \circ fmap_F\ return = return_M \qquad (2)$$
$$prod \circ fmap_F\ join \ \ = join \circ prod \quad (3)$$

Diagrammatically:

(In fact, the multiplication $join_F$ of F is not used; all that is required of F is for it to be a *premonad*.) It turns out that the definition of *prod* is—if not quite forced then at least—very strongly suggested by the necessity of satisfying these three properties.

In our case, M is the monad m, and F is the datatype *Free m* from Section 2. For brevity, let us write M for $fmap_M$, and similarly F for $fmap_F$; let us also write the coproduct type former as $+$ and the coproduct morphism as $\triangledown$, so that we can express the two constructors as one composite,

$$Wrap \triangledown Return :: M\ F + 1 \rightarrow F$$

and name its inverse,

$$out_F :: F \rightarrow M\ F + 1$$

Recall that the unit of the monad *Free m* is the constructor *Return*:

$$return_F = Return$$

Without loss of generality, we let

$$prod \circ (Wrap \triangledown Return) = prod_1 \triangledown prod_2$$

so that

$$prod_1 = prod \circ Wrap$$
$$prod_2 = prod \circ Return$$

The definition of $prod_2$ is forced:

$$
\begin{aligned}
& prod_2 \\
= \quad & \{ \text{ definition of } prod_2 \ \} \\
& prod \circ Return \\
= \quad & \{ \ F \text{ as a premonad: } return_F = Return \ \} \\
& prod \circ return_F \\
= \quad & \{ \text{ property (1) } \} \\
& id
\end{aligned}
$$

Now consider property (2). Expanding the left-hand side, we have

$$
\begin{aligned}
& prod \circ F \ return \\
= \quad & \{ \text{ datatype isomorphism } \} \\
& prod \circ F \ return \circ (Wrap \triangledown Return) \circ out_F \\
= \quad & \{ \text{ naturality of } Wrap \triangledown Return \ \} \\
& prod \circ (Wrap \triangledown Return) \circ (M \ F \ return + return) \circ out_F \\
= \quad & \{ \text{ definitions of } prod_1, prod_2 \ \} \\
& (prod_1 \triangledown prod_2) \circ (M \ F \ return + return) \circ out_F \\
= \quad & \{ \text{ coproducts } \} \\
& ((prod_1 \circ M \ F \ return) \triangledown (prod_2 \circ return)) \circ out_F
\end{aligned}
$$

and for the right-hand side we have

$$
\begin{aligned}
& return_M \\
= \quad & \{ \text{ datatype isomorphism } \} \\
& return_M \circ (Wrap \triangledown Return) \circ out_F \\
= \quad & \{ \text{ coproducts } \} \\
& ((return_M \circ Wrap) \triangledown (return_M \circ Return)) \circ out_F
\end{aligned}
$$

These two expressions must be equal, which implies that the equalities

$$prod_1 \circ M \ F \ return = return_M \circ Wrap$$
$$prod_2 \circ return \qquad = return_M \circ Return$$

must hold. The second equality follows from $prod_2 = id$ and the definition of *return*; we'll use the first one to derive a definition for $prod_1$. We have:

$$
\begin{aligned}
& prod_1 \circ M \ F \ return \\
= \quad & \{ \text{ (A) suppose that } prod_1 = prod_1' \circ M \ prod \ \} \\
& prod_1' \circ M \ prod \circ M \ F \ return \\
= \quad & \{ \text{ functors } \}
\end{aligned}
$$

$$prod'_1 \circ M\,(prod \circ F\,return)$$
$$=\quad \{\,(B)\text{ induction }\}$$
$$prod'_1 \circ M\,return_M$$
$$=\quad \{\,(C)\text{ suppose that }prod'_1 = return_M \circ Wrap \circ join_M\,\}$$
$$return_M \circ Wrap \circ join_M \circ M\,return_M$$
$$=\quad \{\,M\text{ is a monad }\}$$
$$return_M \circ Wrap$$

and so property (2) strongly suggests letting

$$prod_1 = return_M \circ Wrap \circ join_M \circ M\,prod$$

The three steps (A), (B), (C) need a little justification. For (A), we're heading towards a use of induction, so we require an occurrence of *M prod* at the end of $prod_1$. For (B), induction seems plausible, because this calculation takes places under a *Wrap* constructor. For (C), the final *M* $return_M$ is cancellable via $join_M$, so we're done—we let be $prod'_1$ be the final expression we want composed with this cancellation.

   We don't need any stronger justification for induction than mere plausibility, because we are using this calculation only to suggest a possible definition of $prod_1$ that we should then check more rigorously. This still leaves us also with having to check property (3), again by induction. Both of these proofs are presented in Appendix A.

   The final instance declaration for *Nest* is then as follows:

instance (*Functor m*, *Monad m*) $\Rightarrow$ *Monad* (*Nest m*) **where**
　　$return = Nest \circ return_M \circ Return$

　　$join = Nest \circ join_M \circ fmap_M\,prod \circ unNest$
　　　　where
　　　　　$prod\,(Return\,(Nest\,m)) = m$
　　　　　$prod\,(Wrap\,m) \qquad\qquad = (return_M \circ Wrap \circ join_M \circ fmap_M\,prod)\,m$

3.3   Tracing with *Nest*

For any monad *m*, the monad *Nest m* is a tracer. The *lift* function maps *Return* on values. It does not create any *Wrap* constructor, so lifted computations are single-level and are always joined when *join* for *Nest* is performed. The *drop* function traverses the free structure, and collapses levels which were previously separated by *Wrap* constructors. The *mark* gadget explicitly creates a new level by the *Wrap* constructor mapped on *Return*s. All future computations will be confined to the wrapped *Return*s, and so the previous structure is preserved.

instance (*Functor m*, *Monad m*) $\Rightarrow$ *MonadTrans m* (*Nest m*) **where**
　　$lift \quad = Nest \circ fmap\,Return$

　　$drop\,v = unNest\,v \ggg revert$
　　　　where
　　　　　$revert\,(Return\,a) = return\,a$
　　　　　$revert\,(Wrap\,m) \ = m \ggg revert$
instance (*Functor m*, *Monad m*) $\Rightarrow$ *MonadTrace* (*Nest m*) **where**
　　$mark \ \ = (Nest \circ return \circ Wrap \circ return \circ Return)\,()$

The proof that these definitions satisfy the laws of *MonadTrace* from Section 3.1 is given in Appendix B.

As an example of *mark*, consider the following conversation with the Haskell shell. The call of the *mark* operation in the second query maps *Wrap* on every element. This way, every future computation is confined to the inside of those two *Wrap* constructors. Whatever monadic computation is bound to the result, on the top-level there will always be a two-element list, because it is a node with only *Wrap* constructors as children.

$\triangleright$ *lift* $[1,2]$
Nest $[Return\ 1, Return\ 2]$

$\triangleright$ **do** $\{x \leftarrow lift\ [1,2];\ lift\ [0..x]\}$
Nest $[Return\ 0, Return\ 1, Return\ 0, Return\ 1, Return\ 2]$

$\triangleright$ **do** $\{lift\ [1,2];\ mark\}$
Nest $[Wrap\ [Return\ ()], Wrap\ [Return\ ()]]$

$\triangleright$ **do** $\{x \leftarrow lift\ [1,2];\ mark;\ lift\ [0..x]\}$
Nest $[Wrap\ [Return\ 0, Return\ 1], Wrap\ [Return\ 0, Return\ 1, Return\ 2]]$

4 Interpreting traces

The tracing construction has multiple applications. As in the motivating example from the introduction, we can track the course of computation, in order to identify any point of failure. But, stepping back from this particular example, we observe that a trace is simply a data structure, and the computation in the *Nest* monad is performed only if the data structure is forced. This gives us control over the process of execution, which can be useful in the context of non-terminating computations (although the inductive proofs will need strengthening in that context). We can interpret the free parts of a *Nest* value in any way we want, and thus mix in some new effects, not previously available with the original monad. This control over a computation from the outside strongly resembles the paradigm of aspect-oriented programming.

4.1 The partiality monad

Capretta introduced the partiality monad [3] to capture non-termination as an effect; this technique has applications in type theory, to model non-guarded recursion. The original formulation is as follows.

data *Partial a = Later* $(Partial\ a)$ | *Now a*

A structure of this type is a value wrapped in a (possibly infinite) number of *Later* constructors. It represents a computation sliced into layers. We can explicitly force any number of layers. The *Nest* monad can be seen as a monad transformer which allows computations structured by any monad to be sliced. The most basic case, *Nest Identity*, is indeed isomorphic to *Partial*.

In most programming languages, including Haskell, the $\vee$ operator is asymmetric—non-strict in its second argument (*True* $\vee \perp \equiv$ *True*) but strict in its first ($\perp \vee$ *True* $\equiv \perp$). In pure Haskell, it is impossible to define *parallel-or*—that is, a disjunction $\curlyvee$ with the property that *True* $\curlyvee \perp \equiv$ *True* $\equiv \perp \curlyvee$ *True*.

To implement such a disjunction in the real world, we need some kind of parallelism, so that both arguments are evaluated simultaneously; when either one terminates with *True* or both terminate with *False*, the computation of the disjunction is complete. We can purely approximate this behaviour—at

least for the case where undefined arguments arise from non-termination, rather than from any other reason. We do this by explicitly cutting the infinite computation into finite pieces, using the *Nest* tracer.

Consider the following function, which is an implementation of the so-called Collatz problem. It is suspected that for all $n > 0$ it returns *True*, but no proof for this claim is known.

$$
\begin{array}{l}
collatz :: Integer \rightarrow Bool \\
collatz\ 1 \qquad\qquad = True \\
collatz\ n \mid odd\ n \quad = collatz\ (3 \times n + 1) \\
\qquad\quad \mid otherwise = collatz\ (n \div 2)
\end{array}
$$

If we want to check whether at least one of two Collatz sequences ends, it is not the best idea to use the regular Haskell disjunction, since if the first one diverges, the whole function diverges too.

$$
\begin{array}{l}
oneOf :: Integer \rightarrow Integer \rightarrow Bool \\
oneOf\ a\ b = collatz\ a \lor collatz\ b
\end{array}
$$

A much safer solution is to chop the evaluation of the Collatz sequence into pieces. We use the *Nest* tracer for the *Identity* monad. This way, we make the evaluation incremental, which enables us to execute it in parallel. The scheduler is very simple and hidden in the definition of $\curlyvee$. It executes a piece from each argument in turn.

$$
\begin{array}{l}
collatzN :: Integer \rightarrow Nest\ Identity\ Bool \\
collatzN\ 1 \qquad\qquad = return \quad True \\
collatzN\ n \mid odd\ n \quad = mark \gg collatzN\ (3 \times n + 1) \\
\qquad\quad\ \mid otherwise = mark \gg collatzN\ (n \div 2) \\
(\curlyvee) :: Nest\ Identity\ Bool \rightarrow Nest\ Identity\ Bool \rightarrow Bool \\
Nest\ (Identity\ (Leaf\ False))\ \curlyvee\ Nest\ (Identity\ (Leaf\ False)) = False \\
Nest\ (Identity\ (Leaf\ True))\ \ \curlyvee\ \_ \qquad\qquad\qquad\qquad\quad = True \\
\_ \qquad\qquad\qquad\qquad\ \curlyvee\ Nest\ (Identity\ (Leaf\ True))\ \ = True \\
Nest\ (Identity\ (Node\ m))\quad \curlyvee\ x \qquad\qquad\qquad\qquad = x \curlyvee Nest\ m
\end{array}
$$

We can test the $\curlyvee$ operator as follows. Note that *collatz* diverges if applied to 0.

$$
\begin{array}{l}
\triangleright collatzN\ 120 \curlyvee collatzN\ 130 \\
True \\
\triangleright collatzN\ 0 \curlyvee collatzN\ 130 \\
True \\
\triangleright collatzN\ 130 \curlyvee collatzN\ 0 \\
True
\end{array}
$$

This model can easily be extended to different kinds of such thread races. For example, it is possible to simulate McCarthy's ambiguous choice operator $amb :: a \rightarrow b \rightarrow Either\ a\ b$ [18], which has the property that for $a_0 :: a$ and $b_0 :: b$, $amb\ a_0 \perp = Left\ a_0$, and $amb \perp b_0 = Right\ b_0$ (again, assuming that the undefined values arise from non-termination).

4.2 Approximating computations

For some monads, the initial segment of a *Nest* value may be seen as an "approximation" to the computation. Consider the monad of finitely supported probability distributions. Its most common representation is a list of probabilities paired with values.

```
newtype Distr a = Distr { runDistr :: [(Double, a)] }
instance Functor Distr where
    fmap f (Distr xs) = Distr (fmap (λ(p,a) → (p, f a)) xs)
instance Monad Distr where
    return a = Distr [(1,a)]
    join (Distr xs) = Distr [(p × q, a) | (p, Distr ys) ← xs, (q, a) ← ys]
```

However, this representation has a flaw. Consider the following problem: given the uniform distribution of $\{0,1\}$ (a fair coin), select uniformly an element from $\{0,1,2\}$. A solution is to flip the coin twice to get the uniform distribution of $\{0,1,2,3\}$; if you draw 0, 1, or 2, this is your answer, and if you draw 3, flip twice more. In Haskell:

```
coin :: Distr Int
coin = Distr [(0.5,0),(0.5,1)]

third :: Distr Int
third = do x ← coin
           y ← coin
           case (x,y) of
               (1,1) → return 0
               (1,0) → return 1
               (0,1) → return 2
               (0,0) → third
```

(This is a simplification of Knuth and Yao's technique to simulate a fair die using three fair coin tosses [15].) Though the solution is mathematically reasonable, the Haskell implementation is useless, because the *List* monad, and so also the *Distr* monad, gathers the results in a depth-first fashion. Though this may not be obvious at first sight, the recursive call in *third* is in the head of the list. Therefore, *third* actually diverges without producing any usable results: $third = \bot$.

The *Nest* monad can retrieve the situation, if we suspend the recursive call.

```
thirdN :: Nest Distr Int
thirdN = do x ← lift coin
            y ← lift coin
            case (x,y) of
                (1,1) → return 0
                (1,0) → return 1
                (0,1) → return 2
                (0,0) → mark ≫ thirdN
```

What can we do with *thirdN*? One possibility is to get an "approximation" of the structure with the following function, which cuts the subcomputations if the recursion is deeper than the specified argument. All the computations that are too deep are replaced with *Nothing*.

takeN :: (*Functor m*, *Monad m*) ⇒ *Int* → *Nest m a* → *Nest m* (*Maybe a*)
takeN k (*Nest m*) = *Nest* (*fmap* (*aux k*) *m*)
   where
      aux 0 (*Wrap* _) = *Return Nothing*
      aux k (*Wrap m*) = *Wrap* (*fmap* (*aux* (*k* − 1)) *m*)
      aux k (*Return a*) = *Return* (*Just a*)

approx :: (*Functor m*, *Monad m*) ⇒ *Int* → *Nest m a* → *m* (*Maybe a*)
approx k = *drop* ∘ *takeN k*

We can ask the Haskell shell:

▷ *approx* 0 *thirdN*
[(0.25, *Nothing*), (0.25, *Just* 2), (0.25, *Just* 1), (0.25, *Just* 0)]
▷ **let** *simpl* (*Distr xs*) = *Distr* (*map* (λ*x* → (*sum* [*p* | (*p*, *a*) ← *xs*, *x* ≡ *a*], *x*)) (*nub* (*fmap snd xs*)))
▷ *simpl* (*approx* 1 *thirdN*)
[(0.0625, *Nothing*), (0.3125, *Just* 2), (0.3125, *Just* 1), (0.3125, *Just* 0)]

▷ *simpl* (*approx* 2 *thirdN*)
[(0.015635, *Nothing*), (0.32813, *Just* 2), (0.32813, *Just* 1), (0.32813, *Just* 0)]

▷ *simpl* (*approx* 10 *thirdN*)
[(2.38419e−7, *Nothing*), (0.33333, *Just* 2), (0.33333, *Just* 1), (0.33333, *Just* 0)]

4.3   The Prolog *cut* operator

Prolog's *cut* operator allows one to restrict backtracking. The moment it is reached during evaluation of a predicate, it succeeds, but also discards all the possible backtrack choices created by the predicate so far. By exposing the structure of a *List* computation, we can use this effect also in Haskell. We perform a depth-first search on a rose tree of type *Nest* [ ], but once we go down a level, we never go back.

call :: *Nest* [ ] *a* → [*a*]
call (*Nest xs*) = *aux xs*
   where
      aux [ ]   = [ ]
      aux (*Return a* : *xs*) = *a* : *aux xs*
      aux (*Wrap as* : _) = *aux as*

brace :: *Nest* [ ] *a* → *Nest* [ ] *a*
brace = *lift* ∘ *call*

cut :: *Nest* [ ] ()
cut = *mark*

Consider the following example.

▷ *call* (**do** *x* ← *lift* [4, 7, 13, 9]
        y ← *lift* [2, 8, 1]
        when (*x* + *y* ⩾ 15) *cut*
        return (*x* + *y*))
[6, 12, 5, 9, 15]

(Here, *when* is a standard Haskell function, defined by *when b m* = **if** *b* **then** *m* **else** *return* ().) First, we pick 4 from the first list, which after choices from the second list creates $[6,12,5]$. Then, we pick 7 from the first list. We cut on the second element of the second choice (because $7+8 = 15$). So, all the other choices from the second list (that is, 1) are discarded, as well as other choices from the first list (that is, 13 and 9).

We can limit the scope of *cut* by using the *brace* function. Only choices from inside of the brace are now cut.

$$\triangleright call \ (\ \textbf{do} \ x \leftarrow lift \ [4,7,13,9]$$
$$brace \ (\ \textbf{do} \ y \leftarrow lift \ [2,8,1]$$
$$when \ (x+y \geqslant 15) \ cut$$
$$return \ (x+y) \ ) \ )$$
$$[6,12,5,9,15,15,11,17]$$

Different interpretations of the *Nest* data structure enable the definition of different search strategies, such as breadth-first search. Moreover, it can even mix two different strategies in lifted and minded parts.

4.4 Poor man's concurrency transformer, revisited

Claessen's "poor man's" concurrency transformer [5] adds simple concurrency capabilities to any monad. It has two flaws. The first one is that it does not respect the laws presented in Section 3. Every lifted operation is atomic, and execution can only be interrupted in between atomic actions; this means that the evaluation of *lift* $m_1 \gg lift \ m_2$ can be interrupted by an action from another thread, while *lift* $(m_1 \gg m_2)$ cannot. The second flaw is that the return type of its *run* function is *m* (), and so it does not allow one to collect actual results of the computation. Here, we give a version of Claessen's transformer which fixes these flaws, via a conscious use of free structures.

In Claessen's monad, the user first builds a continuation, which produces an expression, which then is interpreted by the *run* function. By augmenting the type of expressions, we can skip the continuation layer. The datatype for concurrent expressions is as follows.

data *Action m a* = *Par* (*Action m a*) (*Action m a*) | *Act* (*m* (*Action m a*)) | *Done a* | *Kill*

Intuitively, the *Par* constructor pairs two expressions for parallel evaluation, *Act* performs a single monadic action, *Done* terminates the computation with an answer[1], and *Kill* terminates the computation with no answer. We treat *Action* as a term algebra with *Done* as a constructor for variables.

instance *Functor m* $\Rightarrow$ *Monad* (*Action m*) **where**
  return   = *Done*
  Par a b $\gg=f$ = *Par* (*a* $\gg=f$) (*b* $\gg=f$)
  Act m $\gg=f$ = *Act* (*fmap* ($\gg=f$) *m*)
  Done a $\gg=f$ = *f a*
  Kill $\gg=f$ = *Kill*

The *Action* datatype describes a program, but does not specify which actions form atomic chunks that should not be interrupted by other operations. This task is delegated to the *Nest* transformer. The

---

[1]The *Kill* constructor is called *Stop* in Claessen's datatype, and *Done* is new.

type of our concurrent monad is then *Nest* (*Action m*) *a*, for a monad *m* and answer type *a*. We define a number of operations, which allow easy construction of concurrent expressions. The functions *done* and *kill* lift the appropriate constructors. A single operation can be embedded in an *Action* data structure and lifted to the concurrent monad with *act*. There are two operators for concurrency, *par* and *fork*. The former constructs a computation from two computations of the same type. The latter starts an auxiliary thread, whose final value is ignored (see the examples below).

> **type** *Concurrent m* = *Nest* (*Action m*)
>
> *done* :: (*Monad m*) ⇒ *a* → *Concurrent m a*
> *done* = *lift* ∘ *Done*
>
> *kill* :: (*Monad m*) ⇒ *Concurrent m a*
> *kill* = *lift Kill*
>
> *act* :: (*Monad m*) ⇒ *m a* → *Concurrent m a*
> *act m* = *lift* (*Act* (*liftM Done m*))
>
> *par* :: (*Monad m*) ⇒ *Concurrent m a* → *Concurrent m a* → *Concurrent m a*
> *par* (*Nest* m_1) (*Nest* m_2) = *Nest* (*Par* (*Done* (*Wrap* m_1)) (*Done* (*Wrap* m_2)))
>
> *fork* :: (*Monad m*) ⇒ *Concurrent m b* → *Concurrent m* ()
> *fork m* = *par* (*m* ≫ *kill*) (*act* (*return* ()))

We schedule such computations with the following *round* function. We can see *Done* as a constructor which either terminates evaluation of an atomic chunk (when it is composed with *Wrap*) or the entire thread (when it is composed with *Return*).

> *round* :: *Monad m* ⇒ [*Nest* (*Action m*) *x*] → *m* [*x*]
> *round* [ ]   = *return* [ ]
> *round* (*Nest w* : *as*) = **case** *w* **of**
>   *Kill*   → *round as*
>   *Done* (*Return x*) → **do** {*xs* ← *round as*; *return* (*x* : *xs*)}
>   *Done* (*Wrap a*)   → *round* (*as* ++ [*Nest a*])
>   *Act m*   → **do** {*a* ← *m*; *round* ([*Nest a*] ++ *as*)}
>   *Par a b*   → *round* ([*Nest b*] ++ *as* ++ [*Nest a*])

We can test our monad as follows. In the first example, we first define two expressions: *cat* writes the string `"cat"` five times, relinquishing control every time the operation is performed. Similarly, *fish* writes `"fish"` seven times.

> **instance** (*Monoid s*) ⇒ *MonadWriter s* (*Concurrent* (*Writer s*)) **where**
>   *tell* = *act* ∘ *tell*
>
> *cat* :: *Concurrent* (*Writer String*) *Int*
> *cat* = *replicateM* 5 (*tell* `"cat"` ≫ *mark*) ≫ *return* 1
>
> *fish* :: *Concurrent* (*Writer String*) *Int*
> *fish* = *replicateM* 7 (*tell* `"fish"` ≫ *mark*) ≫ *return* 2

We can test them, by running them in parallel.

> ▷ *round* [**do** *x* ← *fish* `par` *cat*
>   *tell* `"dog"`

$$\quad\quad\quad\quad\quad\quad return\ x]$$
$$(\texttt{"catfishcatfishcatfishcatfishcatfishdogfishfishdog"}, [1,2])$$

The results of all parallel threads called with *par*, in this example $[1,2]$, are returned in a list. The operation *tell* `"dog"` is bound to both threads.

We can now run *fish* in a separate, auxiliary thread. The thread is run on the side, the following actions are not bound to it, and its result is not returned with the overall result.

$$\triangleright\ round\ [\textbf{do}\ fork\ fish$$
$$\quad\quad\quad\quad\quad x \leftarrow cat$$
$$\quad\quad\quad\quad\quad tell\ \texttt{"dog"}$$
$$\quad\quad\quad\quad\quad return\ x]$$
$$(\texttt{"catfishcatfishcatfishcatfishcatfishdogfishfish"}, [1])$$

We can also define a version of *fish* that is performed atomically.

$$fish' :: Concurrent\ (Writer\ String)\ Int$$
$$fish' = replicateM\ 7\ (tell\ \texttt{"fish"}) \gg return\ 2$$

We can see that *round* does not separate the calls of *tell* `"fish"`.

$$\triangleright\ round\ [\textbf{do}\ fork\ fish'$$
$$\quad\quad\quad\quad\quad x \leftarrow cat$$
$$\quad\quad\quad\quad\quad tell\ \texttt{"dog"}$$
$$\quad\quad\quad\quad\quad return\ x]$$
$$(\texttt{"catfishfishfishfishfishfishfishcatcatcatcatdog"}, [1])$$

5   Related work

The idea of separation of syntax and semantics of monadic computations is not new. It is the very foundation of the success of monads as a tool for encapsulating impure behaviour in pure languages [21]—for example, the way Haskell integrates impure effects such as I/O within a pure language is to make the pure evaluation *construct* a syntactic term, which is subsequently *interpreted* by the run-time system.

The work most related to ours is the Unimo framework introduced by Lin [17], which is an embedded domain-specific language designed to modularise construction of monads in Haskell. Even though Lin's motivation and toolbox significantly differ from ours, he came to the same conclusion that exposing the structure of computations allows more functionality to be added to existing monads. We share a strong flavour of aspect-oriented programming.

There is a resemblance between resumptions, used for modelling semantics of concurrency [8, 11], and the *Nest* monad. A *resumption monad transformer* is used, for example, by Papaspyrou to model semantics of concurrency in domain theory [20], and by Harrison in his "cheap" concurrency [9, 10]. They both use a definition similar to ours from Section 2 (though they fail to mention the connection between free monads and resumptions), and as a result their constructions are not transformers in the sense of Section 3.

Harrison [9] is not quite right in claiming that Claessen's monad is based on first-class continuations. A version of the continuation monad is used only to build a syntactic term, which serves as the backbone

of the concurrent computation. The term is not a monad, since it lacks free variables, but it reveals the structure of resumptions (notice the type of the constructor $Atom :: m\ (Action\ m) \rightarrow Action\ m$). Similarly, in Swierstra and Altenkirch's functional specification of concurrency [23], resumptions are used implicitly, and control is surrendered by a thread whenever it wants to communicate with other parts of the concurrent system (which is denoted by a constructor of the free structure IO_c). As we show in Sections 4.1 and 4.4, the *Nest* monad transformer allows one to separate the concepts of composition of computations and yielding, by an explicit use of *mark*.

A definition of a resumption transformer, which satisfies the laws from Section 3, and is in fact isomorphic to *Nest*, was already given by Cenciarelli and Moggi [4], but practical applications in programming were not studied. Also—as pointed to us by the anonymous reviewers—the *Nest* transformer can be obtained via Hyland, Plotkin and Power's *sum* construction [13], as a sum of a monad and an *Identity*-generated free monad. Though Hyland *et al.*'s construction provides a simpler proof that *Nest* can be given a monad structure (using a distributive law between a monad m and the m-generated free monad), our definition of *join* is not intensionally similar to the one arising from the *sum* construction, and issues like efficiency should also be taken into account. A naive implementation of *join* in the sum construction traverses the structure twice (once to apply the distributive law, and once to join the free structure), while *join* for *Nest* needs to traverse the structure only once. On the other hand, the *sum* construction allows one to include an additional functor—the datatype in question is of the form $M\ (Free\ (F \circ M)\ a)$—which may help to generalize our notion of tracing in the future.

The interleaving between pure data and monadic structure was also considered by Filinski and Støvring [7], and in forthcoming work by Atkey *et al.* [1]. They give proof principles for reasoning about datatypes that include effects, for example a stream in which tails are always guarded by I/O actions.

6   Future work

So far, we failed to mention the mother and father of all purely functional monads, that is the continuation and state monads. We do not have much to say about continuations, but we see a lot of applications in tracing the *State* monad.

Functional specifications of effects.   The idea behind functional specifications of effects in pure languages is to model the logic of an effectful construct in the pure core of the language. For example, such a specification may consist of a datatype representing a model of the outside world and a variation of the *State* monad whose state modifications mirror the actions of the side-effecting monad. This way, we can translate a program into its pure equivalent, test it, and reason about it no differently from how we would reason about any other pure program.

The existing frameworks for specifying "effectful" Haskell in "pure" Haskell—like those proposed by Swierstra and Altenkirch [22, 23] and Butterfield *et al.* [2, 6]—do not concentrate on non-terminating computations, which are of little use in the pure world, but which are back in the spotlight in the presence of effects like I/O and concurrency. For example, Butterfield *et al.* model the interaction between programs and a filesystem by means of the *State* monad, which for an initial state (of the filesystem) produces a final value and a final state. In case of non-terminating programs, no final state exists, so the whole model becomes useless. What we are really interested in is not a final state, but the whole (possibly infinite) sequence of subsequent states of the filesystem, or a trace of all the interactions. In

such a setting, we can use coinduction as a reasoning tool, which coincides with the non-strict semantics of Haskell, in which we try to embed our model.

As mentioned before, our approach can transform monolithic computations into coinductive unfolding of traces. That is why we propose to use a different monad as a basis for functional specifications. It is a monad which produces not only the final state, but the whole (possibly infinite) stream of intermediate states.

> **data** *Trace s a* = *TCons s* (*Trace s a*) | *Nil a*
>
> **newtype** *States s a* = *States* { *runStates* :: *s* → *Trace s a* }

We leave the exact implementation of instances of *Monad* (*States s*), *MonadTrans* (*State s*) (*States s*) and *MonadTrace* (*States s*) to the reader as an exercise. The *mark* operation should accumulate the current state in the *Trace*.

What is the relationship between *States s* and *Nest* (*State s*)? It is possible to interpret free parts of *Nest* (*State s*) in a suitable way, that is to define a monad morphism of type *Nest* (*State s*) → *States s*. We can also suspect a different kind of generality, since *State* is a composition of two adjoint functors, namely *State s* = *Reader s* ∘ *Writer s*, while *States s* = *Reader s* ∘ *Free* (*Writer s*).

References

[1] Robert Atkey, Neil Ghani, Bart Jacobs & Patricia Johann (2012): *Fibrational Induction Meets Effects*. In Lars Birkedal, editor: *Foundations of Software Science and Computational Structures*, *Lecture Notes in Computer Science* 7213, Springer, pp. 42–57. To Appear. 16

[2] Andrew Butterfield & Glenn Strong (2002): *Proving Correctness of Programs with I/O - A Paradigm Comparison*. In: *Selected Papers from the 13th International Workshop on Implementation of Functional Languages*, IFL '02, Springer-Verlag, London, UK, pp. 72–87, doi:10.1007/3-540-46028-4_5. 16

[3] Venanzio Capretta (2005): *General Recursion via Coinductive Types*. *Logical Methods in Computer Science* 1(2), pp. 1–18, doi:10.2168/LMCS-1(2:1)2005. 9

[4] Pietro Cenciarelli & Eugenio Moggi (1993): *A Syntactic Approach to Modularity in Denotational Semantics*. In: *Proceedings of the 5th Biennial Meeting on Category Theory and Computer Science*, CTCS '93, CWI Technical Report, Amsterdam, The Netherlands. 16

[5] Koen Claessen (1999): *A poor man's concurrency monad*. *Journal of Functional Programming* 9, pp. 313–323, doi:10.1017/S0956796899003342. 13

[6] Malcolm Dowse, Glenn Strong & Andrew Butterfield (2003): *Proving make correct: I/O proofs in Haskell and Clean*. In: *Proceedings of the 14th International Conference on Implementation of Functional Languages*, IFL'02, Springer-Verlag, Berlin, Heidelberg, pp. 68–83, doi:10.1007/3-540-44854-3_5. 16

[7] Andrzej Filinski & Kristian Støvring (2007): *Inductive reasoning about effectful data types*. In: *Proceedings of the 12th ACM SIGPLAN International Conference on Functional Programming*, ICFP '07, ACM, New York, NY, USA, pp. 97–110, doi:10.1145/1291151.1291168. 16

[8] Steven E. Ganz, Daniel P. Friedman & Mitchell Wand (1999): *Trampolined style*. In: *Proceedings of the fourth ACM SIGPLAN International Conference on Functional Programming*, ICFP '99, ACM, New York, NY, USA, pp. 18–27, doi:10.1145/317636.317779. 15

[9] William L. Harrison: *Cheap (But Functional) Threads*. Available at http://www.cs.missouri.edu/~harrison/drafts/CheapThreads.pdf. 15

[10] William L. Harrison (2006): *The Essence of Multitasking*. In: *1st International Conference on Algebraic Methodology and Software Technology (AMAST06)*, *LNCS* 4019, Springer-Verlag, doi:10.1007/11784180_14. 15

[11] Matthew Hennessy & Gordon Plotkin (1979): *Full Abstraction for a Simple Parallel Programming Language*. In Jirí Becvár, editor: *MFCS, Lecture Notes in Computer Science* 74, Springer, pp. 108–120, doi:10.1007/3-540-09526-8_8. 15

[12] Ralf Hinze (2000): *Deriving backtracking monad transformers*. In: *Proceedings of the fifth ACM SIGPLAN International Conference on Functional Programming*, ICFP '00, ACM, New York, NY, USA, pp. 186–197, doi:10.1145/351240.351258. 4

[13] Martin Hyland, Gordon Plotkin & John Power (2006): *Combining effects: sum and tensor*. *Theoretical Computer Science* 357, pp. 70–99, doi:10.1016/j.tcs.2006.03.013. 16

[14] Mark P. Jones & Luc Duponcheel (1993): *Composing Monads*. Technical Report, Yale University. 6

[15] Donald Knuth & Andrew Yao (1976): *The Complexity of Nonuniform Random Number Generation*. In J. F. Traub, editor: *Algorithms and Complexity: New Directions and Recent Results*, Academic Press. Also in *Selected Papers on the Analysis of Algorithms*. 11

[16] Sheng Liang, Paul Hudak & Mark Jones (1995): *Monad transformers and modular interpreters*. In: *Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '95, ACM, New York, NY, USA, pp. 333–343, doi:10.1145/199448.199528. 3

[17] Chuan-kai Lin (2006): *Programming monads operationally with Unimo*. In: *Proceedings of the eleventh ACM SIGPLAN International Conference on Functional Programming*, ICFP '06, ACM, New York, NY, USA, pp. 274–285, doi:10.1145/1159803.1159840. 15

[18] John McCarthy (1963): *A Basis for a Mathematical Theory of Computation*. In P. Braffort & D. Hirschberg, editors: *Computer Programming and Formal Systems*, North-Holland, Amsterdam, pp. 33–70, doi:10.1016/S0049-237X(08)72018-4. 10

[19] Eugenio Moggi (1991): *Notions of computation and monads*. *Information and Computation* 93, pp. 55–92, doi:10.1016/0890-5401(91)90052-4. 2

[20] Nikolaos S. Papaspyrou (2001): *A Resumption Monad Transformer and its Applications in the Semantics of Concurrency*. In: *Proceedings of the 3rd Panhellenic Logic Symposium*, Anogia, Crete. 15

[21] Simon L. Peyton Jones & Philip Wadler (1993): *Imperative functional programming*. In: *Proceedings of the 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '93, ACM, New York, NY, USA, pp. 71–84, doi:10.1145/158511.158524. 15

[22] Wouter Swierstra (2008): *A Functional Specification of Effects*. Ph.D. thesis, University of Nottingham. 16

[23] Wouter Swierstra & Thorsten Altenkirch (2007): *Beauty in the beast*. In: *Proceedings of the ACM SIGPLAN Workshop on Haskell*, Haskell '07, ACM, New York, NY, USA, pp. 25–36, doi:10.1145/1291201.1291206. 16

Appendix A   Proof that *Nest* is a monad

We prove the properties (2) and (3) from Section 3.2 by induction, assuming an initial-algebra reading of the datatype *Free m*—that is, we assume a well-founded ordering on the subterms of any value of type *Free m a*. (We have to resort to something like induction, because the definition of *prod* above isn't in the form of a standard recurson pattern—in particular, it is not a fold.) For brevity, we omit the *Nest* constructor.

Property (2): The case for *Return* is straightforward. For the *Wrap* case, we assume that the property holds for each element of data structure *m*; that is, that

$$M\ (prod \circ F\ return)\ m = M\ return_M\ m$$

Then we calculate:

$$
\begin{aligned}
& (prod \circ F\ return)\ (Wrap\ m) \\
=\ & \{\ \text{naturality of } Wrap :: M\ F \to F\ \} \\
& (prod \circ Wrap \circ M\ F\ return)\ m \\
=\ & \{\ \text{definition of } prod\ \} \\
& (return_M \circ Wrap \circ join_M \circ M\ prod \circ M\ F\ return)\ m \\
=\ & \{\ \text{functors}\ \} \\
& (return_M \circ Wrap \circ join_M \circ M\ (prod \circ F\ return))\ m \\
=\ & \{\ \text{induction}\ \} \\
& (return_M \circ Wrap \circ join_M \circ M\ return_M)\ m \\
=\ & \{\ M \text{ as a monad}\ \} \\
& (return_M \circ Wrap)\ m
\end{aligned}
$$

Property (3): The case for *Return* is again straightforward. In the *Wrap* case, we again assume that the property holds for each element of *m*:

$$M\ (join \circ prod)\ m = M\ (prod \circ F\ join)\ m$$

Then we calculate:

$$
\begin{aligned}
& (join \circ prod)\ (Wrap\ m) \\
=\ & \{\ \text{definition of } join\ \} \\
& (join_M \circ M\ prod \circ prod \circ Wrap)\ m \\
=\ & \{\ \text{definition of } prod\ \} \\
& (join_M \circ M\ prod \circ return_M \circ Wrap \circ join_M \circ M\ prod)\ m \\
=\ & \{\ \text{naturality of } return_M\ \} \\
& (join_M \circ return_M \circ prod \circ Wrap \circ join_M \circ M\ prod)\ m \\
=\ & \{\ M \text{ as a monad}\ \} \\
& (prod \circ Wrap \circ join_M \circ M\ prod)\ m \\
=\ & \{\ \text{definition of } prod\ \} \\
& (return_M \circ Wrap \circ join_M \circ M\ prod \circ join_M \circ M\ prod)\ m \\
=\ & \{\ \text{naturality of } join_M\ \} \\
& (return_M \circ Wrap \circ join_M \circ join_M \circ M\ M\ prod \circ M\ prod)\ m \\
=\ & \{\ M \text{ as a monad}\ \} \\
& (return_M \circ Wrap \circ join_M \circ M\ join_M \circ M\ M\ prod \circ M\ prod)\ m
\end{aligned}
$$

$=$ { functors }
$(return_M \circ Wrap \circ join_M \circ M\ (join_M \circ M\ prod \circ prod))\ m$
$=$ { definition of *join*; induction }
$(return_M \circ Wrap \circ join_M \circ M\ (prod \circ F\ join))\ m$
$=$ { functors }
$(return_M \circ Wrap \circ join_M \circ M\ prod \circ M\ F\ join)\ m$
$=$ { definition of *join* }
$(return_M \circ Wrap \circ join \circ M\ F\ join)\ m$
$=$ { definition of *prod* }
$(prod \circ Wrap \circ M\ F\ join)\ m$
$=$ { naturality of $Wrap :: M\ F \dot{\to} F$ }
$(prod \circ F\ join)\ (Wrap\ m)$

Appendix B   Proof that *Nest* is a tracer

Here, we prove that *Nest* is a tracer (see Sections 3.1 and 3.3 for the definitions). The equalities $drop \circ lift = id$ and $drop\ mark = return\ ()$ are straightforward.
   We prove the equality $lift \circ return_M = return_N$ as follows.

$lift \circ return_M$
$=$ { definition of *lift* }
$M\ Return \circ return_M$
$=$ { naturality of $return_M$ }
$return_M \circ Return$
$=$ { definition of $return_N$ }
$return_N$

The fact that $lift\ c \gg\!\!=_N lift \circ f = lift\ (c \gg\!\!=_M f)$ follows from the following.

$lift\ c \gg\!\!=_N lift \circ f$
$=$ { definition of $\gg\!\!=_N$ }
$join_N\ (N\ (lift \circ f)\ (lift\ c))$
$=$ { definition of *lift* }
$join_N\ (N\ (M\ Return \circ f)\ (M\ Return\ c))$
$=$ { definition of N }
$join_N\ (MF\ (M\ Return \circ f)\ (M\ Return\ c))$
$=$ { definition of $join_N$ }
$join_M\ (M\ prod\ (MF\ (M\ Return \circ f)\ (M\ Return\ c)))$
$=$ { functor }
$join_M\ (M\ (prod \circ F\ (M\ Return \circ f) \circ Return)\ c)$
$=$ { definition of F }
$join_M\ (M\ (prod \circ Return \circ M\ Return \circ f)\ c)$
$=$ { definition of *prod* }
$join_M\ (M\ (M\ Return \circ f)\ c)$
$=$ { naturality of $join_M$ }

$$(M\ Return \circ join_M)\ (M f\ c)$$
$$= \quad \{ \text{ definition of } lift \}$$
$$lift\ (join_M\ (M f\ c))$$
$$= \quad \{ \text{ definition of } join_M \}$$
$$lift\ (c \ggeq_M f)$$

To prove the fact that *drop* is a monad morphism we first prove a lemma

$$join_M \circ M\ revert \circ prod \circ F f = join_M \circ join_M \circ MM\ revert \circ M f \circ revert$$

The case for *Return* follows from simple unfolding of the definitions. The case for *Wrap* is as follows.

$$(join_M \circ M\ revert \circ prod \circ F f)\ (Wrap\ m)$$
$$= \quad \{ \text{ definition of } F \}$$
$$(join_M \circ M\ revert \circ prod)\ (Wrap\ (MF f)\ m))$$
$$= \quad \{ \text{ definition of } prod \}$$
$$(join_M \circ M\ revert \circ return_M \circ Wrap \circ join_M \circ M\ prod \circ MF f)\ m$$
$$= \quad \{ \text{ naturality of } return_M \}$$
$$(join_M \circ return_M \circ revert \circ Wrap \circ join_M \circ M\ prod \circ MF f)\ m$$
$$= \quad \{ \text{ monad laws } \}$$
$$(revert \circ Wrap \circ join_M \circ M\ prod \circ MF f)\ m$$
$$= \quad \{ \text{ definition of } revert \}$$
$$(join_M \circ M\ revert \circ join_M \circ M\ prod \circ MF f)\ m$$
$$= \quad \{ \text{ naturality of } join_M \}$$
$$(join_M \circ join_M \circ MM\ revert \circ M\ prod \circ MF f)\ m$$
$$= \quad \{ \text{ monad laws } \}$$
$$(join_M \circ M\ join_M \circ MM\ revert \circ M\ prod \circ MF f))\ m$$
$$= \quad \{ \text{ functor } \}$$
$$(join_M \circ M\ (join_M \circ M\ revert \circ prod \circ F f))\ m$$
$$= \quad \{ \text{ induction } \}$$
$$(join_M \circ M\ (join_M \circ M\ join_M \circ MM\ revert \circ M f \circ revert)\ m$$
$$= \quad \{ \text{ functor } \}$$
$$(join_M \circ M\ join_M \circ MM\ join_M \circ MMM\ revert \circ MM f \circ M\ revert)\ m$$
$$= \quad \{ \text{ monad laws } \}$$
$$(join_M \circ join_M \circ MM\ join_M \circ MMM\ revert \circ MM f \circ M\ revert)\ m$$
$$= \quad \{ \text{ naturality of } join_M \}$$
$$(join_M \circ M\ join_M \circ join\_m \circ MMM\ revert \circ MM f \circ M\ revert)\ m$$
$$= \quad \{ \text{ naturality of } join_M \}$$
$$(join_M \circ M\ join_M \circ MM\ revert \circ join_M \circ MM f \circ M\ revert)\ m$$
$$= \quad \{ \text{ naturality of } join_M \}$$
$$(join_M \circ join_M \circ MM\ revert \circ M f \circ join_M \circ M\ revert)\ m$$
$$= \quad \{ \text{ definition of } revert \}$$
$$(join_M \circ join_M \circ MM\ revert \circ M f \circ revert)\ (Wrap\ m)$$

To prove that $drop\ (c \ggeq_N f) = drop\ c \ggeq_M (drop \circ f)$, for $c :: a$ and $f :: a \to Nest\ m\ b$, we unfold the definition of $(\ggeq)$ and prove the equivalent equality $drop\ (join_N\ (N f\ c)) = join_M\ (M\ (drop \circ f)\ drop\ c)$. It follows from the commutativity of the following diagram (in Hask). For brevity, we write μ_N for the

join of monad *Nest*. The left-most path of the diagram is equal to $join_M \circ M \, (drop \circ f) \circ drop$, while the right-most path is equal to $drop \circ join_N \circ N f$.

The first column of the diagram depicts unfolding of the definitions of *drop*. The second column is equal to the first one due to the fact that $join_M$ is a natural transformation: the morphism $join_M :: MMA \to MA$ form the first column "travelled" down the path to settle as $join_M :: MMMB \to MMB$. Additionally, due to the monad law for $join_M$, we can exchange $MM \, join_M$ with $M \, join_M$. This allows us to use the lemma, the M-image of which commutes columns 2 and 3. Again, we can use the monad law to change $M \, join_M$ into $join_M$. We use the naturality of $join_M$ to switch its position with the mapping of *revert*, which justify the fourth column.