

FUNCTIONAL PEARL

Applicative programming with effects

CONOR MCBRIDE¹

University of Nottingham

ROSS PATERSON

City University, London

Abstract

In this article, we introduce *Applicative* functors – an abstract characterisation of an applicative style of effectful programming, weaker than *Monads* and hence more widespread. Indeed, it is the ubiquity of this programming pattern that drew us to the abstraction. We retrace our steps in this article, introducing the applicative pattern by diverse examples, then abstracting it to define the *Applicative* type class and introducing a bracket notation that interprets the normal application syntax in the idiom of an *Applicative* functor. Furthermore, we develop the properties of applicative functors and the generic operations they support. We close by identifying the categorical structure of applicative functors and examining their relationship both with *Monads* and with *Arrows*.

1 Introduction

This is the story of a pattern that popped up time and again in our daily work, programming in Haskell (Peyton Jones, 2003), until the temptation to abstract it became irresistible. Let us illustrate with some examples.

1.1 Sequencing commands

One often wants to execute a sequence of commands and collect the sequence of their responses, and indeed there is such a function in the Haskell Prelude (here specialised to the IO monad):

```
sequence :: [IO a] → IO [a]
sequence [] = return []
sequence (c : cs) = do
  x ← c
  xs ← sequence cs
  return (x : xs)
```

¹ McBride is funded by EPSRC grant EP/C512022/1.

In the $(c : cs)$ case, we collect the values of some effectful computations, which we then use as the arguments to a pure function $(:)$. We could avoid the need for names to wire these values through to their point of usage if we had a kind of ‘effectful application’. Fortunately, exactly such a thing lives in the standard Monad library:

```

ap :: Monad m => m (a -> b) -> m a -> m b
ap mf mx = do
  f <- mf
  x <- mx
  return (f x)

```

Using this function, we could rewrite `sequence` as follows:

```

sequence :: [IO a] -> IO [a]
sequence [] = return []
sequence (c : cs) = return (:) 'ap' c 'ap' sequence cs

```

where the `return` operation, which every Monad must provide, lifts pure values to the effectful world, whilst `ap` provides ‘application’ within it.

Except for the noise of the `returns` and `aps`, this definition is in a fairly standard applicative style, even though effects are present.

1.2 Transposing ‘matrices’

Suppose that we represent matrices (somewhat approximately) by lists of lists. A common operation on matrices is transposition:²

```

transpose :: [[a]] -> [[a]]
transpose [] = repeat []
transpose (xs : xss) = zipWith (:) xs (transpose xss)

```

Now, the binary `zipWith` is one of a family of operations that ‘vectorise’ pure functions. As Fridlender and Indrika (2000) point out, the entire family can be generated from `repeat`, which generates an infinite stream from its argument, and `zapp`, a kind of ‘zippy’ application:

```

repeat :: a -> [a]
repeat x = x : repeat x
zapp :: (a -> b) -> [a] -> [b]
zapp (f : fs) (x : xs) = f x : zapp fs xs
zapp _ _ = []

```

The general scheme is as follows:

```

zipWithn :: (a1 -> ... -> an -> b) -> [a1] -> ... -> [an] -> [b]
zipWithn f xs1 ... xsn = repeat f 'zapp' xs1 'zapp' ... 'zapp' xsn

```

In particular, transposition becomes

```

transpose :: [[a]] -> [[a]]
transpose [] = repeat []
transpose (xs : xss) = repeat (:) 'zapp' xs 'zapp' transpose xss

```

² This function differs from the one in the standard library in its treatment of ragged lists.

Except for the noise of the repeats and zapps, this definition is in a fairly standard applicative style, even though we are working with vectors.

1.3 Evaluating expressions

When implementing an evaluator for a language of expressions, it is customary to pass around an environment, giving values to the free variables. Here is a very simple example:

```

data Exp v = Var v
           | Val Int
           | Add (Exp v) (Exp v)

eval :: Exp v → Env v → Int
eval (Var x)  γ = fetch x γ
eval (Val i)  γ = i
eval (Add p q) γ = eval p γ + eval q γ

```

where $\text{Env } v$ is some notion of environment and $\text{fetch } x$ projects the value for the variable x .

We can eliminate the clutter of the explicitly threaded environment with a little help from some very old friends, designed for this purpose:

```

eval :: Exp v → Env v → Int
eval (Var x)  = fetch x
eval (Val i)  = K i
eval (Add p q) = K (+) ‘S’ eval p ‘S’ eval q

```

where

```

K :: a → env → a      S :: (env → a → b) → (env → a) → (env → b)
K x γ = x              S ef es γ = (ef γ) (es γ)

```

Except for the noise of the **K** and **S** combinators,³ this definition of `eval` is in a fairly standard applicative style, even though we are abstracting an environment.

2 The Applicative class

We have seen three examples of this ‘pure function applied to funny arguments’ pattern in apparently quite diverse fields – let us now abstract out what they have in common. In each example, there is a type constructor f that embeds the usual notion of value, but supports its *own peculiar way* of giving meaning to the usual applicative language – its *idiom*. We correspondingly introduce the Applicative class:

```

infixl 4 ⊗

class Applicative f where
  pure :: a → f a
  (⊗) :: f (a → b) → f a → f b

```

³ Also known as the `return` and `ap` of the environment Monad.

This class generalises **S** and **K** from threading an environment to threading an effect. (As we shall see later, these effects include, but are not limited to, monadic effects.)

We shall require the following laws for applicative functors:

identity $\text{pure id} \otimes u = u$
composition $\text{pure } (\cdot) \otimes u \otimes v \otimes w = u \otimes (v \otimes w)$
homomorphism $\text{pure } f \otimes \text{pure } x = \text{pure } (f\ x)$
interchange $u \otimes \text{pure } x = \text{pure } (\lambda f \rightarrow f\ x) \otimes u$

The idea is that `pure` embeds pure computations into the pure fragment of an effectful world – the resulting computations may thus be shunted around freely, as long as the order of the genuinely effectful computations is preserved.

One can easily check that applicative functors are indeed functors, with the following action on functions:

```
fmap :: Applicative f => (a -> b) -> f a -> f b
fmap f u = pure f \otimes u
```

Moreover, any expression built from the `Applicative` combinators can be transformed to a canonical form in which a single pure function is ‘applied’ to the effectful parts in depth-first order:

$$\text{pure } f \otimes u_1 \otimes \dots \otimes u_n$$

This canonical form captures the essence of `Applicative` programming: computations have a fixed structure, given by the pure function, and a sequence of subcomputations, given by the effectful arguments. We therefore find it convenient, at least within this article, to write this form using a special bracket notation:

$$\llbracket f\ u_1\ \dots\ u_n \rrbracket$$

This signals a shift into the idiom of an `Applicative` functor, where a pure function is applied to a sequence of effectful arguments using the appropriate \otimes . Our intention is to give an indication that effects are present, whilst retaining readability of code.

Given Haskell extended with multi-parameter type classes, enthusiasts for overloading may replace ‘ \llbracket ’ and ‘ \rrbracket ’ by appropriately defined identifiers `ll` and `lr`⁴ with appropriate parentheses.

The `IO` monad, and indeed any `Monad`, can be made `Applicative` by taking `pure = return` and $(\otimes) = \text{ap}$. (We could alternatively use the variant of `ap` that performs the computations in the opposite order, but we shall keep to the left-to-right order in this article.) Here is a direct implementation of the instance derived from the (\rightarrow) `env` monad:

```
instance Applicative ((->) env) where
  pure x = \gamma -> x           -- K
  ef \otimes ex = \gamma -> (ef \gamma) (ex \gamma) -- S
```

With these instances, `sequence` and `eval` become

⁴ *Hint*: Define an overloaded function $\text{apply } u\ v_1\ \dots\ v_n\ \mathbb{I} = u \otimes v_1 \otimes \dots \otimes v_n$.

```

sequence :: [IO a] → IO [a]
sequence [] = [[]]
sequence (c : cs) = [(:) c (sequence cs)]
eval :: Exp v → Env v → Int
eval (Var x) = fetch x
eval (Val i) = [i]
eval (Add p q) = [(+) (eval p) (eval q)]

```

If we want to do the same for our transpose example, we shall have to avoid the library's 'list of successes' (Wadler, 1985) monad and take instead an instance `Applicative []` that supports 'vectorisation', where `pure = repeat` and `(\otimes) = zapp`, yielding

```

transpose :: [[a]] → [[a]]
transpose [] = [[]]
transpose (xs : xss) = [(:) xs (transpose xss)]

```

While there is a `Monad` of *infinite* lists for which `repeat` and `zapp` are `return` and `ap`, there is no such `Monad` on *finite* lists: the associativity law necessarily fails.

3 Traversing data structures

Have you noticed that `sequence` and `transpose` now look rather alike? The details that distinguish the two programs are inferred by the compiler from their types. Both are instances of the *applicative distributor* for lists:

```

dist :: Applicative f ⇒ [f a] → f [a]
dist [] = [[]]
dist (v : vs) = [(:) v (dist vs)]

```

Distribution is often used together with 'map'. For example, given the monadic 'failure-propagation' applicative functor for `Maybe`, we can map some failure-prone operation (a function in $a \rightarrow \text{Maybe } b$) across a list of inputs in such a way that any individual failure causes failure overall.

```

flakyMap :: (a → Maybe b) → [a] → Maybe [b]
flakyMap f ss = dist (fmap f ss)

```

As you can see, `flakyMap` traverses `ss` twice – once to apply `f`, and again to collect the results. More generally, it is preferable to define this applicative mapping operation directly, with a single traversal:

```

traverse :: Applicative f ⇒ (a → f b) → [a] → f [b]
traverse f [] = [[]]
traverse f (x : xs) = [(:) (f x) (traverse f xs)]

```

This is just the way you would implement the ordinary `fmap` for lists, but with the right-hand sides wrapped in `[...]` , shifting them into the idiom. Just like `fmap`, `traverse` is a useful gadget to have for many data structures, hence we introduce

the type class `Traversable`, capturing functorial data structures through which we can thread an applicative computation:

```
class Traversable t where
  traverse :: Applicative f => (a → f b) → t a → f (t b)
  dist     :: Applicative f =>           t (f a) → f (t a)
  dist     = traverse id
```

Of course, we can recover an ordinary ‘map’ operator by taking f to be the identity – the simple applicative functor (corresponding to the identity monad) in which all computations are pure:

```
newtype Id a = An { an :: a }
```

Haskell’s **newtype** declarations allow us to shunt the syntax of types around without changing the run-time notion of value or incurring any run-time cost. The ‘labelled field’ notation defines the projection $\text{an} :: \text{Id } a \rightarrow a$ at the same time as the constructor $\text{An} :: a \rightarrow \text{Id } a$. The usual applicative functor has the usual application:

```
instance Applicative Id where
  pure      = An
  An f ⊗ An x = An (f x)
```

So, with the **newtype** signalling which Applicative functor to thread, we have

$$\text{fmap } f = \text{an} \cdot \text{traverse } (\text{An} \cdot f)$$

Meertens (1998) defined generic `dist`-like operators, families of functions of type $t (f a) \rightarrow f (t a)$ for every regular functor t (i.e. ‘ordinary’ uniform datatype constructors with one parameter, constructed by recursive sums of products). His conditions on f are satisfied by applicative functors, so the regular type constructors can all be made instances of `Traversable`. The rule-of-thumb for `traverse` is ‘like `fmap` but with $\llbracket \cdot \cdot \cdot \rrbracket$ on the right’. For example, here is the definition for trees:

```
data Tree a = Leaf | Node (Tree a) a (Tree a)
instance Traversable Tree where
  traverse f Leaf      = \llbracket Leaf \rrbracket
  traverse f (Node l x r) = \llbracket Node (traverse f l) (f x) (traverse f r) \rrbracket
```

This construction even works for non-regular types. However, not every Functor is `Traversable`. For example, the functor $(\rightarrow) \text{env}$ cannot, in general, be `Traversable`. To see why, take $\text{env} = \text{Integer}$ and try to distribute the `Maybe` functor!

Although Meertens did suggest that threading monads might always work, his primary motivation was to generalise reduction or ‘crush’ operators, such as flattening trees and summing lists. We shall turn to these in the next section.

4 Monoids are phantom Applicative functors

The data that one may sensibly accumulate have the structure of a Monoid:

class Monoid o where

```

 $\emptyset$   ::  $o$ 
 $(\oplus)$  ::  $o \rightarrow o \rightarrow o$ 

```

such that ‘ \oplus ’ is an associative operation with identity \emptyset . The functional programming world is full of monoids – numeric types (with respect to zero and plus, or one and times), lists with respect to $[]$ and $++$, and many others – so generic technology for working with them could well prove to be useful. Fortunately, every monoid induces an applicative functor, albeit in a slightly peculiar way:

newtype Accy o a = Acc { acc :: o }

Accy o a is a *phantom type* (Leijen & Meijer, 1999) – its representation is independent of a – but it does yield the applicative functor of accumulating computations:

instance Monoid $o \Rightarrow$ Applicative (Accy o) where

```

pure _      = Acc  $\emptyset$ 
Acc  $o_1 \otimes$  Acc  $o_2$  = Acc ( $o_1 \oplus o_2$ )

```

Now reduction or ‘crushing’ is just a special kind of traversal, in the same way as with any other applicative functor, just as Meertens suggested:

```

accumulate :: (Traversable  $t$ , Monoid  $o$ )  $\Rightarrow$  ( $a \rightarrow o$ )  $\rightarrow t$   $a \rightarrow o$ 
accumulate  $f$  = acc  $\cdot$  traverse (Acc  $\cdot$   $f$ )
reduce :: (Traversable  $t$ , Monoid  $o$ )  $\Rightarrow$   $t$   $o \rightarrow o$ 
reduce = accumulate id

```

Operations like flattening and concatenation become straightforward:

```

flatten :: Tree  $a \rightarrow [a]$           concat :: [[ $a$ ]]  $\rightarrow [a]$ 
flatten = accumulate (:[])          concat = reduce

```

We can extract even more work from instance inference if we use the type system to distinguish different monoids available for a given datatype. Here, we use the disjunctive structure of `Bool` to test for the presence of an element satisfying a given predicate:

newtype Mighty = Might { might :: Bool }

instance Monoid Mighty where

```

 $\emptyset$           = Might False
Might  $x \oplus$  Might  $y$  = Might ( $x \vee y$ )
any :: Traversable  $t \Rightarrow$  ( $a \rightarrow$  Bool)  $\rightarrow t$   $a \rightarrow$  Bool
any  $p$  = might  $\cdot$  accumulate (Might  $\cdot$   $p$ )

```

Now `any` (\equiv) behaves just as the `elem` function for lists, but it can also tell whether a variable from v occurs free in an `Exp v`. Of course, `Bool` also has a conjunctive `Musty` structure, which is just as easy to exploit.

5 Applicative versus Monad?

We have seen that every Monad can be made Applicative via `return` and `ap`. Indeed, two of our three introductory examples of applicative functors involved the IO monad and the environment monad $(\rightarrow) \text{ env}$. However the Applicative structure we defined on lists is not monadic, and nor is `Accy o` (unless `o` is the trivial 1-point monoid): `return` can deliver \emptyset , but if you try to define

$$(\gg) :: \text{Accy } o \ a \rightarrow (a \rightarrow \text{Accy } o \ b) \rightarrow \text{Accy } o \ b$$

you will find it tricky to extract an `a` from the first argument to supply to the second – all you get is an `o`. The \otimes for `Accy o` is not the `ap` of a monad.

So now we know: there are strictly more Applicative functors than Monads. Should we just throw the Monad class away and use Applicative instead? Of course not! The reason there are fewer monads is just that the Monad structure is more powerful. Intuitively, the $(\gg) :: m \ a \rightarrow (a \rightarrow m \ b) \rightarrow m \ b$ of some Monad `m` allows the value returned by one computation to influence the choice of another, whereas \otimes keeps the structure of a computation fixed, just sequencing the effects. For example, one may write

```
miffy :: Monad m => m Bool -> m a -> m a -> m a
miffy mb mt me = do
  b <- mb
  if b then mt else me
```

so that the value of `mb` will choose between the *computations* `mt` and `me`, performing only one, whilst

```
iffy :: Applicative f => f Bool -> f a -> f a -> f a
iffy fb ft fe = [cond fb ft fe] where
  cond b t e = if b then t else e
```

performs the effects of all three computations, using the value of `fb` to choose only between the *values* of `ft` and `fe`. This can be a bad thing: for example,

```
iffy [True] [t] Nothing = Nothing
```

because the ‘else’ computation fails, even though its value is not needed, but

```
miffy [True] [t] Nothing = [t]
```

However, if you are working with `miffy`, it is probably because the condition is an expression with effectful components, so the idiom syntax provides quite a convenient extension to the monadic toolkit:

```
miffy [(<=) getSpeed getSpeedLimit] stepOnIt checkMirror
```

The moral is this: if you have got an Applicative functor, that is good; if you have also got a Monad, that is even better! And the dual of the moral is this: if you need a Monad, that is fine; if you need only an Applicative functor, that is even better!

One situation in which the full power of monads is not always required is parsing, for which Røjemo (1995) proposed an interface including the equivalents of `pure` and `'⊗'` as an alternative to monadic parsers (Hutton & Meijer, 1998). Several ingenious non-monadic implementations have been developed by Swierstra and colleagues (Swierstra & Duponcheel, 1996; Baars *et al.*, 2004). Because the structure of these parsers is independent of the results of parsing, these implementations are able to analyse the grammar lazily and generate very efficient parsers.

5.1 Composing applicative functors

The weakness of applicative functors makes them easier to construct from components. In particular, although only certain pairs of monads are composable (Barr & Wells, 1984), the `Applicative` class is *closed under composition*,

newtype $(f \circ g) a = \text{Comp}\{\text{comp} :: (f (g a))\}$

just by lifting the inner `Applicative` operations to the outer layer:

instance $(\text{Applicative } f, \text{Applicative } g) \Rightarrow \text{Applicative } (f \circ g)$ **where**
`pure x` = `Comp [(pure x)]`
`Comp fs ⊗ Comp xs` = `Comp [(⊗) fs xs]`

As a consequence, the composition of two monads may not be a monad, but it is certainly applicative. For example, both `Maybe ∘ IO` and `IO ∘ Maybe` are applicative: `IO ∘ Maybe` is an applicative functor in which computations have a notion of ‘failure’ and ‘prioritised choice’, even if their ‘real world’ side effects cannot be undone. Note that `IO` and `Maybe` may also be composed as monads (though not *vice versa*), but the applicative functor determined by the composed monad differs from the composed applicative functor: the binding power of the monad allows the second `IO` action to be *aborted* if the first returns a failure.

We began this section by observing that `Accy o` is not a monad. However, given `Monoid o`, it can be defined as the composition of two applicative functors derived from monads – which two, we leave as an exercise.

5.2 Accumulating exceptions

The following type may be used to model exceptions:

data `Except err a = OK a | Failed err`

A `Monad` instance for this type must abort the computation on the first error, as there is then no value to pass to the second argument of `'>>='`. However, with the `Applicative` interface we can continue in the face of errors:

```

instance Monoid err  $\Rightarrow$  Applicative (Except err) where
  pure                = OK
  OK f  $\otimes$  OK x      = OK (f x)
  OK f  $\otimes$  Failed err = Failed err
  Failed err  $\otimes$  OK x  = Failed err
  Failed err1  $\otimes$  Failed err2 = Failed (err1  $\oplus$  err2)

```

This could be used to collect errors by using the list monoid (Coutts, 2002), or to summarise them in some way.

6 Applicative functors and Arrows

To handle situations in which monads were inapplicable, Hughes (2000) defined an interface that he called *arrows*, defined by the following class with nine axioms:

```

class Arrow ( $\rightsquigarrow$ ) where
  arr :: (a  $\rightarrow$  b)  $\rightarrow$  (a  $\rightsquigarrow$  b)
  ( $\gg$ ) :: (a  $\rightsquigarrow$  b)  $\rightarrow$  (b  $\rightsquigarrow$  c)  $\rightarrow$  (a  $\rightsquigarrow$  c)
  first :: (a  $\rightsquigarrow$  b)  $\rightarrow$  ((a, c)  $\rightsquigarrow$  (b, c))

```

Examples include ordinary ' \rightarrow ', Kleisli arrows of monads and comonads, and stream processors. Equivalent structures called *Freyd-categories* had been independently developed to structure denotational semantics (Power & Robinson, 1997).

There are similarities to the Applicative interface, with *arr* generalising *pure*. The ' \gg ' operation arranges for the result of the first computation to be fed to the second; as with ' \otimes ', the plumbing is independent of the data flowing through it.

By fixing the first argument of an arrow type, we obtain an applicative functor, generalising the environment functor we saw earlier:

```

newtype EnvArrow ( $\rightsquigarrow$ ) env a = Env (env  $\rightsquigarrow$  a)
instance Arrow ( $\rightsquigarrow$ )  $\Rightarrow$  Applicative (EnvArrow ( $\rightsquigarrow$ ) env) where
  pure x          = Env (arr (const x))
  Env u  $\otimes$  Env v = Env (u  $\Delta$  v  $\gg$  arr ( $\lambda$ (f, x)  $\rightarrow$  f x))
  where u  $\Delta$  v = arr dup  $\gg$  first u  $\gg$  arr swap  $\gg$  first v  $\gg$  arr swap
  dup a = (a, a)
  swap (a, b) = (b, a)

```

In the other direction, each applicative functor defines an arrow constructor that adds static information to an existing arrow:

```

newtype StaticArrow f ( $\rightsquigarrow$ ) a b = Static (f (a  $\rightsquigarrow$  b))
instance (Applicative f, Arrow ( $\rightsquigarrow$ ))  $\Rightarrow$  Arrow (StaticArrow f ( $\rightsquigarrow$ )) where
  arr f          = Static  $\llbracket$  (arr f)  $\rrbracket$ 
  Static f  $\gg$  Static g = Static  $\llbracket$  ( $\gg$ ) f g  $\rrbracket$ 
  first (Static f)    = Static  $\llbracket$  first f  $\rrbracket$ 

```

To date, most applications of the extra generality provided by arrows over monads have been either various forms of process, in which components may consume multiple inputs, or computing static properties of components. Indeed one of Hughes's

motivations was the parsers of Swierstra and Duponcheel (1996). It may turn out that applicative functors are more convenient for applications of the second class.

7 Applicative functors, categorically

The Applicative class features the asymmetrical operation ‘ \otimes ’, but there is an equivalent symmetrical definition:

```
class Functor  $f \Rightarrow$  Monoidal  $f$  where
  unit ::  $f ()$ 
  ( $\star$ ) ::  $f a \rightarrow f b \rightarrow f (a, b)$ 
```

These operations are clearly definable for any Applicative functor:

```
unit    :: Applicative  $f \Rightarrow f ()$ 
unit    =  $\llbracket () \rrbracket$ 
( $\star$ )   :: Applicative  $f \Rightarrow f a \rightarrow f b \rightarrow f (a, b)$ 
 $f a \star f b = \llbracket (,) f a f b \rrbracket$ 
```

Moreover, we can recover the Applicative interface from Monoidal as follows:

```
pure     :: Monoidal  $f \Rightarrow a \rightarrow f a$ 
pure  $x = \text{fmap } (\lambda\_ \rightarrow x) \text{ unit}$ 
( $\otimes$ )   :: Monoidal  $f \Rightarrow f (a \rightarrow b) \rightarrow f a \rightarrow f b$ 
 $mf \otimes mx = \text{fmap } (\lambda(f, x) \rightarrow f x) (mf \star mx)$ 
```

The laws of Applicative given in Section 2 are equivalent to the usual Functor laws, plus the following laws of Monoidal:

```
naturality of  $\star$     $\text{fmap } (f \times g) (u \star v) = \text{fmap } f u \star \text{fmap } g v$ 
left identity      $\text{fmap } \text{snd } (\text{unit} \star v) = v$ 
right identity     $\text{fmap } \text{fst } (u \star \text{unit}) = u$ 
associativity      $\text{fmap } \text{assoc } (u \star (v \star w)) = (u \star v) \star w$ 
```

for the functions

```
( $\times$ ) ::  $(a \rightarrow b) \rightarrow (c \rightarrow d) \rightarrow (a, c) \rightarrow (b, d)$ 
( $f \times g$ )  $(x, y) = (f x, g y)$ 
assoc ::  $(a, (b, c)) \rightarrow ((a, b), c)$ 
assoc  $(a, (b, c)) = ((a, b), c)$ 
```

Fans of category theory will recognise the above laws as the properties of a *lax monoidal functor* for the monoidal structure given by products. However the functor composition and naturality equations are actually stronger than their categorical counterparts. This is because we are working in a higher-order language, in which function expressions may include variables from the environment, as in the above definition of pure for Monoidal functors. In the first-order language of category theory, such data flow must be explicitly plumbed using functors with *tensorial strength*, an arrow:

$$t_{AB} : A \times F B \longrightarrow F (A \times B)$$

satisfying standard equations. The natural transformation m corresponding to ‘ \star ’ must also respect the strength:

$$\begin{array}{ccc}
 (A \times B) \times (F C \times F D) & \cong & (A \times F C) \times (B \times F D) \\
 (A \times B) \times m \downarrow & & \downarrow t \times t \\
 (A \times B) \times F(C \times D) & & F(A \times C) \times F(B \times D) \\
 t \downarrow & & \downarrow m \\
 F((A \times B) \times (C \times D)) & \cong & F((A \times C) \times (B \times D))
 \end{array}$$

Note that B and FC swap places in the above diagram: strong naturality implies commutativity with pure computations.

Thus in categorical terms applicative functors are *strong lax monoidal functors*. Every strong monad determines two of them, as the definition is symmetrical. The Monoidal laws and the above definition of pure imply that pure computations commute past effects:

$$\text{fmap swap (pure } x \star u) = u \star \text{pure } x$$

The proof (an exercise) makes an essential use of higher-order functions.

8 Conclusions

We have identified Applicative functors, an abstract notion of effectful computation lying between Arrow and Monad in strength. Every Monad is an Applicative functor, but significantly, the Applicative class is closed under composition, allowing computations such as accumulation in a Monoid to be characterised in this way.

Given the wide variety of Applicative functors, it becomes increasingly useful to abstract Traversable functors – container structures through which Applicative actions may be threaded. Combining these abstractions yields a small but highly generic toolkit whose power we have barely begun to explore. We use these tools by writing types that structure not merely the *storage* of data, but also the *properties* of data that we intend to exploit.

Library modules based on these classes are included in the libraries shared by all Haskell implementations.

The explosion of categorical structure in functional programming – monads, comonads, arrows and now applicative functors – should not, we suggest, be a cause for alarm. Why should we not profit from whatever structure we can sniff out, abstract and re-use? The challenge is to avoid a chaotic proliferation of peculiar and incompatible notations. If we want to rationalise the notational impact of all these structures, perhaps we should try to recycle the notation we already possess. Our $\llbracket f u_1 \dots u_n \rrbracket$ notation does minimal damage, showing when the existing syntax for applicative programming should be interpreted with an effectful twist.

Acknowledgements

We thank Thorsten Altenkirch, Duncan Coutts, Jeremy Gibbons, Peter Hancock, Simon Peyton Jones, Doaitse Swierstra and Phil Wadler for their help and encouragement.

References

- Baars, A., Löh, A. & Swierstra, S. D. (2004) Parsing permutation phrases. *Journal of Functional Programming*, **14**(6), 635–646.
- Barr, M. & Wells, C. (1984) *Toposes, Triples and Theories*. Grundlehren der Mathematischen Wissenschaften, no. 278. New York: Springer, Chap. 9.
- Coutts, D. (2002) *Arrows for Errors: Extending the Error Monad*. Unpublished presentation at the Summer School on Advanced Functional Programming.
- Fridlender, D. & Indrika, M. (2000) Do we need dependent types? *Journal of Functional Programming*, **10**(4), 409–415.
- Hughes, J. (2000) Generalising monads to arrows. *Science of Computer Programming*, **37**(1–3), 67–111.
- Hutton, G. & Meijer, E. (1998) Monadic parsing in Haskell. *Journal of Functional Programming*, **8**(4), 437–444.
- Leijen, D. & Meijer, E. (1999) Domain specific embedded compilers. *2nd Conference on Domain-Specific Languages (DSL)*. USENIX, Austin, TX, USA.
- Meertens, L. (1998) Functor pulling. *Workshop on Generic Programming (WGP'98)*. Marstrand, Sweden: Chalmers University of Technology.
- Peyton Jones, S. (ed). (2003) *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press.
- Power, J. & Robinson, E. (1997) Premonoidal categories and notions of computation. *Mathematical Structures in Computer Science*, **7**(5), 453–468.
- Röjemo, N. (1995) *Garbage Collection and Memory Efficiency in Lazy Functional Languages*. Ph.D. thesis, Chalmers University of Technology and Göteborg University.
- Swierstra, S. D. & Duponcheel, L. (1996) Deterministic, error-correcting combinator parsers. *Pages 184–207 of*: Launchbury, J., Meijer, E. & Sheard, T. (eds), *Advanced Functional Programming*. LNCS, vol. 1129. Springer.
- Wadler, P. (1985) How to replace failure by a list of successes. *Pages 113–128 of*: Jouannaud, J.-P. (ed), *Functional Programming Languages and Computer Architecture*. LNCS, vol. 201. Springer.