

Theory and Practice of Fusion

Ralf Hinze, Thomas Harper, and Daniel W. H. James

Computing Laboratory, University of Oxford
Wolfson Building, Parks Road, Oxford, OX1 3QD, England
`ralf.hinze,tom.harper,daniel.james@comlab.ox.ac.uk`

Abstract. There are a number of approaches for eliminating intermediate data structures in functional programs—this elimination is commonly known as *fusion*. Existing fusion strategies are built upon various, but related, recursion schemes, such as folds and unfolds. We use the concept of *recursive coalgebras* as a unifying theoretical and notational framework to explore the foundations of these fusion techniques. We first introduce the calculational properties of recursive coalgebras and demonstrate their use with proofs and derivations in a calculational style, then provide an overview of fusion techniques by bringing them together in this setting. We also showcase these developments with examples in Haskell.

1 Introduction

Functional programmers love modular programs. It is easy for them to create clear, concise, and reusable code by composing functions. Consider the following Haskell program as an example:

$$\begin{aligned} f &: (\text{Integer}, \text{Integer}) \rightarrow \text{Integer} \\ f &= \text{sum} \cdot \text{map sq} \cdot \text{filter odd} \cdot \text{between} \end{aligned}$$

The program takes a pair of integers representing an interval and returns the sum of the squared odd integers in this interval. We have expressed this a composition of four functions: *between* generates an enumeration between two natural numbers as a list, *filter odd* removes any even numbers, *map sq* squares the remaining (odd) numbers, and *sum* adds them together. Unfortunately, the clarity of this program comes at a cost. The constituent functions of this program communicate with each other using intermediate data structures, the production and immediate consumption of which carries an obvious performance penalty. Yet, because these definitions are recursive, eliminating the need for these transient structures is beyond the reach of a typical compiler.

Nonetheless, such a transformation is possible. We can manually construct a program that is equivalent to *f*, but without the intermediate data structures

$$\begin{aligned} f'(m, n) &= \text{go } m \\ &\text{where } \text{go } m \mid m > n \quad = 0 \\ &\quad \mid \text{otherwise} = \text{go } (m + 1) + \text{if odd } m \text{ then } \text{sq } m \text{ else } 0 \end{aligned}$$

This new program has lost the desirable qualities of the original—our concise, modular and declarative code has been hammered into a long, opaque and specialised function. In doing so, however, we have accomplished our goal of removing the intermediate data structures by transforming the numerous recursive traversals into a single one. This process is called *fusion*.

Fusing programs by hand quickly becomes infeasible for those of non-trivial length. Furthermore, it can be difficult to manually pinpoint all the opportunities for fusion. Instead, such a transformation should be performed automatically. Difficulties arise, however, in automatically fusing functions defined using general recursion. Specifically, such transformations often have proof obligations that cannot be discharged by the compiler.

One remedy is to standardise the way data structures are produced or consumed by encapsulating the recursion scheme in a higher-order function. The arguments to these functions are the non-recursive ‘steps’. Simple syntactic transformations can then fuse many recursive traversals into a single one, and then non-recursive steps can be optimised using conventional methods. This approach is known as *shortcut fusion*. Different incarnations of this technique utilise different recursion schemes, e.g. *folds* for consumers or *unfolds* for producers. The steps of such a scheme are known as algebras or coalgebras, respectively.

The implementation of these fusion techniques is usually described syntactically, by giving a definition of the production and consumption combinators and accompanying rewrite rules. This alone does not really explain the underlying fusion mechanism. Furthermore, it is difficult to construct correctness proofs, or relate various fusion approaches to one another, despite the fact that such close relations exist. In this paper, we move fusion to a clearer setting, where the syntactic details of fusion fall away.

Category theory provides the tools we need to tackle the semantics of the recursion schemes. While some fusion techniques have been individually given this treatment before, our focus is to bring them all under one roof. In this paper, we propose using *recursive coalgebras* as that roof. We will show how recursive coalgebras enable us to explain the fusion rules underlying the various fusion techniques and give short, simple proofs of correctness.

Some proofs and source code examples have been elided from this paper. Any reader who desires further material should consult the associated technical report [10].

2 Background: Algebras and Coalgebras

The category theory concept of an *initial algebra* is key in the theory of functional programming languages, specifically for giving a semantics to recursive datatypes [15]. The remainder of this section will refresh the salient details. For the more functional-programming-minded reader, we will parallel these developments with examples in Haskell, where possible.

Let $F : \mathbb{C} \rightarrow \mathbb{C}$ be a functor. An *F-algebra* is a pair $\langle a, A \rangle$ consisting of an object $A : \mathbb{C}$ and an arrow $a : F A \rightarrow A : \mathbb{C}$. An *F-algebra homomorphism* between

algebras $\langle a, A \rangle$ and $\langle b, B \rangle$ is an arrow $h : A \rightarrow B : \mathbb{C}$ such that $h \cdot a = b \cdot F h$. The fact that functors preserve identity and composition entails that identity is a homomorphism and that homomorphisms compose. Consequently, F-algebras and their homomorphisms form a category, called $\mathbf{Alg}(F)$. We abbreviate a homomorphism, $h : \langle a, A \rangle \rightarrow \langle b, B \rangle : \mathbf{Alg}(F)$, by $h : a \rightarrow b : \mathbf{Alg}(F)$ if the objects are obvious from the context, or simply by $h : a \rightarrow b$ if the functor F is also obvious.

In Haskell, we model a functor as a datatype whose constructors describe its action on data (i.e. objects). Its action on arrows is defined by making that datatype an instance of the *Functor* typeclass

```
class Functor f where
  fmap : (a → b) → (f a → f b) .
```

We can then simply treat the concept of an F-algebra as a function, where F is a datatype that is an instance of the *Functor* class. The F-algebra $\langle a, A \rangle$ is simply a function $a : F A \rightarrow A$. An F-algebra homomorphism between a and $b : F B \rightarrow B$ is a function $h : A \rightarrow B$ that satisfies the side condition $h \cdot a = b \cdot fmap h$. This property cannot, however, be deduced from the type and must be checked by the programmer.

If the category $\mathbf{Alg}(F)$ has an initial object, then we call it $\langle in, \mu F \rangle$. Initiality means that there is a *unique* arrow from $\langle in, \mu F \rangle$ to any other F-algebra $\langle a, A \rangle$. This arrow, called *fold*, is written $(\!|a|) : in \rightarrow a$. We construct elements of μF using *in* and deconstruct them using $(\!|a|)$. We can think of $(\!|a|) : in \rightarrow a$ as replacing constructors by functions, represented by the algebras *in* and *a*, respectively. Initiality is captured by the *uniqueness property* of folds:¹

$$h = (\!|a|) \iff h : in \rightarrow a \iff h \cdot in = a \cdot F h . \quad (1)$$

It is important to note that we have omitted the quantification of the names that appear in property (1); we have done so for presentational succinctness and we will continue in this manner. In this case we will spell out implicit quantification: the uniqueness property holds for all functors F , where the category $\mathbf{Alg}(F)$ has an initial object named $\langle in, \mu F \rangle$, and for all F-algebras $\langle a, A \rangle$, and for all F-algebra homomorphisms $h : in \rightarrow a$.

This property provides us with a general definition of $(\!|-|)$ in Haskell. First, we define the μ datatype, which takes a functor to its least fixed point:

```
data  $\mu f = in \{ out : f (\mu f) \}$ 
```

The constructor *in* allows us to construct a structure of type μf out of something of type $f (\mu f)$ and *out* deconstructs it. The relationship between *in* and *out* in our setting is discussed further in Sections 3 and 4.

¹ The formula $P \iff Q \iff R$ has to be read *conjunctively* as $P \iff Q \wedge Q \iff R$. Likewise, $P \iff Q \iff R$ is shorthand for $P \iff Q \wedge Q \iff R$.

We can define $\langle - \rangle$ as a higher-order function that takes an algebra $f a \rightarrow a$ and returns a function $\mu f \rightarrow a$, according to the uniqueness property:

$$\begin{aligned} \langle - \rangle &: (\text{Functor } f) \Rightarrow (f a \rightarrow a) \rightarrow (\mu f \rightarrow a) \\ \langle a \rangle &= a \cdot \text{fmap } \langle a \rangle \cdot \text{out} \end{aligned}$$

By allowing us to substitute $\langle a \rangle$ for h , we see that the uniqueness property provides us with a definition of $\langle - \rangle$ that recursively replaces occurrences of in by some algebra a . The placement of the recursive call for a given structure μf is determined by the definition fmap .

We will not employ the uniqueness property in an example proof just yet. In fact the uniqueness property is rarely used in its raw form; instead, there are a number of specific forms that we will introduce now.

If we set h to the identity id and a to the initial algebra in , then we obtain the *reflection law*: $\langle \text{in} \rangle = \text{id}$. If we substitute the left-hand side into the right-hand side, then we obtain the *computation law*: $\langle a \rangle : \text{in} \rightarrow a$, or expressed in terms of the base category, $\langle a \rangle \cdot \text{in} = a \cdot \mathbf{F} \langle a \rangle$.

The most important consequence of the uniqueness property is the *fusion law* for fusing an arrow with a fold to form a new fold.

$$h \cdot \langle a \rangle = \langle b \rangle \quad \Leftarrow \quad h : a \rightarrow b \quad \iff \quad h \cdot a = b \cdot \mathbf{F} h \quad (2)$$

As its name would suggest, the fusion law is closely related to the program transformation techniques described in the introduction. It allows a fold to absorb a function on its left, thereby producing a single fold. The law also shows the difficulty of mechanising this process; in order to produce the fused program, we must invent a new algebra b that satisfies the precondition.

Folds enjoy an additional fusion law. Whereas fusion allows us to absorb an additional function on the left, the *functor fusion law* allows us to absorb a function on the right. In order to formulate it, we have to turn μ into a higher-order functor of type $\mathbb{C}^{\mathbb{C}} \rightarrow \mathbb{C}$. The object part of this functor maps a functor to its initial algebra. (This is only well-defined for functors that have an initial algebra.) The arrow part maps a natural transformation $\alpha : \mathbf{F} \rightarrow \mathbf{G}$ to an arrow $\mu\alpha : \mu\mathbf{F} \rightarrow \mu\mathbf{G} : \mathbb{C}$. It is defined as $\mu\alpha = \langle \text{in} \cdot \alpha(\mu\mathbf{F}) \rangle$. To reduce clutter, we will henceforth omit the argument of the natural transformation α . From these definitions we obtain the *functor fusion law* (we have annotated $\langle - \rangle$ with the underlying functors):

$$\langle b \rangle_{\mathbf{G}} \cdot \mu\alpha = \langle b \cdot \alpha \rangle_{\mathbf{F}} . \quad (3)$$

It states that a fold after a map can be fused into a single fold — the map $\mu\alpha$ can be seen as a “base changer”.

We can also provide a Haskell definition of μ as a functor. The action on data is given by its datatype declaration. The action on functions is given by:

$$\begin{aligned} \mu- &: (\text{Functor } f) \Rightarrow (\forall a . f a \rightarrow g a) \rightarrow (\mu f \rightarrow \mu g) \\ \mu\alpha &= \langle \text{in} \cdot \alpha \rangle \end{aligned}$$

Note that we use a rank-2 polymorphic type to express the idea that μ maps natural transformation from f to g to a function between their fixpoints.

Finally, the initial algebra μF is the least fixed point of F — this is known as Lambek’s Lemma [13]. One direction of the isomorphism $F(\mu F) \cong \mu F$ is given by in , its inverse is $in^\circ = \langle F in \rangle$. Lambek’s Lemma is the key to giving a semantics to recursively defined datatypes. To illustrate this, the recursive definition of lists of natural numbers

data List = Nil | Cons (ℕ, List)

implicitly defines an underlying functor $LX = 1 + \mathbb{N} \times X$, the so-called base functor of List. (This notation is a categorical rendering of sum-of-products algebraic datatypes, and defines a functor L with an argument X , where 1 denotes the terminal object of the underlying category.) Since the initial object μL satisfies the equation $X \cong LX$, we can use it to assign meaning to the recursive datatype definition. (As an aside, the fold of the List datatype is a specialisation of Haskell’s library function *foldr*.)

The initial F -algebra is the least solution of the equation $X \cong FX$. If we dualise the development above, we obtain another canonical solution, namely the greatest one. In category theory, dualisation is denoted by the prefix “co-”.

An F -coalgebra is a pair $\langle C, c \rangle$ consisting of an object $C : \mathbb{C}$ and an arrow $c : C \rightarrow FC : \mathbb{C}$. An F -coalgebra homomorphism between coalgebras $\langle C, c \rangle$ and $\langle D, d \rangle$ is an arrow $h : C \rightarrow D : \mathbb{C}$ such that $Fh \cdot c = d \cdot h$. Coalgebras and coalgebra homomorphisms also form a category, called **Coalg**(F). The dual of the initial algebra is the final coalgebra, whose carrier νF is the greatest fixed point of F . Finality means that, for any other coalgebra, there is a unique arrow from it to the final coalgebra. Whereas a fold consumes a data structure, an unfold produces some data structure from a given seed.

Unfortunately, least and greatest fixed points are different beasts in general. In the category **Set** of sets and total functions, μL is the set of finite lists, whereas νL also contains infinite lists. This means that folds and unfolds are incompatible, in general. In the following section, we will focus on a restricted species of coalgebras, enabling us to work with folds and unfolds under the same roof.

3 Recursive Coalgebras

In this section we will introduce *recursive* coalgebras. We follow the work of Capretta et al. [2], who motivate the use of hylomorphisms based on recursive coalgebras as a structured recursion scheme. We shall continue to parallel our developments with examples in Haskell.

A coalgebra $\langle C, c \rangle$ is called *recursive* if for *every* algebra $\langle a, A \rangle$ the equation in the unknown $h : A \leftarrow C$,

$$h = a \cdot Fh \cdot c, \tag{4}$$

has a *unique* solution. The equation captures *divide-and-conquer*: a problem is divided into sub-problems (c), the sub-problems are solved recursively (Fh), and finally the sub-solutions are combined into a single solution (a). The uniquely

defined function h is called a *hylomorphism* or *hylo* for short and is written $\langle a \leftarrow c \rangle_F : A \leftarrow C$. The notation is meant to suggest that h takes a coalgebra to an algebra. We omit the subscripted functor name if it is obvious from the context. Uniqueness of h is captured by the following property.

$$h = \langle a \leftarrow c \rangle \iff h = a \cdot F h \cdot c \quad (5)$$

In Haskell, $\langle - \leftarrow - \rangle$ becomes a function that takes an algebra and a recursive coalgebra (which, dual to algebras, is a function of type $c \rightarrow f c$) as arguments and returns resulting hylo according to the definition in the universal property:

$$\begin{aligned} \langle - \leftarrow - \rangle : (\text{Functor } f) \Rightarrow (f a \rightarrow a) \rightarrow (c \rightarrow f c) \rightarrow (c \rightarrow a) \\ \langle a \leftarrow c \rangle = a \cdot \text{fmap } \langle a \leftarrow c \rangle \cdot c . \end{aligned}$$

This function takes an algebra and a recursive coalgebra, yielding a hylo. Note that the type of this function does not guarantee that c is a *recursive* coalgebra and therefore does not guarantee that the resulting hylo has a unique solution; the programmer needs to discharge this obligation by some other means.

The category of recursive coalgebras and coalgebra homomorphisms forms a full subcategory of $\mathbf{Coalg}(F)$, called $\mathbf{Rec}(F)$. If the latter category has a final object $\langle F, \text{out} \rangle$, then there is a unique arrow from any other *recursive* coalgebra $\langle C, c \rangle$ to $\langle F, \text{out} \rangle$. This arrow, called *unfold*, is written $\llbracket c \rrbracket : c \rightarrow \text{out}$. Finality is captured by the following uniqueness property.

$$h = \llbracket c \rrbracket \iff h : c \rightarrow \text{out} \iff F h \cdot c = \text{out} \cdot h \quad (6)$$

This is the usual property of unfolds, except that we are working in the category $\mathbf{Rec}(F)$, *not* $\mathbf{Coalg}(F)$. As with folds, we can draw out a Haskell definition of unfolds from the uniqueness property:

$$\begin{aligned} \llbracket - \rrbracket : (\text{Functor } f) \Rightarrow (c \rightarrow f c) \rightarrow (c \rightarrow \mu f) \\ \llbracket c \rrbracket = \text{in} \cdot \text{fmap } \llbracket c \rrbracket \cdot c . \end{aligned}$$

In contrast to folds, we are *creating* a structure of type μf from a seed value. The recursion, similarly, is determined by the form of the underlying functor f through the use of *fmap*. The uniqueness property for unfolds, like the one for folds, implies the *reflection law*, $\llbracket \text{out} \rrbracket = \text{id}$, the *computation law*, $F \llbracket c \rrbracket \cdot c = \text{out} \cdot \llbracket c \rrbracket$, and the *fusion law*:

$$\llbracket c \rrbracket = \llbracket d \rrbracket \cdot h \iff h : c \rightarrow d \iff F h \cdot c = d \cdot h . \quad (7)$$

The definition of a hylomorphism does not assume that the initial F-algebra exists. The powerset functor, for instance, admits no fixed points. However, if the initial algebra exists, then it coincides with the final recursive coalgebra and, furthermore, folds and unfolds emerge as special cases of hylos. We can state this more formally:

Theorem 1. *Initial F-algebras and final recursive F-coalgebras coincide: (1) If $\langle C, \text{out} \rangle$ is the final recursive F-coalgebra, then $\langle \text{out}^\circ, C \rangle$ is the initial F-algebra. Furthermore, $\langle a \rangle = \langle a \leftarrow \text{out} \rangle$. (2) If $\langle \text{in}, A \rangle$ is the initial F-algebra, then $\langle A, \text{in}^\circ \rangle$ is the final recursive F-coalgebra. Furthermore, $\llbracket c \rrbracket = \langle \text{in} \leftarrow c \rangle$.*

Theorem 1 allows us to treat folds and unfolds in the same setting—note that an unfold produces an element of an initial algebra! An alternative is to work in a setting where μF and νF coincide; an *algebraically compact* category is such a setting [8]. Haskell’s ambient category \mathbf{Cpo}_\perp serves as the standard example. This is the usual approach [7], however, the downside is that the hylo equation (4) only has a canonical, least solution, not a unique solution, so (5) does not hold.

4 Calculational Properties

In this section we will cover the calculational properties of our hylomorphisms. In a similar fashion to folds and unfolds, hylomorphisms have an identity law and a computation law, and they follow similarly from the uniqueness property (5).

Identity law Setting $h := id$, we obtain the *identity law*

$$(a \leftarrow c) = id \iff a \cdot c = id . \quad (8)$$

Computation law Substituting the left-hand side into the right-hand side gives the *computation law*:

$$(a \leftarrow c) = a \cdot F(a \leftarrow c) \cdot c . \quad (9)$$

For hylomorphisms, we have *three* fusion laws: algebra fusion, coalgebra fusion, and composition.

Algebra fusion An algebra homomorphism after a hylo can be fused to form a single hylo.²

$$h \cdot (a \leftarrow c) = (b \leftarrow c) \iff h : a \rightarrow b \iff h \cdot a = b \cdot F h \quad (10)$$

For the proof we appeal to the uniqueness property and show that $h \cdot (a \leftarrow c)$ satisfies the recursion equation of $(b \leftarrow c)$. The obligation is discharged as follows:

$$\begin{aligned} & h \cdot (a \leftarrow c) \\ = & \{ \text{hylo computation (9)} \} \\ & h \cdot a \cdot F(a \leftarrow c) \cdot c \\ = & \{ \text{assumption: } h : a \rightarrow b \} \\ & b \cdot F h \cdot F(a \leftarrow c) \cdot c \\ = & \{ F \text{ functor} \} \\ & b \cdot F(h \cdot (a \leftarrow c)) \cdot c . \end{aligned}$$

Coalgebra fusion Dually, we can fuse a coalgebra homomorphism before a hylo to form a single hylo.

$$(a \leftarrow c) = (a \leftarrow d) \cdot h \iff h : c \rightarrow d \iff F h \cdot c = d \cdot h \quad (11)$$

² Note that h appears as both an algebra homomorphism in $\mathbf{Alg}(F)$ and as the underlying arrow in the underlying category.

Like the law, the proof is the dual of that for algebra fusion.

Composition law A composition of hylos can be merged into a single one if the coalgebra of the hylo on the left inverts the algebra of the right hylo.

$$\langle a \leftarrow c \rangle \cdot \langle b \leftarrow d \rangle = \langle a \leftarrow d \rangle \quad \Leftarrow \quad c \cdot b = id \quad (12)$$

Composition is, in fact, a simple consequence of algebra fusion as the hylomorphism $\langle a \leftarrow c \rangle : b \rightarrow a$ is simultaneously an F -algebra homomorphism.

$$\begin{aligned} & \langle a \leftarrow c \rangle \cdot b \\ = & \{ \text{hylo computation (9)} \} \\ & a \cdot F \langle a \leftarrow c \rangle \cdot c \cdot b \\ = & \{ \text{assumption: } c \cdot b = id \} \\ & a \cdot F \langle a \leftarrow c \rangle \end{aligned}$$

Alternatively, we can derive the composition law from coalgebra fusion by showing that $\langle b \leftarrow d \rangle : d \rightarrow c$ is an F -coalgebra homomorphism. The composition law, together with the next law, generalises the functor fusion law of folds.

Hylo shift law or base change law If we have a natural transformation $\alpha : G \rightarrow F$, then

$$\langle a \cdot \alpha A \leftarrow c \rangle_G = \langle a \leftarrow \alpha C \cdot c \rangle_F . \quad (13)$$

In fact, the statement can be strengthened: if c is recursive, then $\alpha C \cdot c$ is recursive, as well.

$$\begin{aligned} & h = a \cdot F h \cdot \alpha C \cdot c \\ \iff & \{ \alpha \text{ natural} \} \\ & h = a \cdot \alpha A \cdot G h \cdot c \\ \iff & \{ \text{uniqueness property of hylos (5)} \} \\ & h = \langle a \cdot \alpha A \leftarrow c \rangle_G \end{aligned}$$

It is worth pointing out that the laws stated thus far are independent of the existence of initial algebras. Only the following law makes this assumption.

Fold/unfold law A fold after an unfold is a hylo.

$$\langle a \rangle \cdot \llbracket c \rrbracket = \langle a \leftarrow c \rangle \quad (14)$$

From left to right we are performing fusion and thus deforesting an intermediate data structure. From right to left we are turning a control structure into a data structure. The *fold/unfold law* is a direct consequence of Theorem 1 and any of the fusion laws.

5 Fusion

In the previous sections we have introduced the fusion laws that we will now use to help us explain a collection of specific fusion techniques. We collectively

brand these techniques *shortcut fusion*, as they share the common characteristic of standardising the way data structures are recursively consumed and produced. Where shortcut fusion techniques differ is in their choice of recursion scheme. By using recursive coalgebras, we can clearly lay out and compare these approaches within the *same* framework.³ This allows us to examine the relationships among these fusion approaches which are not readily apparent when examining their individual implementations.

5.1 Warm-up: Type Functors

We have seen in §2 that μ is a functor, whose action on arrows is defined $\mu\alpha = (in \cdot \alpha)$. Using Theorem 1 and the hylo shift law (13) we can actually express $\mu\alpha$ as a fold, an unfold or a hylo.

$$\mu\alpha = (in \cdot \alpha) = (in \cdot \alpha \leftarrow out) = (in \leftarrow \alpha \cdot out) = \llbracket \alpha \cdot out \rrbracket .$$

In §5.5 we shall see a key use of μ for stream fusion. For now, let us show a use of μ with the base functor of parametrized List

data L $a\ b = Nil \mid Cons\ (a, b)$.

This is a higher-order functor of type $L : \mathbb{C} \rightarrow \mathbb{C}^{\mathbb{C}}$ that takes objects to functors and arrows to natural transformations. In Haskell, we can make this datatype an instance of the *Functor* class:

```
instance Functor (L a) where
  fmap f Nil           = Nil
  fmap f (Cons (a, b)) = Cons (a, f b) .
```

We define this instance for the functor obtained by applying L to some type a . Haskell allows us to define this polymorphically for all a . The list datatype defined in terms of its base functor is $List\ A = \mu(L\ A)$. The parametric type List is itself a functor, a so-called type functor, whose action on arrows is Haskell's *map* function, defined in this setting by $List\ f = \mu(L\ f)$. Note that μ expects a natural transformation and that L delivers one.

5.2 Generalised *foldr/build* Fusion

We now move on to the main target of our new setting: shortcut fusion. The original shortcut fusion technique is a fold-centric approach called *foldr/build* fusion [9]. As its name would suggest, its original intention was to provide fusion for list functions written in terms of *foldr* and an additional combinator *build*. In this section, we will explore the foundations of this technique.

The mother of all fusion rules is algebra fusion (10). It allows us to fuse a hylo followed by an algebra homomorphism into a single hylo. It is similar to

³ Previously these recursion schemes were only compatible for analysis by restricting the working category to one that is algebraically compact, such as \mathbf{Cpo}_{\perp} .

fold fusion in the sense that to use this law, we must construct a new algebra that satisfies a pre-condition. To illustrate this, the pipeline $sum \cdot filter \cdot odd$ can be expressed as a composition of two folds: $\langle \mathfrak{s} \rangle \cdot \langle \mathfrak{f} \rangle$. The algebras \mathfrak{s} and \mathfrak{f} are given by

$$\begin{array}{ll} \mathfrak{f} : \mathbb{L}\mathbb{N}(\mu(\mathbb{L}\mathbb{N})) \rightarrow \mu(\mathbb{L}\mathbb{N}) & \mathfrak{s} : \mathbb{L}\mathbb{N}\mathbb{N} \rightarrow \mathbb{N} \\ \mathfrak{f} Nil = in Nil & \mathfrak{s} Nil = 0 \\ \mathfrak{f}(Cons(x, y)) = \mathbf{if} \textit{ odd } x & \mathfrak{s}(Cons(x, y)) = x + y . \\ \quad \mathbf{then} \textit{ in } (Cons(x, y)) & \\ \quad \mathbf{else} \textit{ y} & \end{array}$$

To be able to apply algebra fusion (10), we have to show that $\langle \mathfrak{s} \rangle$ is an algebra homomorphism from \mathfrak{f} to some unknown algebra \mathfrak{sf} . By hand, it is not hard to derive \mathfrak{sf} so that $\langle \mathfrak{s} \rangle \cdot \mathfrak{f} = \mathfrak{sf} \cdot F \langle \mathfrak{s} \rangle$.

$$\begin{array}{ll} \mathfrak{sf} : \mathbb{L}\mathbb{N}\mathbb{N} \rightarrow \mathbb{N} & \\ \mathfrak{sf} Nil = \mathfrak{s} Nil & \\ \mathfrak{sf}(Cons(x, y)) = \mathbf{if} \textit{ odd } x & \mathbf{then} \mathfrak{s}(Cons(x, y)) \mathbf{else} \textit{ y} \end{array}$$

Since $\langle \mathfrak{s} \rangle$ replaces in by \mathfrak{s} , we simply have to replace the occurrences of in in \mathfrak{f} by \mathfrak{s} . While this is an easy task to perform by hand, it is potentially difficult to mechanise as it requires analysis of the body of \mathfrak{f} ; within it, the constructor in could easily have any name and conversely any function could be named in . Also, \mathfrak{f} could contain unrelated occurrences of in . This transformation is therefore not purely syntactic, but also involves some further analysis of the source program; this is not an approach we wish to pursue.

The central idea of *foldr/build* fusion is to expose in so that replacing it by the algebra a is simple to implement. Consider fold fusion (2) again.

$$h \cdot \langle a \rangle = \langle b \rangle \quad \Leftarrow \quad h : a \rightarrow b$$

A fold $\langle - \rangle$ is a transformation that takes an algebra to a homomorphism. Assume that we have another such transformation, say, β that satisfies

$$h \cdot \beta a = \beta b \quad \Leftarrow \quad h : a \rightarrow b . \tag{15}$$

The generalisation of *foldr/build* from lists to arbitrary datatypes, the so-called *acid rain rule* [19], is then

$$\langle a \rangle \cdot \beta in = \beta a . \tag{16}$$

Using β we expose in so that replacing in by a is achieved through a simple function application. Instead of building a structure and then folding over it, we eliminate the in and pass a directly to β . The proof of correctness is painless.

$$\begin{array}{l} \langle a \rangle \cdot \beta in = \beta a \\ \Leftarrow \{ \text{assumption (15)} \} \\ \langle a \rangle : in \rightarrow a \end{array}$$

But, have we made any progress? After all, before we can apply (16), we have to prove (15). Fold satisfies this property, but this instance of (16) is trivial:

$\langle a \rangle \cdot \langle in \rangle = \langle a \rangle$. Now, it turns out that in a *relationally parametric* programming language [16], the proof obligation (15) amounts to the *free theorem* [21] of the polymorphic type

$$\beta : \forall A . (\mathbf{F} A \rightarrow A) \rightarrow (B \rightarrow A) , \quad (17)$$

where B is some fixed type. In other words, in such a language the proof obligation can be discharged by the type checker.

Returning to our example, we redefine *filter odd* as $(\lambda a . \langle \phi a \rangle) in$ where

$$\begin{aligned} \phi &: (\mathbf{L} \mathbb{N} b \rightarrow b) \rightarrow (\mathbf{L} \mathbb{N} b \rightarrow b) \\ \phi a \text{ Nil} &= a \text{ Nil} \\ \phi a (\text{Cons } (x, y)) &= \mathbf{if } \text{odd } x \mathbf{ then } a (\text{Cons } (x, y)) \mathbf{ else } y . \end{aligned}$$

We derived ϕ from the algebra \mathfrak{f} by abstracting away from *in*. The reader should convince herself that $\lambda a . \langle \phi a \rangle$ has indeed the desired polymorphic type (17). We can then invoke the acid rain rule (16) to obtain

$$\langle \mathfrak{s} \rangle \cdot (\lambda a . \langle \phi a \rangle) in = (\lambda a . \langle \phi a \rangle) \mathfrak{s} = \langle \phi \mathfrak{s} \rangle .$$

The example also shows that the *acid rain* rule is somewhat unstructured in that a *hylo* is hidden inside the abstraction λa . Without performing an additional beta-reduction, we can apply the rule only once. We obtain a more structured rule if we shift the abstraction to the algebra and achieve *cata-hylo fusion*: If τ is a transformation that takes \mathbf{F} -algebras to \mathbf{G} -algebras satisfying

$$h : \tau a \rightarrow \tau b : \mathbf{Alg}(\mathbf{G}) \quad \Leftarrow \quad h : a \rightarrow b : \mathbf{Alg}(\mathbf{F}) , \quad (18)$$

then

$$\langle a \rangle_{\mathbf{F}} \cdot \langle \tau in \leftarrow c \rangle_{\mathbf{G}} = \langle \tau a \leftarrow c \rangle_{\mathbf{G}} . \quad (19)$$

If τ is $\lambda a . a$, then this is just the fold/unfold law (14). For $\tau a = a \cdot \alpha$, this is essentially functor fusion (3). The proof of correctness is straightforward.

$$\begin{aligned} &\langle a \rangle_{\mathbf{F}} \cdot \langle \tau in \leftarrow c \rangle_{\mathbf{G}} = \langle \tau a \leftarrow c \rangle_{\mathbf{G}} \\ &= \{ \text{algebra fusion (10)} \} \\ &\langle a \rangle_{\mathbf{F}} : \tau in \rightarrow \tau a : \mathbf{Alg}(\mathbf{G}) \\ &= \{ \text{assumption (18)} \} \\ &\langle a \rangle_{\mathbf{F}} : in \rightarrow a : \mathbf{Alg}(\mathbf{F}) \end{aligned}$$

The proof obligation (18) once again amounts to a theorem for free, this time of the polymorphic type

$$\tau : \forall A . (\mathbf{F} A \rightarrow A) \rightarrow (\mathbf{G} A \rightarrow A) .$$

Using *cata-hylo fusion*, the running example simplifies to

$$\langle \mathfrak{s} \rangle \cdot \langle \phi in \rangle = \langle \phi \mathfrak{s} \rangle .$$

We can now also fuse a composition of folds:

$$\langle a \rangle \cdot \langle \tau_1 in \rangle \cdot \dots \cdot \langle \tau_n in \rangle \cdot \llbracket c \rrbracket = \langle (\tau_n \cdot \dots \cdot \tau_1) a \leftarrow c \rangle .$$

This demonstrates how the rewrite rule is able to achieve fusion over an entire pipeline of functions.

5.3 Generalised *destroy/unfoldr* Fusion

The *foldr/build* brand of shortcut fusion, and its generalisation to algebraic datatypes, is *fold-centric*. This limits the kind of functions that we can fuse, simply because some functions such as *zip* or *take* are not folds, or are not naturally written as folds. We can dualise *foldr/build* fusion to achieve an *unfold-centric* approach, called *destroy/unfoldr* [18]. To illustrate, consider the simple pipeline *take 5 · between*, where *take n* takes *n* elements (if available) from a list. It can be written as an unfold after an initialisation step: *take n = [[t]] · start n*, where *start n = (λ l . (n, l))*, and where the coalgebra *t* is given by

```

type State a = (ℕ, a)
t : State (μ(L a)) → L a (State (μ(L a)))
t (0, x)      = Nil
t (n + 1, x) = case out x of Nil → Nil; Cons (a, y) → Cons (a, (n, y)) .

```

Here we make explicit the notion that an unfold models the steps of a stateful computation. The coalgebra takes a state as an argument and uses it to produce a value and a new state. In this example, the state type pairs the input list with a natural number, enabling us to track the overall number of values produced. The number of elements to take, paired with the list where the values are to be taken from, forms the initial state.

We can dualise the *acid rain rule* to fuse the pipeline. If β is a transformation that satisfies

$$\beta c = \beta d \cdot h \quad \Leftarrow \quad h : c \rightarrow d \quad , \quad (20)$$

then

$$\beta c = \beta \text{out} \cdot \llbracket c \rrbracket . \quad (21)$$

Previously we exposed *in*, now we expose *out*. To apply the dual of acid rain we redefine *take n* as $(\lambda c . \llbracket \gamma c \rrbracket \cdot \text{start } n) \text{out}$, where

```

γ : (c → L a c) → (State c → L a (State c))
γ c (0, x) = Nil
γ c (n, x) = case c x of Nil → Nil; Cons (a, y) → Cons (a, (n - 1, y)) .

```

The transformation γ is derived from *t* by abstracting away from *out*. We can now tackle our example:

$$(\lambda c . \llbracket \gamma c \rrbracket \cdot \text{start } 5) \text{out} \cdot \llbracket \mathbf{b} \rrbracket = (\lambda c . \llbracket \gamma c \rrbracket \cdot \text{start } 5) \mathbf{b} = \llbracket \gamma \mathbf{b} \rrbracket \cdot \text{start } 5 .$$

The proof obligation (20) corresponds to the free theorem of

$$\beta : \forall C . (C \rightarrow F C) \rightarrow (C \rightarrow D) \quad , \quad (22)$$

where D is fixed. And, indeed, $\lambda c . \llbracket \gamma c \rrbracket \cdot \text{start } 5$ has the required type.

Similarly, we can dualise our more structured *cata-hylo fusion* to achieve *hylo-ana fusion*: If τ is a transformation that takes recursive F-coalgebras to recursive G-coalgebras satisfying

$$h : \tau c \rightarrow \tau d : \mathbf{Rec}(G) \quad \Leftarrow \quad h : c \rightarrow d : \mathbf{Rec}(F) , \quad (23)$$

then

$$\langle a \leftarrow \tau c \rangle_G = \langle a \leftarrow \tau \text{out} \rangle_G \cdot \llbracket c \rrbracket_F . \quad (24)$$

This time the proof obligation (23) cannot be discharged by the type checker alone as τ has to transform a *recursive* coalgebra into a *recursive* coalgebra! As an aside, the new rule cannot handle our running example as the two unfolds are separated by the initialisation function *start*.

Our example has focused on fusing the list parameter of *take*, yet if we admit to the fact that natural numbers are an inductive datatype, then *take* is really a function that consumes *two* data structures. The aforementioned *zip* is another function that consumes two data structures, and therefore has the potential to be fused with both of these inputs. Let us employ the expression $\text{zip} \cdot (\text{between} \times \text{between})$ as another example that can be written in terms of unfolds: $\llbracket \mathfrak{z} \rrbracket \cdot (\llbracket \mathfrak{b} \rrbracket \times \llbracket \mathfrak{b} \rrbracket)$. The algebra \mathfrak{z} is given by

$$\begin{aligned} \mathfrak{z} : (\mu(L a_1), \mu(L a_2)) &\rightarrow L(a_1, a_2) (\mu(L a_1), \mu(L a_2)) \\ \mathfrak{z}(x_1, x_2) &= \mathbf{case}(\text{out } x_1, \text{out } x_2) \mathbf{of} \\ &\quad (\text{Cons}(a_1, b_1), \text{Cons}(a_2, b_2)) \rightarrow \text{Cons}((a_1, a_2), (b_1, b_2)) \\ &\quad \text{otherwise} \rightarrow \text{Nil} . \end{aligned}$$

Our rules (21) and (24) are not applicable as we have *two* producers to the right of *zip*. Now, to fuse such a function, we need to employ *parallel hylo-ana fusion*: If τ satisfies,

$$\begin{aligned} h_1 \times h_2 : \tau(c_1, c_2) \rightarrow \tau(d_1, d_2) : \mathbf{Rec}(G) \\ \Leftarrow \quad h_1 : c_1 \rightarrow d_1 : \mathbf{Rec}(F_1) \wedge h_2 : c_2 \rightarrow d_2 : \mathbf{Rec}(F_2) , \quad (25) \end{aligned}$$

then

$$\langle a \leftarrow \tau(c_1, c_2) \rangle_G = \langle a \leftarrow \tau(\text{out}, \text{out}) \rangle_G \cdot (\llbracket c_1 \rrbracket_{F_1} \times \llbracket c_2 \rrbracket_{F_2}) . \quad (26)$$

Using this rule, we are now able to fuse the *zip* example:

$$\llbracket \zeta(\text{out}, \text{out}) \rrbracket \cdot (\llbracket \mathfrak{b} \rrbracket \times \llbracket \mathfrak{b} \rrbracket) = \llbracket \zeta(\mathfrak{b}, \mathfrak{b}) \rrbracket ,$$

where the transformation ζ is defined

$$\begin{aligned} \zeta : (b_1 \rightarrow L a_1 b_1, b_2 \rightarrow L a_2 b_2) &\rightarrow (b_1, b_2) \rightarrow L(a_1, a_2)(b_1, b_2) \\ \zeta(c_1, c_2)(x_1, x_2) &= \mathbf{case}(c_1 x_1, c_2 x_2) \mathbf{of} \\ &\quad (\text{Cons}(a_1, b_1), \text{Cons}(a_2, b_2)) \rightarrow \text{Cons}((a_1, a_2), (b_1, b_2)) \\ &\quad \text{otherwise} \rightarrow \text{Nil} . \end{aligned}$$

The proofs of correctness for (parallel) hylo-ana fusion are contained in extended version of this paper.

5.4 Church and Co-Church Encodings

In the two previous sections we have studied generalisations of *foldr/build* and *destroy/unfoldr* fusion. We have noted that $\llbracket - \rrbracket$ generalises the list function *foldr*, and, likewise, $\llbracket - \rrbracket$ generalises *unfoldr*. We have been silent, however, about their counterparts *build* and *destroy*. It is time to break that silence, and in the process, provide a fresh perspective on recursive datatypes. For simplicity, we assume that we are working in **Set**.⁴

Consider again the polymorphic type of β (17) repeated below.

$$\forall A . (F A \rightarrow A) \rightarrow (B \rightarrow A) \cong B \rightarrow (\forall A . (F A \rightarrow A) \rightarrow A)$$

We have slightly massaged the type to bring B to the front. The universally quantified type on the right is known as the *Church encoding* of μF [4]. The type is quite remarkable as it encodes a recursive type without using recursion. One part of the isomorphism $\mu F \cong \forall A . (F A \rightarrow A) \rightarrow A$ is given by the acid rain rule (16). The following derivation, which infers the isomorphisms, makes this explicit—the initial equation is (16) with the arguments of β swapped.

$$\begin{aligned} & \forall a . \llbracket a \rrbracket (\beta b \text{ in}) = \beta b a \\ \iff & \{ \text{change of variables } \beta b = \gamma \} \\ & \forall a . \llbracket a \rrbracket (\gamma \text{ in}) = \gamma a \\ \iff & \{ \text{extensionality} \} \\ & \lambda a . \llbracket a \rrbracket (\gamma \text{ in}) = \gamma \\ \iff & \{ \text{define } \text{toChurch } x = \lambda a . \llbracket a \rrbracket x \} \\ & \text{toChurch } (\gamma \text{ in}) = \gamma \\ \iff & \{ \text{define } \text{fromChurch } \gamma = \gamma \text{ in} \} \\ & \text{toChurch } (\text{fromChurch } \gamma) = \gamma \end{aligned}$$

The isomorphism *toChurch*, creates a function whose argument is an algebra and which folds that algebra over the given data structure. Its converse *fromChurch*, commonly called *build*, applies this function to the *in* algebra. Going back and forth, we get back the original structure: $\text{fromChurch } (\text{toChurch } s) = s$. This is the other part of the isomorphism, which follows directly from fold reflection.

As to be expected, everything nicely dualises. The polymorphic type (22) gives rise to the *co-Church encoding*.

$$\forall C . (C \rightarrow F C) \rightarrow (C \rightarrow D) \cong (\exists C . (C \rightarrow F C) \times C) \rightarrow D$$

Think of the co-Church encoding $\exists C . (C \rightarrow F C) \times C$ as the type of state machines encapsulating a transition function $C \rightarrow F C$ and the current state C .

The conversion to (co-)Church-encoding types are central to the concept of shortcut fusion. By changing representations to one with the recursion “built-in”, we can write our transformations as non-recursively-defined (co-)algebras. Unlike recursive programs, compositions of these (co-)algebras can be optimised

⁴ The development can be generalised using ends and coends [14].

by the compiler to remove any intermediate allocations. All that remains is for the programmer to instruct the compiler to remove any unnecessary conversions, i.e. cases of *toChurch* · *fromChurch*. Removing these transformations preserves the semantics of the program because we can prove the isomorphism between these representations. More importantly, however, prevents us from producing a data structure only to immediately consume it. The co-Church encoding also underlies the original formulation of stream fusion, which we consider next.

5.5 Stream Fusion

The *foldr/build* flavour of fusion is fold-centric, in that it requires all functions that are intended to be fusible to be written as folds; similarly, *destroy/unfoldr* is unfold-centric. The boundaries of these world views are fuzzy. A *zip* can be written as a fold, the snag is that only one of the two inputs can be fused [9, §9]. Along a similar vein, a *filter* for the odd natural numbers, which we wrote before as a fold, can also be written as an unfold: $\llbracket \mathbf{f} \rrbracket$ where

$$\begin{aligned} \mathbf{f} &: \mu(\mathbb{L}\mathbb{N}) \rightarrow \mathbb{L}\mathbb{N}(\mu(\mathbb{L}\mathbb{N})) \\ \mathbf{f}x &= \mathbf{case\ out\ } x \mathbf{ of\ } Nil \rightarrow Nil; Cons(x, y) \rightarrow \mathbf{if\ odd\ } x \mathbf{ then\ } Cons(x, y) \mathbf{ else\ } \mathbf{f}y \ . \end{aligned}$$

The coalgebra \mathbf{f} is recursive and thus theoretically fine, but it is also recursive in its definition and this is a practical problem. A coalgebra must be non-recursively defined for it to be fused with others. We have two definitions and are caught between two worlds; is it possible to free ourselves?

Perhaps surprisingly, the answer is yes. Let us first try to eliminate the recursion from the definition above—the rest will then fall out. The idea is to use a different base functor, one that allows us to skip list elements. We draw inspiration from stream fusion [5] here:

data $\mathbb{S} a b = Done \mid Yield(a, b) \mid Skip b \ .$

instance *Functor* ($\mathbb{S} a$) **where**
 $fmap\ f\ Done = Done$
 $fmap\ f\ (Skip\ b) = Skip\ (f\ b)$
 $fmap\ f\ (Yield\ (a, b)) = Yield\ (a, f\ b)$

The *filter* coalgebra can now be written as a composition of *out* with

$$\begin{aligned} \mathbf{f} &: \mathbb{S}\mathbb{N} b \rightarrow \mathbb{S}\mathbb{N} b \\ \mathbf{f}\ Done &= Done \\ \mathbf{f}(Skip\ y) &= Skip\ y \\ \mathbf{f}(Yield\ (x, y)) &= \mathbf{if\ odd\ } x \mathbf{ then\ } Yield\ (x, y) \mathbf{ else\ } Skip\ y \ . \end{aligned}$$

So, $filter = \llbracket \mathbf{f} \cdot out \rrbracket$. Something interesting has happened: since \mathbf{f} is a natural transformation, we also have $filter = \langle in \cdot \mathbf{f} \rangle$. We are unstuck; *filter* is both a fold and an unfold. Moreover, it is an application of a mapping function: $filter = \mu\mathbf{f}$.

In general, consumers are folds, transformers are maps, and producers are unfolds. An entire pipeline of these can be fused into a single hylomorphism:

$$\langle a \rangle \cdot \mu\alpha_1 \cdots \mu\alpha_n \cdot \llbracket c \rrbracket = \langle a \cdot \alpha_1 \cdots \alpha_n \leftarrow c \rangle \ .$$

Inspecting the types, the rule is clear:

$$A \xleftarrow{\llbracket a \rrbracket} \mu F_0 \xleftarrow{\mu \alpha_1} \mu F_1 \quad \dots \quad \mu F_{n-1} \xleftarrow{\mu \alpha_n} \mu F_n \xleftarrow{\llbracket c \rrbracket} C .$$

In a sense, the introduction of *Skip* keeps the recursion in sync. Each transformation consumes a token and produces a token. Before, *filter* possibly consumed several tokens before producing one. We are finally in a position to deal with the example from the introduction, written in terms of the combinators we have

$$\llbracket \mathfrak{s} \rrbracket \cdot \mu(\mathfrak{m} \, sq) \cdot \mu(\mathfrak{f} \, odd) \cdot \llbracket \mathfrak{b} \rrbracket = \llbracket \mathfrak{s} \cdot \mathfrak{m} \, sq \cdot \mathfrak{f} \, odd \leftarrow \mathfrak{b} \rrbracket .$$

Utilising streams in this fashion is an instance of data abstraction; although we wish to present the `List` type using $\mu(\mathbb{L} \, a)$, we intend to do all the work using $\mu(\mathbb{S} \, a)$. We have functions $\rightarrow\mathbb{S}$ and $\leftarrow\mathbb{S}$ to convert to and from streams, respectively. They are defined as an algebra and a coalgebra that allow us to consume streams using a fold and produce them using an unfold:

$$\begin{aligned} \leftarrow\mathbb{S} : \mathbb{S} \, a \, (\mu(\mathbb{L} \, a)) &\rightarrow (\mu(\mathbb{L} \, a)) \\ \leftarrow\mathbb{S} \, Done &= in \, Nil \\ \leftarrow\mathbb{S} \, (Skip \, xs) &= xs \\ \leftarrow\mathbb{S} \, (Yield \, (x, xs)) &= in \, (Cons \, (x, xs)) \\ \rightarrow\mathbb{S} : \mu(\mathbb{L} \, a) &\rightarrow \mathbb{S} \, a \, (\mu(\mathbb{L} \, a)) \\ \rightarrow\mathbb{S} \, (in \, Nil) &= Done \\ \rightarrow\mathbb{S} \, (in \, (Cons \, (x, xs))) &= Yield \, (x, xs) . \end{aligned}$$

We must prove that our stream implementations, together with the conversion functions, fulfil the same specification as the analogous functions over $\mu(\mathbb{L} \, a)$ (cf. Lemma 1 and Theorem 3 in [23]). This is called the *data abstraction property*. In our framework, this obligation is expressed as a simple equality between a conventional list function definition and its associated stream version composed with our conversion functions. For example, for *filter* we must prove

$$filter = \llbracket \leftarrow\mathbb{S} \rrbracket \cdot \mu \mathfrak{f} \cdot \llbracket \rightarrow\mathbb{S} \rrbracket ,$$

Because we can phrase these functions as folds, unfolds, and natural transformations, the proof is straightforward, using the laws we have set out in previous sections. We leave it as an exercise to the reader.

Just as for lists, every datatype can be extended with a *Skip*. Although stream fusion is the first to make use of this augmentation, we note its relation to Capretta’s representation of general recursion in type theory [1], which proposes adding a “computation step” constructor to coinductive types.

6 Related Work

Wadler first introduced the idea of simplifying the fusion problem with his deforestation algorithm [22]. This was limited to so-called *treeless* programs, a subset of first-order programs. The fusion transformation proposed by Chin [3] generalises Wadler’s deforestation. It uses a program annotation scheme to recognise

the terms that can be fused and skip the terms that cannot. Sheard and Fegaras focus on the use of folds over algebraic types as a recursion scheme [17]. Their algorithm for normalising the nested application of folds is based on the fold fusion law. Their recursion schemes are suitably general to handle functions such as *zip* that recurse over multiple data structures simultaneously [6].

Gill et al. first introduced the notion of shortcut fusion with *foldr/build* fusion [9] for Haskell. This allowed programs written as folds to be fused. It was subsequently introduced into the List library for Haskell in GHC. Takano and Meijer [19] provided a calculational view of fusion and generalised it to arbitrary data structures. It generalised the fusion law by using hylomorphisms and also noted the possibility of dualising *foldr/build* fusion. They worked in the setting of **Cpo**, however, where hylomorphisms do not have unique solutions, only canonical ones. Takano and Meijer claimed that, even when restricted to lists, their method is more powerful than that of Gill et al. as theirs could fuse both parameters of *zip*. This was incorrect, and the need for an additional parallel rule for *zip* was pointed out later by Hu et al. [11]. Their extension is what we present as the parallel hylo-ana rule.

Svenningsson provided an actual implementation of *destroy/unfoldr* fusion [18], where he showed how *filter*-like functions could be expressed as unfolds. Svenningsson did not, however, solve the issue of recursion in the coalgebras of such functions, which could therefore not be fused even though they could be written as unfolds. This was addressed by Coutts et al., who presented stream fusion [5], which introduced the *Skip* constructor as a way to encode non-productive computation steps, similar to Capretta’s work on encoding general recursion in type theory [1].

The correctness and generalisation of fusion has been explored in many different settings. In addition to the work of Takano and Meier, Ghani et al. generalised *foldr/build* to work with datatypes “induced by inductive monads”. Johann and Ghani further showed how to apply initial algebra semantics, and thus *foldr/build* fusion, to nested datatypes [12]. Voigtländer has also used free theorems to show correctness, specifically of the *destroy/build* rule [20].

7 Conclusions

We have presented a framework that has allowed us to bring three fusion techniques into the same setting. We have exploited recursive coalgebras and hylomorphisms as ‘the rug that ties the room together’. This enabled us to formally describe and reason about these fusion techniques. In doing so, we have exposed their underlying foundations, including the importance of Church and co-Church encodings. The fact that our hylomorphisms have unique solutions plays a central rôle. The knock-on effect is that we gain clear, short proofs thanks to the calculational properties available to us.

References

1. Capretta, V.: General recursion via coinductive types. *Logical Methods in Computer Science* 1(2), 1–28 (2005)
2. Capretta, V., Uustalu, T., Vene, V.: Recursive coalgebras from comonads. *Information and Computation* 204(4), 437–468 (2006)
3. Chin, W.N.: Safe Fusion of Functional Expressions. In: *LISP and functional programming*. pp. 11–20 (1992)
4. Church, A.: *The calculi of lambda-conversion*. Annals of Mathematics Studies No. 6, Princeton University Press (1941)
5. Coutts, D., Leshchinskiy, R., Stewart, D.: Stream Fusion: From Lists to Streams to Nothing At All. In: *ICFP '07*. pp. 315–326 (2007)
6. Fegaras, L., Sheard, T., Zhou, T.: Improving Programs which Recurse over Multiple Inductive Structures. In: *PEPM'94* (June 1994)
7. Fokkinga, M.M., Meijer, E.: Program calculation properties of continuous algebras. Technical Report CS-R9104, CWI, Amsterdam (Jan 1991)
8. Freyd, P.J.: Remarks on algebraically compact categories. In: Fourman, M.P., Johnstone, P.T., Pitts, A.M. (eds.) *Applications of Categories in Computer Science*, LMS Lecture Note Series, vol. 177, pp. 95–106. Cambridge University Press (1992)
9. Gill, A., Launchbury, J., Peyton Jones, S.L.: A Short Cut to Deforestation. In: *Functional programming languages and computer architecture*. pp. 223–232 (1993)
10. Hinze, R., Harper, T., James, D.W.H.: *Theory and Practice of Fusion*. Tech. Rep. CS-RR-11-01, Oxford University Computing Laboratory (2011)
11. Hu, Z., Iwasaki, H., Takeichi, M.: An Extension of The Acid Rain Theorem. In: *Functional and Logic Programming*. pp. 91–105 (1996)
12. Johann, P., Ghani, N.: Initial algebra semantics is enough! In: *Typed Lambda Calculi and Applications*. LNCS, vol. 4583, pp. 207–222 (2007)
13. Lambek, J.: A fixpoint theorem for complete categories. *Math. Zeitschr.* 103, 151–161 (1968)
14. Mac Lane, S.: *Categories for the Working Mathematician*. Graduate Texts in Mathematics, Springer-Verlag, Berlin, 2nd edn. (1998)
15. Meijer, E., Fokkinga, M., Paterson, R.: Functional programming with bananas, lenses, envelopes and barbed wire. In: *FPCA*. LNCS, vol. 523, pp. 124–144 (1991)
16. Reynolds, J.C.: Types, abstraction and parametric polymorphism. In: Mason, R.E.A. (ed.) *Information Processing '83*, pp. 513–523. North-Holland, Amsterdam (1983)
17. Sheard, T., Fegaras, L.: A Fold for All Seasons. In: *Functional programming languages and computer architecture*. pp. 233–242 (1993)
18. Svenningsson, J.: Shortcut fusion for Accumulating Parameters & Zip-like Functions. In: *ICFP '02*. pp. 124–132 (2002)
19. Takano, A., Meijer, E.: Shortcut deforestation in calculational form. In: *Functional programming languages and computer architecture*. pp. 306–313 (1995)
20. Voigtländer, J.: Proving correctness via free theorems: the case of the destroy/-build-rule. In: *Partial Eval. and Semantics-Based Prog. Manip.* pp. 13–20 (2008)
21. Wadler, P.: Theorems for free! In: *FPCA*. pp. 347–359 (1989)
22. Wadler, P.: Deforestation: transforming programs to eliminate trees. *Theoretical Computer Science* 73(2), 231 – 248 (1990)
23. Wang, M., Gibbons, J., Matsuda, K., Hu, Z.: Gradual refinement: Blending pattern matching with data abstraction. In: Bolduc, C., Desharnais, J., Ktari, B. (eds.) *Mathematics of Program Construction*. LNCS, vol. 6120, pp. 397–426 (2010)