Chapter 8

Short Cut Fusion of Recursive Programs with Computational Effects

Neil Ghani, Patricia Johann[1]
Category: Research

Abstract: Fusion is the process of improving the efficiency of modularly constructed programs by transforming them into monolithic equivalents. This paper defines a generalization of the standard `build` combinator which expresses uniform production of functorial contexts containing data of inductive types. It also proves correct a fusion rule which generalizes the `fold/build` and `fold/buildp` rules from the literature, and eliminates intermediate data structures of inductive types without disturbing the contexts in which they are situated. An important special case arises when this context is monadic. When it is, a second rule for fusing combinations of producers and consumers via monad operations, rather than via composition, is also available. We give examples illustrating both rules, and consider their coalgebraic duals as well.

8.1 GENERALIZING SHORT CUT FUSION

8.1.1 Introducing Short Cut Fusion

Fusion is the process of improving the efficiency of modularly constructed programs by transforming them into monolithic equivalents. Short cut fusion [6] is concerned with eliminating list traversals from compositions of components that are "glued" together via intermediate lists. Short cut fusion uses a local transformation — known as the `foldr/build` rule — to fuse computations which can be written as compositions of applications of the uniform list-consuming function `foldr` and the uniform list-producing function `build` given by

---

[1] University of Strathclyde, Glasgow, Scotland
{Neil.Ghani,Patricia.Johann}@cis.strath.ac.uk.

```
newtype Mu f = In {unIn :: f (Mu f)}

fold :: Functor f => (f a -> a) -> Mu f -> a
fold h (In k) = h (fmap (fold h) k)

build :: Functor f =>
   (forall a. (f a -> a) -> c -> a) -> c -> Mu f
build g = g In

fold k . build g = g k
```

FIGURE 8.1.   The `fold` **and** `build` **combinators and** `fold`/`build` **rule.**

```
foldr :: (b -> a -> a) -> a -> [b] -> a
foldr c n [] = n
foldr c n (x:xs) = c x (foldr c n xs)

build :: (forall a. (b -> a -> a) -> a -> a) -> [b]
build g = g (:) []
```

The function `foldr` is standard in the Haskell prelude. Intuitively, `foldr c n xs` produces a value by replacing all occurrences of `(:)` in `xs` by `c` and the occurrence of `[]` in `xs` by `n`. Thus, `sum xs = foldr (+) 0 xs` sums the (numeric) elements of the list `xs`. Uniform production of lists, on the other hand, is accomplished using the combinator `build`, which takes as input a type-independent template for constructing "abstract" lists and produces a corresponding "concrete" list. Thus, `build (\c n -> c 4 (c 7 n))` produces the list `[4,7]`. Uniform list transformers can be written in terms of both `foldr` and `build`. For example, the function `map` can be implemented as

```
map :: (a -> b) -> [a] -> [b]
map f xs = build (\c n -> foldr (c . f) n xs)
```

The `foldr`/`build` rule capitalizes on the uniform production and consumption of lists to improve the performance of list-manipulating programs. It says

$$\text{foldr c n (build g) = g c n} \qquad (8.1)$$

If `sqr x = x * x`, then this rule can be used, for example, to transform the modular function `sum . map sqr :: [Int] -> Int` which produces an intermediate list into an optimized form which does not:

```
sum (map sqr xs) = foldr (+) 0
                     (build (\c n -> foldr (c . sqr) n xs))
                 = (\c n -> foldr (c . sqr) n xs) (+) 0
                 = foldr ((+) . sqr) 0 xs
```

8.1.2   Short Cut Fusion for Inductive Types

Inductive datatypes are fixed points of functors. Functors can be implemented in Haskell as type constructors supporting `fmap` functions as follows:

```
buildp :: Functor f =>
   (forall a. (f a -> a) -> c -> (a,z)) -> c -> (Mu f, z)
buildp g = g In

fmap (fold k) . buildp g = g k
```

FIGURE 8.2. The `buildp` **combinator and** `fold`/`buildp` **fusion rule.**

```
class Functor f where
  fmap :: (a -> b) -> f a -> f b
```

The function `fmap` is expected to satisfy the two semantic functor laws stating that `fmap` preserves identities and composition.It is well-known that analogues of `foldr` exist for every inductive datatype. As shown in [4, 5], every inductive type also has an associated generalized `build` combinator; the extra type `c` in the type of `build` is motivated in those papers and to lesser extent in Section 8.3 below. These combinators can be implemented generically in Haskell as in Figure 8.1. There, `Mu f` represents the least fixed point of the functor `f`, and `In` represents the structure map for `f`, i.e., the "bundled" constructors for the datatype `Mu f`. The `fold`/`build` fusion rule for inductive types can be used to eliminate data structures of type `Mu f` from computations. The `foldr` and `build` combinators for lists can be recovered by taking `f` to be the functor whose fixed point is `[b]`. The `foldr`/`build` rule can be recovered by taking `c` to be the unit type as well. As usual, `fold` and `build` implement the isomorphisms between inductive types and their Church encodings.

8.1.3   Short Cut Fusion in Context

Short cut fusion handles compositions `g . f` in which the data structure produced by `f` is passed from `f` to `g`. But what if `f` produces not just a single data structure, but multiple such structures, embeds these data structures in a non-trivial context, and passes the result to `g` for consumption of these data structures "in context"? Is it possible to eliminate these intermediate data structures from `g . f` while keeping the context information, which `g` may need to compute its result, intact? Standard fusion techniques cannot achieve this: the intermediate data structures produced by `f` cannot be decoupled from the context in which they are situated. In [2], Fernandes et al. introduce a technique for fusing compositions `g . f` in which `f` passes to `g` not only the intermediate data structure produced by `f`, but an additional datum as well. Although `g` requires this datum to compute its result, it is not used when processing the intermediate data structure, and so only the data structure itself needs to be eliminated from `g . f`. To do this, [2] uses a variant of the standard `fold`/`build` rule based on the combinator `buildp`, which captures the extra datum by returning a data structure embedded in a pair context. The datatype-generic `buildp` combinator and its associated `fold`/`buildp` fusion rule are given in Figure 8.2. There, `fmap` is the map function

```
fmap :: (a -> b) -> (a,z) -> (b,z)
fmap f (a,z) = (f a, z)
```

```
superbuild :: (Functor f, Functor h) =>
    (forall a. (f a -> a) -> c -> h a) -> c -> h (Mu f)
superbuild g = g In


fmap (fold k) . (superbuild g) = g k
```

FIGURE 8.3.   The `superbuild` **combinator and** `fold/superbuild` **fusion rule.**

which witnesses the fact that the type constructor `h` given by `h x = (x,z)` is a
functor. The context information produced by `buildp` and used by the consumer
in the left-hand side of the `fold/buildp` fusion rule is reflected in the pair return
types of `buildp` and its template argument, as well as in the mapping of `fold`
across the pair in the associated `fold/buildp` rule. This rule eliminates interme-
diate data structures within the context of pairing with an additional datum.

   But now suppose we want to write a function

```
gsplitWhen :: (b -> Bool) -> [b] -> [[b]]
```

which splits a list into sublists at every element that satisfies a given `p`. Note that
the function `gsplitWhen` splits lists into arbitrary numbers of sublists, depend-
ing on the data they contain, and that the type `z` in the type of `buildp` cannot
be instantiated to allow the return of a number of lists which has the potential to
change on each program run. This means that `gsplitWhen` cannot be written in
terms of `buildp`. Moreover, compositions of `gsplitWhen` with functions that
consume each of the individual "inner" lists produced by `gsplitWhen` but re-
quire the information inherent in its "context list" to compute their results cannot
be fused using the `fold/buildp` rule. But why try to structure programs only
with contexts of the form `(-,z)`? That is, why not consider a generalization
of the `buildp` combinator, and a generalization of the `fold/buildp` fusion rule
which can be used to eliminate intermediate data structures, like those returned
by `gsplitWhen`, which appear in contexts other than just pairs? That is pre-
cisely what this paper does. We call these generalizations `superbuild` and the
`fold/superbuild` rule, respectively. Like `buildp` and the `fold/buildp` rule,
our `superbuild` combinator and `fold/superbuild` fusion rule are available at
every inductive datatype. Datatype-generic versions are given in Figure 8.3; note
that the type of superbuild is actually generic in both `f` and `h`. The generaliza-
tion of the pair context in the type of `buildp` is captured by the replacement in
the type of `superbuild` of the type `(x,z)` by the type `h x` for a more general
"context functor" `h`. This generalization is further reflected in the replacement of
the `fmap` function for pairs in the `fold/buildp` rule by the `fmap` function for the
more general context functor `h` in the `fold/superbuild` rule. The `fold` com-
binator in the `fold/superbuild` rule is the one for `Mu f`, as usual. These `fmap`
and `fold` functions are guaranteed to be defined precisely because the type of
`superbuild` requires both `f` and `h` to be functors. We argue in Section 8.3 that
the `fold/superbuild` rule holds for a large class of functors `h`.

   Taking `h x = x` gives the generalized `build` combinator and `fold/build`
rule from Figure 8.1, while taking `h x = (x,z)` gives the `buildp` combina-

tor and `fold/ buildp` rule from Figure 8.2. In general, the `fold/superbuild` rule can fuse compositions in which context information describable by non-pair functors is passed, along with intermediate data structures, from producer to consumer. Indeed, the `fold/superbuild` rule eliminates intermediate structures of type `Mu f` obtained by mapping a consumer expressed as a `fold` over the data of type `Mu f` stored in a context specified by a functor `h`. Thus, setting `c = [b]`, `h x = [x]`, and `f` to be the functor whose least fixed point is `[b]`, we can write

```
gsplitWhen p = superbuild go where
go c n z = case z of
            []       -> []
            [w]      -> [c w n]
            (w : ws) -> let xs = go c n ws
                        in if p w then (c w n) : xs
                            else (c w (head xs)) : (tail xs)
```

If `lgh = foldr (\x -> (1+)) 0` then using the `fold/superbuild` rule to fuse the composition `evLghs = map lgh . gsplitWhen even` gives

```
evLghs' z = case z of
            []       -> []
            [w]      -> [1]
            (w : ws) -> let xs = evLghs' ws
                          in if even w then 1 : xs
                              else (head xs + 1) : (tail xs)
```

Note that `evLghs'` trades production and consumption of the list of intermediate lists returned by `gsplitWhen even` in `evLghs` for production of the corresponding list of values obtained by applying `lgh` to each such list.

8.1.4 Short Cut Fusion in Effectful Contexts

The ability to fuse intermediate data structures in context turns out to be the key to extending short cut fusion to the effectful setting. Although fusion in the presence of computational effects has been studied by other researchers (see, e.g., [10, 12, 14, 16]), short cut fusion in particular has not previously been formally explored in this context. To perform short cut fusion in an effectful context, the functional argument to `superbuild`, and thus `superbuild` itself, must have a monadic return type. Monads can be implemented in Haskell as type constructors supporting `>>=` and `return` operations as follows; these operations are expected to satisfy the semantic monad laws.

```
class Monad m where
  return :: a -> m a
  (>>=)  :: m a -> (a -> m b) -> mb
```

If `m` is a monad, setting `h` to `m` in `superbuild`'s type gives the `msuperbuild` combinator in Figure 8.4. The accompanying `fold/msuper build` rule is the natural "monadification" of the `fold/superbuild` rule; we give an example of its use in Section 8.2. As we see in Section 8.3.3 the `fold/msuper build` rule

```
msuperbuild :: (Functor f, Monad m) =>
    (forall a. (f a -> a) -> c -> m a) -> c -> m (Mu f)
msuperbuild g = g In


msuperbuild g c >>= fold k = g k c >>= id
```

FIGURE 8.4.   The `msuperbuild` **combinator and** `fold/msuperbuild` **fusion rule.**

follows from the `fold/super build` rule and standard properties of monad multiplication. This rule does not eliminate the monadic context described by `m`, but does eliminate intermediate data structures of type `Mu f` within that monadic context. Moreover, although the rule does not change the type of the context containing the data structure, it can change the context itself, and so is more sophisticated than its non-monadic counterpart.

The remainder of this paper is structured as follows. In Section 8.2 we apply our new `fold/superbuild` and `fold/msuperbuild` rules to substantive examples. In Section 8.3 we show how the `superbuild` and `msuperbuild` combinators are derived from initial algebra semantics, and prove the correctness of their associated fusion rules. In Section 8.4 we give non-monadic and monadic `superdestroy/unfold` rules dual to our non-monadic and monadic `fold/super build` rules; our results for `superbuild` and `msuperbuild` are easily dualized to prove them correct. In Section 8.5 we discuss related work, and in Section 8.6 we conclude and offer directions for future research. A Haskell implementation of our results and an additional example highlighting the versatility of our rules are available at `http://www.cs.nott.ac.uk/~nxg`.

8.2   EXAMPLES

In this section we give some more sophisticated examples showcasing the power of the `fold/superbuild` and `fold/msuperbuild` fusion rules. Our first example shows that the `fold/superbuild` rule can be used to eliminate intermediate data structures other than lists. Our second example shows that the `fold/msuperbuild` rule can eliminate data structures within the state monad.

Example 8.1. Consider the simple arithmetic expression datatype given by

```
data Opr = Add | Mul | Sub deriving (Eq, Show)
```

```
data Expr  = Lit Int | Op Opr Expr Expr deriving (Eq, Show)
```

The `fold` combinator for expressions, the instance of `superbuild` for expressions where `c` is `Expr` and `h  x` is `[x]`, and the associated fusion rule are

```
foldExpr :: (Int -> a) -> (Opr -> a -> a -> a) -> Expr -> a
foldExpr l o e = case e of
                    Lit i -> l i
                    Op op e1 e2 -> o op (foldExpr l o e1)
```

```
                                      (foldExpr l o e2)

superbuildExpr :: (forall a. (Int -> a) ->
    (Opr -> a -> a -> a) -> Expr -> [a]) -> Expr -> [Expr]
superbuildExpr g = g Lit Op

map (foldExpr l o e) (superbuildExpr g) = g l o e
```

If we define `opToHas Add = (+)`, `opToHas Mul = (*)`, and `opToHas Sub = (-)`, then we can implement an interpreter which traces the evaluation steps taken in computing the integer values represented by expressions as

```
trace :: Expr -> [Expr]
trace = superbuildExpr g

g :: (Int -> a) -> (Opr -> a -> a -> a) -> Expr -> [a]
g l o e = case e of
            Lit i -> [l i]
            Op op e1 e2 -> let b1 = foldExpr l o e1
                               b2 = foldExpr l o e2
                               e' = o op b1 b2
                               Lit k = last (g Lit Op e1)
                               Lit j = last (g Lit Op e2)
                               b1s = g l o e1
                               b2s = g l o e2
                           in  (e' : (map (\x -> o op x b2)
                                     (tail b1s))
                                  ++ (map (o op (last b1s))
                                     (tail b2s))
                                  ++ [l (opToHas op k j)])
```

For example, if

```
myexp = Op Mul (Op Add (Lit 5) (Lit 6))
               (Op Sub (Lit 7) (Lit 4))
```

then `trace myexp` generates the trace

```
[ Op Mul (Op Add (Lit 5) (Lit 6)) (Op Sub (Lit 7) (Lit 4)),
  Op Mul (Lit 11) (Op Sub (Lit 7) (Lit 4)),
  Op Mul (Lit 11) (Lit 3), Lit 33 ]
```

Once an interpreter trace is built, we can perform various analyses of it. For example, we can measure the computational effort required to compute the value represented by each expression arising in the evaluation of a given expression. For this we use `count`, which counts 0 units of effort to compute a literal, 2 to perform an addition, 3 to perform a subtraction, and 5 to perform a multiplication.

```
count Add x y = x + y + 2
count Sub x y = x + y + 3
count Mul x y = x + y + 5
```

We then have

```
costExprs :: Expr -> [Int]
costExprs expr = map (foldExpr (\x -> 0) count)
                     (superbuildExpr g expr)
```

Thus `costExprs myexp` generates the result `[10,8,5,0]`. Fusion using the
`fold/superbuild` rule gives `costExprs' g (\x -> 0) count` — an equiv-
alent function in which no intermediate list of expressions is constructed.

Example 8.2. Pardo [14] shows that graph traversal algorithms, such as depth-first
and breadth-first traversal, can be written as calls to a monadic `unfold` combi-
nator. We show that these algorithms can be written in terms of `msuper build`.
The relationship between monadic and non-monadic `unfold` combinators, and
between `superbuild` and `msuperbuild`, is discusssed in Section 8.5 below..

A graph traversal is represented as a function which takes as input a list of
root vertices of a graph and returns a list containing the vertices met in order as
the graph is traversed. We can represent the vertices of a graph by integers, and a
graph by an adjacency list function for vertices as follows:

```
type V     = Int
type Graph = V -> [V]
```

In a graph traversal, each vertex is visited at most once. To avoid repeated visits
to vertices we can use the state monad [13, 15] to maintain a list of vertices visited
previously in the computation and thread this list through the traversal. We there-
fore define a datatype of *visit-dependent data*, each element of which is a function
taking a list of vertices already visited as input and returning a datum depending
on that list together with an updated list of visited vertices. We have

```
data State s a = State {runstate :: s -> (s,a)}

instance Monad (State s) where
  return a = State (\s -> (s,a))
  t >>= f  = State (\s -> let (s',v) = runstate t s
                          in runstate (f v) s')

type Vis a = State [V] a
```

Visit-dependent data support the following useful auxiliary functions:

```
data Unit = Unit

emp :: Vis a -> a
emp xs = snd (runstate xs [])

sunion :: V -> Vis Unit
sunion v = State (\vs -> (v:vs, Unit))

mem :: V -> Vis Bool
mem v = State (\vs -> (vs, elem v vs))
```

With this machinery we can define depth-first traversal as in Figure 8.5. There,
`dft` first allocates an empty list of visited vertices, then runs `depthFirst`, yield-
ing a final list of visited vertices, and then de-allocates this visitation list and

```
dft :: (V -> [V]) -> [V] -> [V]
dft g vs = emp (depthFirst g vs)


depthFirst :: (V -> [V]) -> [V] -> Vis [V]
depthFirst g = msuperbuild (df g)


df :: (V -> [V]) ->
        forall a. (V -> a -> a) -> a -> [V] -> Vis a
df g c n vs = case vs of
  []     -> return n
  (x:xs) -> mem x >>=
              (\b -> if b then df g c n xs
                      else sunion x >>=
                           (\z -> df g c n (g x ++ xs) >>=
                           (\ys -> return (c x ys))))
```

FIGURE 8.5. Depth-first graph traversal functions.

returns the list resulting from the traversal. At each iteration of the traversal, `df` explores the current list of roots in `vs` to find a vertex it has not reached before. This is done by removing from the front of `vs` all vertices for which `mem x` is true until either an unvisited vertex or the end of `vs` is encountered. When an unvisited vertex `x` is encountered, `df` adds `x` to the list of vertices visited, recursively computes the depth-first traversals of the graphs rooted at `x`'s children, as well as those specified by the rest of the vertices in `vs`, and then returns the list of vertices obtained by adding `x` to the list of vertices recording the order in which the rest of the vertices are traversed. The code for breath-first search is identical, except that the function `bf` corresponding to `df` uses `xs ++ g x` rather than `g x ++ xs`. To traverse a particular graph we specify the desired traversal, the graph's adjacency list function, and its root vertices. For example, if the graph G is modeled by `g 0 = [2,1]`, `g 1 = []`, and `g x = [x+1]`, then `depthFirst g [0]` computes the depth-first search of G starting at root vertex `0`.

For example, to consume the result of a traversal with `filtergph odd` where

```
filtergph :: (V -> Bool) -> [V] -> Vis [V]
filtergph p = foldr (\v i -> if p v then return (v : emp i)
                              else return (emp i))(return [])
```

we can write one of the following, depending on the desired traversal

```
dfFil g = emp (depthFirst g [0] >>= filtergph odd)

bfFil g = emp (breadthFirst g [0] >>= filtergph odd)
```

To perform the same computations without constructing the intermediate lists of visit-dependent vertices, we can use the `fold/msuperbuild` rule to get

```
fn v i = if odd v then return (v : emp i)
                  else return (emp i)
```

```
dfFil' g = emp ((df g) fn (return []) [0] >>= id)

bfFil' g = emp ((bf g) fn (return []) [0] >>= id)
```

Note that the lists obtained by taking any non-empty initial segments of the results of `dfFil g` and `bfFil g` — and thus of `dfFil' g` and `bfFil' g` — reflect the distinction between the underlying depth-first and breadth-first traversals.

8.3 CORRECTNESS

8.3.1 Categorical Preliminaries

Let $\mathscr{C}$ be a category and F be an endofunctor on $\mathscr{C}$. An *F-algebra* is a morphism $h : FA \rightarrow A$ in $\mathscr{C}$. The object A is called the *carrier* of the F-algebra. The F-algebras for a functor F are the objects of a category called the *category of F-algebras* and denoted $F\text{-}\mathscr{A}lg$. In the category of F-algebras, a morphism from $h : FA \rightarrow A$ to $g : FB \rightarrow B$ is a morphism $f : A \rightarrow B$ such that the following diagram commutes:

$$
\begin{array}{ccc}
FA & \xrightarrow{\;Ff\;} & FB \\
{\scriptstyle h}\downarrow & & \downarrow{\scriptstyle g} \\
A & \xrightarrow{\;f\;} & B
\end{array}
$$

We call such a morphism an *F-algebra morphism*. If the category of F-algebras has an initial object then Lambek's Lemma ensures that this *initial F-algebra* is an isomorphism, and thus that its carrier is a fixed point of F. Initiality ensures that the carrier of the initial F-algebra is actually a *least* fixed point of F. If it exists, the least fixed point for F is unique up to isomorphism. Henceforth we write μF for the least fixed point for F and $in : F(\mu F) \rightarrow \mu F$ for the initial F-algebra.

Within the paradigm of initial algebra semantics, every datatype is the carrier μF of the initial algebra of a suitable endofunctor F on a suitable category $\mathscr{C}$. The unique F-algebra morphism from *in* to any other F-algebra $h : FA \rightarrow A$ is given by the interpretation *fold* of the `fold` combinator for the interpretation μF of the datatype `Mu F`. The *fold* operator for μF thus makes the following commute:

$$
\begin{array}{ccc}
F(\mu F) & \xrightarrow{\;F(fold\,h)\;} & FA \\
{\scriptstyle in}\downarrow & & \downarrow{\scriptstyle h} \\
\mu F & \xrightarrow{\;fold\,h\;} & A
\end{array}
$$

From this diagram, we see that *fold* has type $(FA \rightarrow A) \rightarrow \mu F \rightarrow A$ and that *fold h* satisfies *fold h* (*in t*) = h (*F* (*fold h*) *t*). The uniqueness of the mediating map ensures that, for every F-algebra h, the map *fold h* is defined uniquely.

As shown in [5], the carrier of the initial algebra of an endofunctor F on $\mathscr{C}$ can be seen not only as the carrier of the initial F-algebra, but also as the limit of the forgetful functor $U_F : F\text{-}\mathscr{A}lg \rightarrow \mathscr{C}$ mapping each F-algebra $h : FA \rightarrow A$ to A.

If $G : \mathscr{C} \to \mathscr{D}$ is a functor, then a *cone* $\tau : D \to G$ to the base G with vertex D is an object D of $\mathscr{D}$ and a family of morphisms $\tau_C : D \to GC$, one for every object C of $\mathscr{C}$, such that for every arrow $\sigma : A \to B$ in $\mathscr{C}$, $\tau_B = G\sigma \circ \tau_A$ holds.

$$
\begin{array}{ccc}
GA & \xrightarrow{\;G\sigma\;} & GB \\
\big\uparrow\!\scriptstyle{\tau_A} & \nearrow\!\scriptstyle{\tau_B} & \\
D & &
\end{array}
$$

We usually refer to a cone simply by its family of morphisms, rather than the pair comprising the vertex together with the family of morphisms. A *limit* for $G : \mathscr{C} \to \mathscr{D}$ is an object $\lim G$ of $\mathscr{D}$ and a limiting cone $\nu : \lim G \to G$, i.e., a cone $\nu : \lim G \to G$ with the property that if $\tau : D \to G$ is any cone, then there is a unique morphism $\theta : D \to \lim G$ such that $\tau_C = \nu_C \circ \theta$ for all $C \in \mathscr{C}$.

$$
\begin{array}{ccc}
GA & \xrightarrow{\;F\sigma\;} & GB \\
\big\uparrow\!\scriptstyle{\tau_A} \quad \scriptstyle{\tau_B}\!\!\nwarrow\;\nearrow & & \big\uparrow\!\scriptstyle{\nu_B} \\
& \scriptstyle{\nu_A}\!\times & \\
D & \xrightarrow[\;\theta\;]{} & \lim G
\end{array}
$$

The characterization of μF as $\lim U_F$ provides a principled derivation of the interpretation *build* of the `build` combinator for μF which complements the derivation of its *fold* operator from standard initial algebra semantics. It also guarantees the correctness of the standard `fold`/`build` rules. Indeed, the universal property that the carrier μF of the initial F-algebra enjoys as $\lim U_F$ ensures:

- The projection from the limit μF to the carrier of each F-algebra defines the *fold* operator with type $(FA \to A) \to \mu F \to A$.

- Given a cone $\theta : C \to U_F$, the mediating morphism from it to the limiting cone $\nu : \lim U_F \to U_F$ defines a map from C to $\lim U_F$. Since a cone to U_F with vertex C has type $\forall x.(Fx \to x) \to C \to x$, this mediating morphism defines the *build* operator with type $(\forall x.\,(Fx \to x) \to C \to x) \to C \to \mu F$.

- The correctness of the `fold`/`build` fusion rule then follows from the fact that *fold* after *build* is a projection after a mediating morphism, and thus is equal to the cone applied to the specific algebra. Diagrammatically, we have

$$
\begin{array}{ccc}
& & A \\
& \scriptstyle{g\,k}\nearrow & \big\uparrow\!\scriptstyle{fold\,k} \\
C & \xrightarrow[\;build\,g\;]{} & \mu F
\end{array}
$$

8.3.2   Correctness of the `fold/superbuild` Rule

To prove correctness of our `fold/superbuild` rule we are actually interested in the following variation of the preceding diagram:

$$
\begin{array}{ccc}
 & H\,A & \\
\overset{g\,k}{\nearrow} & & \uparrow{\scriptstyle H(fold\,k)} \\
C \xrightarrow[superbuild\,g]{} & H\,(\mu F) &
\end{array}
$$

Here, *superbuild* is the interpretation in $\mathscr{C}$ of `superbuild`. If the functor $H : \mathscr{D} \to \mathscr{E}$ preserves limits — i.e., if, for every functor $G : \mathscr{C} \to \mathscr{D}$ and every limiting cone $\nu : \lim G \to G$, the cone $H\nu : H(\lim G) \to H \circ G$ is also a limit, henceforth denoted $\lim(H \circ G)$ — then this is the diagram for the universal property of $\lim(H \circ U_F)$. We thus ask which functors H preserve limits. It is well-known that right adjoints preserve limits, but this is a more restrictive class of functors than we would like. On the other hand, H needn't preserve *all* limits, just $\lim U_F$.

A *connected category* is a non-empty category whose underlying graph is connected. A *connected limit* is a limit of a functor whose domain is a connected category. The limit $\lim U_F : F\text{-}\mathscr{A}lg \to \mathscr{C}$ is a connected limit since the category of F-algebras is connected (there is a morphism from the initial F-algebra $in : F(\mu F) \to \mu F$ to any other F-algebra), so knowing that the functor H interpreting the type constructor `h` in the type of `superbuild` preserves connected limits is sufficient to ensure correctness of the `fold/superbuild` rule. It is well-known that strictly positive functors preserve connected limits [3, 7]; in particular, all polynomial functors preserve them. More generally, all functors created by containers preserve connected limits [7]. The class of containers includes functors, such as those whose least fixed points are nested types, which are not strictly positive; the above proof thus covers many situations that are interesting in practice. To prove correctness of the `fold/superbuild` rule for functors H which do not preserve connected limits, it should be possible to give a formal argument based on logical relations [1]. However, a proof based upon logical relations would not cover examples such such as nested types which preserve connected limits but are not definable in the underlying type theory of the logical relation.

8.3.3   Correctness of the `fold/msuperbuild` Rule

To see that the `fold/msuperbuild` rule is correct, we consider the diagram

$$
\begin{array}{ccc}
 & M\,(MA) \xrightarrow{\;id^*\;} MA & \\
\overset{g\,k}{\nearrow} & \Big\uparrow{\scriptstyle M(fold\,k)} \quad \nearrow{\scriptstyle (fold\,k)^*} & \\
C \xrightarrow[msuperbuild\,g]{} & M\,(\mu F) &
\end{array}
$$

where M is the interpretation of `m` in the type of `msuperbuild`, *bind* and *return* are the interpretations of the `>>=` and `return` operations for `m`, respectively, and

$f^* x = bind\, x\, f$. Correctness of the `fold`/`msuperbuild` rule is exactly commutativity of the diagram's outer parallelogram. The diagram's left-hand triangle commutes because it is an instance of the previous diagram, and standard properties of monads ensure that its right-hand side commutes as well. Then

$$
\begin{aligned}
g\, k\, c \gg= id \quad &= \quad id^* \,(g\, k\, c) \\
&= \quad (id^* \circ g\, k)\, c \\
&= \quad ((\textit{fold}\, k)^* \circ \textit{msuperbuild}\, g)\, c \\
&= \quad (\textit{fold}\, k)^* (\textit{msuperbuild}\, g\, c) \\
&= \quad \textit{msuperbuild}\, g\, c \gg= \textit{fold}\, k
\end{aligned}
$$

It is worth noting here that many monads that arise in applications — including the exceptions monad, the state monad, and the list monad — preserve connected limits. The continuations monad, however, does not.

8.4 DUALITY

Our `fold`/`superbuild` and `fold`/`msuperbuild` rules dualize to the coinductive setting. Shortage of space prevents us from giving the corresponding constructs and results in detail here, so we simply present their implementation. We have

```
unfold :: Functor f => (a -> f a) -> a -> Mu f
unfold k x = In (fmap (unfold k) (k x))

superdestroy :: (Functor f, Functor h) =>
   (forall a. (a -> f a) -> h a -> c) -> h (Mu f) -> c
superdestroy g = g unIn

superdestroy g . fmap (unfold k) = g k
```

When `c` is `Mu f`, `superdestroy` returns an `h`-algebra which stores coalgebraic `f`-data. When `h` is a comonad, i.e., an instance of the `Comonad` class

```
class Comonad cm where
 coreturn :: cm a -> a
 (=<<)    :: cm b -> (cm b -> a) -> cm a
```

we have

```
cmsuperdestroy :: (Functor f, Comonad cm) =>
   (forall a. (a -> f a) -> cm a -> c) -> cm (Mu f) -> c
cmsuperdestroy g = g unIn

cmsuperdestroy g (x =<< unfold k) = g k (x =<< id)
```

8.5 RELATED WORK

The work most closely related to ours is that of Pardo and his coauthors. Like this paper, [14] also investigates conditions under which the composition of a function producing an expression of type $M(\mu F)$ for M a monad and F a functor, and a

function *fold k* of type $\mu F \rightarrow A$ can be fused to produce an expression of type MA. But there are several crucial differences with our work. First, Pardo uses `unfold` rather than `msuperbuild` to construct the intermediate expression. This gives his fusion rule some additional logical generality over ours, since `unfold` can construct elements of its associated functor `f`'s final coalgebra which are not in `f`'s initial algebra, whereas `msuperbuild` can construct only elements of `f`'s initial algebra. But when the initial and final algebras of each functor coincide, as in Haskell, this added logical generality yields no advantage in practice.

Secondly, Pardo's monadic hylofusion (and hylofusion in general) is only known to be correct in algebraically compact categories, i.e., categories in which the initial algebra and final coalgebra for each functor coincide. By contrast, our `fold`/`superbuild` rule is correct in any category supporting a parametric interpretation of `forall`, and this condition is independent of any compactness condition. The requirement that the interpreting category be algebraically compact is unfortunate since it generates strictness conditions that must be satisfied, and also requires the underlying monad to be strictness-preserving. This results in strictness condition propagation. By contrast, neither our `fold`/`superbuild` nor our `fold`/`msuperbuild` rules require the satisfaction of side conditions.

Thirdly, Pardo trades a composition of an `unfold` and a monadic `fold` for the computation of an equivalent fixed point. By contrast, our `fold`/`msuperbuild` rule trades a bind of a call to `msuperbuild` with a monadic `fold` for the bind of the application of the function argument to `msuperbuild` to the `fold`'s algebra with the identity function. Like all generalizations of the `fold`/`build` rule, our `fold`/`msuperbuild` rule requires "payment up front" in that the producer in a composition to be fused must be expressed in terms of `msuperbuild`. (This is not very different from the price paid by expressing consumers in terms of `unfold`). But our rule delivers a fused result which is simpler than that obtained using Pardo's technique. In particular, the functions obtained from our fusion rules involve only binds of applications involving data structure "templates", rather than fixed point calculations. Their computation is thus guaranteed to terminate.

Finally, Pardo requires the existence of a distributivity law of the underlying monad over the underlying functor to construct the lifting of functors to the Kleisli category on which his monadic hylofusion rule depends. But distributivity laws for arbitrary functors, even those admitting fixed points, need not exist.

Recently, Manzino and Pardo [11] have proposed a fusion rule similar to the `fold`/`msuperbuild` rule given here. This rule seems to be interderivable with ours in the presence of distributivity. Meijer and Jeuring [12] have also developed fusion laws in the monadic setting, including a short cut fusion law for eliminating intermediate structures of type FA in a monadic context M. Many fusion methods, including those of [12] and [14], eliminate data structures in the carriers of initial algebras for only restricted classes of functors. By contrast, our method can eliminate data structures of *any* inductive type, and can handle non-monadic contexts as well. In addition, Jürgensen [10] and Voigtländer [16] have each defined fusion combinators based on the uniqueness of the map from a free monad to any other monad. These techniques give very different forms of fusion from ours.

8.6 CONCLUSION AND DIRECTIONS FOR FUTURE WORK

In this paper we defined a `superbuild` combinator which generalizes the standard `build` combinator and expresses uniform production of functorial contexts containing data of inductive types. We also proved correct a `fold/super build` fusion rule which generalizes the `fold/build` and `fold/buildp` rules from the literature, and eliminates intermediate data structures of inductive types without disturbing the contexts in which they are situated. An important special case arises when this context is monadic. When it is, our `fold/msuperbuild` rule fuses combinations of producers and consumers via monad operations, rather than via composition. We have given examples illustrating both the `fold/superbuild` and `fold/msuperbuild` rules, and considered their coalgebraic duals as well.

The standard `fold` combinator can consume data structures in any context describable by a functor, but the algebra it uses cannot depend on the context in a non-trivial way. By contrast, context information can be used by algebras to partially determine how the `pfold` combinator given in [2] will consume the data structures, but unfortunately the contexts are limited to pairs. Interestingly, the `pfold/buildp` rule given there for context-dependent `fold`s derives from the `fold/buildp` rule from Figure 8.2 for standard `fold`s. As already noted, it is the `fold/buildp` rule that our `fold/superbuild` and `fold/msuperbuild` rules generalize. One direction for future work is to generalize these rules even further to accommodate *both* context-dependent algebras and non-pair contexts.

Another direction for future work is suggested by considering an even more monadic fusion rule based on `fold-` and `build`-like combinators which manipulate algebra-like functions of type `f a -> m a`. Such a rule would produce intermediate data structures using "templates" based on so-called monadic algebras and, in the presence of a distributivity rule `delta` for `m` over `f`, would consume data structures using them via a monadic `mafold` combinator. We'd have

```
mafold :: (Functor f, Monad m) =>
                             (f a -> m a) -> Mu f -> m a
mafold k = fold (\x -> fmap k (delta x) >>= id)

masuperbuild :: (Functor f, Monad m) =>
   (forall a. (f a -> m a) -> c -> m a) -> c -> m (Mu f)
masuperbuild g = g (return . In)

masuperbuild c >>= mafold k = g k c
```

Although a datatype-generic `masuperbuild` combinator is not defined in [12], several instances of the above fusion rule are given. Yet no correctness proofs for any of these specific instances — let alone any formulation of, or correctness proof for, a datatype-generic fusion rule — are given. We believe an independent proof of the `mafold/masuperbuild` rule similar to those in Section 8.3 is possible. Although it is not entirely clear how such a proof would go, a proof for monads which preserve connected limits will likely require independent verification that $\lim(M U_{F,M}) = M(\mu F)$ for the forgetful functor $U_{F,M}$ mapping each

monadic algebra $h : F\,a \to M\,a$ to a, and a proof for monads which do not preserve connected limits will likely be based on logical relations.

The facts that `mafold` is defined in terms of `fold` and that `masuperbuild g` can be expressed as `msuperbuild (\k -> g (return . k))` together suggest that the `mafold`/`masuperbuild` rule might be derivable from (distributivity and) the `fold`/`msuperbuild` rule. However, we believe the two rules to offer distinct fusion options in the presence of distributivity; it would be interesting to see which is more useful for programs that arise in practice.

REFERENCES

[1] T. Altenkirch, P. Levy, and M. Hasagawa, 2008. Personal communication, and message 1192 from the TYPES mailing list archive.

[2] J. P. Fernandes, A. Pardo, and J. Saraiva. A shortcut fusion rule for circular program calculation. In *Proceedings, Haskell Workshop*, pages 95–106, 2007.

[3] N. Ghani, M. Abbott, and T. Altenkirch. Containers - constructing strictly positive types. *Theoretical Computer Science*, 341(1):3–27, 2005.

[4] N. Ghani, P. Johann, T. Uustalu, and V. Vene. Monadic augment and generalised short cut fusion. In *Proceedings, International Conference on Functional Programming*, pages 294–305, 2005.

[5] N. Ghani, T. Uustalu, and V. Vene. Build, augment and destroy. Universally. In *Proceedings, Asian Symposium on Programming Languages*, pages 327–347, 2003.

[6] A. Gill, J. Launchbury, and S. L. Peyton Jones. A short cut to deforestation. In *Proceedings, Functional Programming Languages and Computer Architecture*, pages 223–232, 1993.

[7] R. Hasegawa. Two applications of analytic functors. *Theoretical Computer Science*, 272(1-2):113–175, 2002.

[8] P. Johann and N. Ghani. Initial algebra semantics is enough! In *Proceedings, Typed Lambda Calculus and Applications*, pages 207–222, 2007.

[9] P. Johann and N. Ghani. Foundations for structured programming with GADTs. In *Proceedings, Principles of Programming Languages*, pages 297–308, 2008.

[10] C. Jürgensen. Using monads to fuse recursive programs (extended abstract), 2002.

[11] C. Manzino and Pardo. Shortcut fusion of monadic programs. Presented at *Simposio Brasiliero de Linguagens de Programacao*, 2008.

[12] E. Meijer and J. Juering. Merging monads and folds for functional programming. In *Proceedings, Advanced Functional Programming*, pages 228–266, 1995.

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

[14] A. Pardo. Fusion of recursive programs with computational effects. *Theoretical Computer Science*, 260(1-2):165–207, 2001.

[15] S. L. Peyton Jones, editor. *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press, 2003.

[16] J. Voigtländer. Asymptotic improvement of computations over free monads. In *Proceedings, Mathematics of Program Construction*, pages 388–403, 2008.