

Factorising Folds for Faster Functions

GRAHAM HUTTON
University of Nottingham, UK

MAURO JASKELIOFF
Universidad Nacional de Rosario, Argentina

ANDY GILL
University of Kansas, USA

Abstract

The worker/wrapper transformation is a general technique for improving the performance of recursive programs by changing their types. The previous formalisation (Gill & Hutton, 2009) was based upon a simple fixed point semantics of recursion. In this article we develop a more structured approach, based upon initial algebra semantics. In particular, we show how the worker/wrapper transformation can be applied to programs defined using the structured pattern of recursion captured by fold operators, and illustrate our new technique with a number of examples.

1 Introduction

The worker/wrapper transformation is a general technique for changing the type of a recursive program to improve its performance. The basic idea is simple and pervasive: given a recursive program of some type, we aim to factorise it into a more efficient *worker* program of a different type, together with a *wrapper* program that acts as an interface between the original program and the new worker.

Special cases of the worker/wrapper transformation have been used for many years, particularly in optimizing compilers. For example, the technique has been used in the Glasgow Haskell Compiler since its inception, to replace the use of boxed data structures by more efficient unboxed data structures when safe to do so (Peyton Jones & Launchbury, 1991). However, it is only recently that the transformation has been formalised, proved correct, and presented as a general technique for improving the performance of programs by improving the choice of data structures (Gill & Hutton, 2009).

The previous formalisation was based upon a simple fixed point semantics of recursive programs. In this article we take a more structured approach, based upon initial algebra semantics. In particular, we develop a general worker/wrapper theory for changing the type of recursive programs defined using fold operators, and show how it can be used in practice as an equational reasoning technique for improving the performance of programs. More precisely, the article makes the following contributions:

- We show how the worker/wrapper transformation applies to programs defined using folds, by generalising to a categorical view of types as initial algebras.

- We identify four conditions for the correctness of the transformation, and show that these conditions form a simple lattice structure.
- We illustrate our technique with a number of examples, including a correctness proof for a new approach to implementing substitution efficiently (Voigtländer, 2008).

The use of initial algebras also means that our worker/wrapper technique for folds is generic in the underlying recursive type to which it applies (Backhouse *et al.*, 1999). That is, the technique is defined and proved once, for an arbitrary recursive type, and can then simply be instantiated as required for each new type.

The article is aimed at readers who are familiar with the basics of initial algebra semantics (in particular, the concepts of categories, functors, products, co-products, and initial algebras), say to the level of chapter two of Bird & de Moor (1997), but no previous experience with the worker/wrapper transformation is assumed. An extended version of the article that includes all the proofs is available from the authors' web pages.

2 Initial algebra semantics

The recursion operator *fold* encapsulates a common pattern for defining functions that process values of a recursively defined type (Hutton, 1999). In this section we review the categorical treatment of *fold*, and introduce our notation. For further details, see for example (Malcolm, 1990; Meijer *et al.*, 1991; Bird & de Moor, 1997).

Suppose that we fix a category \mathbf{C} and a functor $F : \mathbf{C} \rightarrow \mathbf{C}$ on this category. Then the notion of an *algebra* is defined as a pair (A, f) comprising an object A and an arrow $f : FA \rightarrow A$. In turn, a *homomorphism* $h : (A, f) \rightarrow (B, g)$ from one such algebra to another is an arrow $h : A \rightarrow B$ such that the following diagram commutes:

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

Algebras and homomorphisms themselves form a category, with composition and identities inherited from \mathbf{C} . An *initial algebra* is an initial object in this new category, and we write $(\mu F, in)$ for an initial algebra, and *fold* f for the unique homomorphism $h : (\mu F, in) \rightarrow (A, f)$ from the initial algebra to any other algebra (A, f) . That is, *fold* f is defined as the unique arrow that makes the following diagram commute:

$$\begin{array}{ccc} F\mu F & \xrightarrow{F(\text{fold } f)} & FA \\ \downarrow in & & \downarrow f \\ \mu F & \xrightarrow{\text{fold } f} & A \end{array}$$

In the literature, *fold* f is sometimes written using the banana brackets notation $(\lfloor f \rfloor)$, and termed a *catamorphism*. The above definition for *fold* f can also be expressed as the

following equivalence, known as the *universal property of fold*:

$$h = \text{fold } f \Leftrightarrow h \circ \text{in} = f \circ Fh$$

The \Rightarrow direction states that $\text{fold } f$ is a homomorphism from the initial algebra $(\mu F, \text{in})$ to another algebra (A, f) , while the \Leftarrow direction states that any other such homomorphism h must be equal to $\text{fold } f$. Taken as a whole, the universal property expresses in an equational manner the fact that $\text{fold } f$ is the unique homomorphism from $(\mu F, \text{in})$ to (A, f) .

The universal property can be used to verify the well-known *fusion* property of *fold*, which states that the composition of a function and a *fold* can always be re-expressed as a single *fold*, provided the function is a homomorphism of the appropriate type:

$$h \circ f = g \circ Fh \Rightarrow h \circ \text{fold } f = \text{fold } g$$

As a simple example of initial algebra semantics, suppose that we define a functor F on the category **SET** by $FA = 1 + A$. Then F has an initial algebra, given by the set \mathbb{N} of natural numbers, together with a function $[\text{zero}, \text{succ}] : 1 + \mathbb{N} \rightarrow \mathbb{N}$ comprising two constructors $\text{zero} : 1 \rightarrow \mathbb{N}$ and $\text{succ} : \mathbb{N} \rightarrow \mathbb{N}$ for this set. In turn, given any other set A and functions $v : 1 \rightarrow A$ and $f : A \rightarrow A$, the function $\text{fold } [v, f] : \mathbb{N} \rightarrow A$ is uniquely defined by:

$$\begin{aligned} h(\text{zero } ()) &= v() \\ h(\text{succ } n) &= f(hn) \end{aligned}$$

That is, $\text{fold } [v, f]$ processes a natural number by replacing the *zero* constructor by the function v , and each *succ* constructor by the function f . For example, a doubling function can be defined by $\text{double} = \text{fold } [\text{zero}, \text{succ} \circ \text{succ}]$, and fusion can then be used to show that $\text{double} \circ \text{double} = \text{fold } [\text{zero}, \text{succ} \circ \text{succ} \circ \text{succ} \circ \text{succ}]$.

3 Worker/wrapper for folds

Consider the problem of changing the return type of a *fold* to improve its performance. More precisely, suppose we are given a function $\text{fold } f : \mu F \rightarrow A$ for some $f : FA \rightarrow A$, and we wish to change the return type from A to some other type B . The worker/wrapper approach to this problem is based upon the use of conversion functions


$$\begin{array}{ccc} & \xrightarrow{\text{rep}} & \\ A & \xleftrightarrow{\hspace{1cm}} & B \\ & \xleftarrow{\text{abs}} & \end{array}$$

with the property that

$$\text{abs} \circ \text{rep} = \text{id}_A$$


This equation states that abs is a left inverse (or retraction) of rep , or in more practical terms that converting a value of the original type into the new type and then back again does not change the value. In the terminology of data representation (Hoare, 1972), this means that the *abstract* type A can be faithfully represented by the *concrete* type B . For example, in the case of the category **SET**, the equation ensures that the set A is isomorphic to the subset of B given by the image of rep . Given the above assumptions, we now seek

conditions under which the following diagram commutes:



That is, in worker/wrapper terminology, we seek conditions that allow the original recursive function *fold f* that produces a result of type *A* to be factorised as the composition of a recursive *worker* function *fold g* that produces a result of type *B*, and a *wrapper* function *abs* that converts the result back to the original type *A*.


One approach to solving this problem is to simply apply fusion. Even though this property is normally viewed as being concerned with combining a function with a *fold*, it can also be viewed in the opposite direction as providing a sufficient condition for the factorisation or *fission* (Gibbons, 2006) of a *fold* in the manner above, namely that $f \circ F \text{ abs} = \text{abs} \circ g$. However, given the assumption that $\text{abs} \circ \text{rep} = \text{id}_A$, we can in fact identify four relevant conditions, given by the four possible ways of completing the following commuting diagram that relates the argument algebras *f* and *g*



by replacing each ? in the diagram with either $\text{rep} : A \rightarrow B$ or $\text{abs} : B \leftarrow A$:

- (1) $g = \text{rep} \circ f \circ F \text{ abs}$
- (2) $\text{rep} \circ f = g \circ F \text{ rep}$
- (3) $f \circ F \text{ abs} = \text{abs} \circ g$
- (4) $f = \text{abs} \circ g \circ F \text{ rep}$

What do these conditions express, and how do they relate? Equation (1) provides an explicit definition for *g* in terms of *f*; (2) states that *rep* is a homomorphism from *f* to *g*; (3) states that *abs* is a homomorphism from *g* to *f* (the condition that directly arises from the use of fusion); and (4) provides a definition for *f* in terms of *g*. Together, they form a simple lattice, with (1) as the strongest condition and (4) as the weakest:



It is now straightforward to verify that each of the first three conditions implies the desired factorisation result, namely that $\text{fold } f = \text{abs} \circ \text{fold } g$. The situation regarding (4) is more involved, and we will return to this shortly. In the meantime, let us consider how the first three conditions are used in practice.

For some applications, the definition for the function g that forms the body of the worker $\text{fold } g$ will already be given, and our aim then is to *verify* that one of the three conditions is satisfied, to ensure that the worker/wrapper factorisation holds. For many applications, however, our aim will be to *construct* a suitable function g . In such cases, condition (1) provides an explicit definition $g = \text{rep} \circ f \circ F \text{abs}$ for the body of the worker in a similar manner to (Gill & Hutton, 2009), and our aim then is to *simplify* the definition. This simplification process is typically driven by the desire to fuse together instances of rep and abs , to eliminate the overhead of repeatedly converting between the concrete and abstract types. In contrast, conditions (2) and (3) provide a specification for g , and our aim is then to *calculate* a definition that satisfies the specification, again with the desire to fuse together instances of the conversion functions between the two types.

Given that (1) is the strongest condition and provides an explicit definition for g as a starting point, why would we ever wish to use the other conditions? In our experience, using one of the weaker conditions often results in a simpler verification or calculation process. In combination with the fact that (3) corresponds to the familiar case of fusion, for the purposes of examples we will primarily focus on (2). Nonetheless, it is interesting to consider the other conditions, and their relationships.

Let us now return to the remaining condition in our lattice:

$$(4) \quad f = \text{abs} \circ g \circ F \text{rep}$$

Unfortunately, in general this condition does not imply that $\text{fold } f = \text{abs} \circ \text{fold } g$, and is only sufficient to ensure the following more specialised worker/wrapper factorisation in which the body g of the worker is composed with an additional term:

$$\text{fold } f = \text{abs} \circ \text{fold } (g \circ F (\text{rep} \circ \text{abs}))$$

The additional term $F (\text{rep} \circ \text{abs})$ in the worker plays the role of a *normalisation* function that is applied after each recursive call. In general, $\text{rep} \circ \text{abs} \neq \text{id}_B$, but we can think of $\text{rep} \circ \text{abs}$ as normalising a value of type B by first converting to the type A , which is typically a ‘smaller’ type, and then converting back to B . It is natural to ask when (4) does imply $\text{fold } f = \text{abs} \circ \text{fold } g$. The answer is given by the following condition, which states that $\text{rep} \circ \text{abs}$ is a homomorphism from g to itself:

$$(5) \quad \text{rep} \circ \text{abs} \circ g = g \circ F (\text{rep} \circ \text{abs})$$

In particular, we then have the following equivalence:

$$(4) \wedge (5) \Leftrightarrow (2) \wedge (3)$$

That is, the combination of (4) and (5) is equivalent to the combination of (2) and (3), either condition of which implies the worker/wrapper factorisation.

We conclude this section by noting that condition (5) also implies the following property, which is precisely the worker/wrapper fusion property from (Gill & Hutton, 2009) for the

into our worker/wrapper paradigm based upon *fold*, and leads to a simpler derivation than the previous worker/wrapper approach based upon *fix*.

Using explicit recursion, a reverse function can be defined by

$$\begin{aligned} rev &:: [a] \rightarrow [a] \\ rev [] &= [] \\ rev (x : xs) &= rev xs ++ [x] \end{aligned}$$

or equivalently, using the *fold* operator for lists:

$$\begin{aligned} rev &:: [a] \rightarrow [a] \\ rev &= fold snoc [] \\ snoc &:: a \rightarrow [a] \rightarrow [a] \\ snoc x xs &= xs ++ [x] \end{aligned}$$

However, because of the use of append (*++*), this definition for *rev* takes quadratic time. We now show how our worker/wrapper technique for *fold* can be used to derive a more efficient worker that uses an extra argument to accumulate the result, together with a wrapper that takes care of the initial setup. Using the notion of currying, the introduction of an accumulator argument corresponds to changing the return type of *rev* from a list to a function on lists, i.e. changing from the original return type $[a]$ to the new return type $[a] \rightarrow [a]$. The necessary conversion functions between the two types, the latter of which is sometimes called *Hughes lists* (Hughes, 1986), are defined as follows:

$$\begin{aligned} \mathbf{type} \ H \ a &= [a] \rightarrow [a] \\ rep &:: [a] \rightarrow H \ a \\ rep \ xs &= (xs \ ++) \\ abs &:: H \ a \rightarrow [a] \\ abs \ h &= h \ [] \end{aligned}$$

Note that *rep* is just a synonym for (*++*). It is straightforward to verify the worker/wrapper assumption $abs \circ rep = id_{[a]}$. We also have the important property that *rep* forms a monoid homomorphism from lists to Hughes lists, in the sense that:

$$\begin{aligned} rep \ (xs \ ++ \ ys) &= rep \ xs \circ rep \ ys \\ rep \ [] &= id_{[a]} \end{aligned}$$

In the case of reverse, it turns out that the most convenient condition to use as the basis for constructing the worker function is condition (2):

$$rep \ (snoc \ x \ xs) = g \ x \ (rep \ xs)$$

We calculate a function *g* satisfying this equation as follows:

$$\begin{aligned} &rep \ (snoc \ x \ xs) \\ = &\quad \{ \text{applying } snoc \} \\ &rep \ (xs \ ++ \ [x]) \\ = &\quad \{ rep \ \text{is a homomorphism} \} \\ &rep \ xs \circ rep \ [x] \end{aligned}$$

$$\begin{aligned}
&= \{ \text{applying } \mathit{rep} \} \\
&\quad \mathit{rep} \, \mathit{xs} \circ (x:) \\
&= \{ \text{define } g \, x \, h = h \circ (x:) \} \\
&\quad g \, x \, (\mathit{rep} \, \mathit{xs})
\end{aligned}$$

Now that we have satisfied the necessary preconditions, applying the worker/wrapper transformation for *fold* gives the following new definitions:

$$\begin{aligned}
\mathit{rev} &:: [a] \rightarrow [a] \\
\mathit{rev} &= \mathit{abs} \circ \mathit{work} \\
\mathit{work} &:: [a] \rightarrow H \, a \\
\mathit{work} &= \mathit{fold} \, g \, (\mathit{rep} \, [])
\end{aligned}$$

Finally, if we make the list arguments explicit, and expand out the component functions, we obtain the expected linear time version of reverse that uses an accumulator:

$$\begin{aligned}
\mathit{rev} &:: [a] \rightarrow [a] \\
\mathit{rev} \, \mathit{xs} &= \mathit{work} \, \mathit{xs} \, [] \\
\mathit{work} &:: [a] \rightarrow [a] \rightarrow [a] \\
\mathit{work} \, [] \, \mathit{ys} &= \mathit{ys} \\
\mathit{work} \, (x:\mathit{xs}) \, \mathit{ys} &= \mathit{work} \, \mathit{xs} \, (x:\mathit{ys})
\end{aligned}$$

We conclude with a number of observations about the above derivation. First of all, in common with the previous derivation of fast reverse using the worker/wrapper technique for *fix* (Gill & Hutton, 2009), once we have made the decision to use Hughes' representation of lists, the rest of the derivation proceeds using simple equational reasoning, without the need for induction. However, in contrast to the previous derivation, the additional structure made explicit by using *fold* avoids the need for the additional functions *wrap* and *unwrap*, the use of worker/wrapper fusion, and the need to expand out the worker as an essential step in the derivation, resulting in a simpler derivation.

4.2 Example: fast reverse revisited

It is interesting now to return to our earlier question of why we don't always use condition (1), which provides an explicit definition for *g* as a starting point. In the case of the reverse example, the initial definition would then be as follows:

$$g \, x \, y = \mathit{rep} \, (\mathit{snoc} \, x \, (\mathit{abs} \, y))$$

The problem comes when we try and simplify this definition:

$$\begin{aligned}
&g \, x \, y \\
&= \{ \text{applying } g \} \\
&\quad \mathit{rep} \, (\mathit{snoc} \, x \, (\mathit{abs} \, y)) \\
&= \{ \text{applying } \mathit{snoc} \} \\
&\quad \mathit{rep} \, (\mathit{abs} \, y \, ++ [x]) \\
&= \{ \mathit{rep} \text{ is a homomorphism} \} \\
&\quad \mathit{rep} \, (\mathit{abs} \, y) \circ \mathit{rep} \, [x]
\end{aligned}$$

$$= \quad \{ \text{applying } \mathit{rep} \} \\ \mathit{rep} (\mathit{abs} \ y) \circ (x:)$$

Now we appear to be stuck. We would like to fuse together rep and abs in the final expression to give the definition $g \ x \ y = y \circ (x:)$, but unfortunately it is not the case that $\mathit{rep} \circ \mathit{abs} = \mathit{id}_{Ha}$. In order to make progress, we begin by rewriting the worker

$$\mathit{work} = \mathit{fold} \ g \ (\mathit{rep} \ [])$$

by making the first list argument explicit, expanding out the fold , and using the above simplification of g to give the following definition using explicit recursion:

$$\begin{aligned} \mathit{work} \ [] &= \mathit{rep} \ [] \\ \mathit{work} \ (x:xs) &= \mathit{rep} \ (\mathit{abs} \ (\mathit{work} \ xs)) \circ (x:) \end{aligned}$$

While $\mathit{rep} \circ \mathit{abs} = \mathit{id}_{Ha}$ is not true in general, for the special case of values produced by worker itself we do have $\mathit{rep} \circ \mathit{abs} \circ \mathit{work} = \mathit{work}$, the worker/wrapper fusion property (6), which allows use to rewrite the worker as

$$\begin{aligned} \mathit{work} \ [] &= \mathit{rep} \ [] \\ \mathit{work} \ (x:xs) &= \mathit{work} \ xs \circ (x:) \end{aligned}$$

which can then be expanded to give the expected definition:

$$\begin{aligned} \mathit{work} \ [] \ ys &= \ ys \\ \mathit{work} \ (x:xs) \ ys &= \mathit{work} \ xs \ (y:ys) \end{aligned}$$

However, an unsatisfactory aspect of the above derivation is the need to rewrite the worker using explicit recursion in order to make progress by applying worker/wrapper fusion. Can the derivation also be performed at the fold level, without expanding out the recursion? The key to achieving this is to observe that in this context the second argument of g will always be of the form $\mathit{rep} \ z$ for some list z , since both the base and recursive case for the worker have an application of rep at the outer level. Using this assumption, the definition for g can then be simplified as follows:

$$\begin{aligned} &g \ x \ y \\ = &\quad \{ \text{previous simplification} \} \\ &\mathit{rep} \ (\mathit{abs} \ y) \circ (x:) \\ = &\quad \{ \text{assuming } y = \mathit{rep} \ z \} \\ &\mathit{rep} \ (\mathit{abs} \ (\mathit{rep} \ z)) \circ (x:) \\ = &\quad \{ \mathit{abs} \circ \mathit{rep} = \mathit{id} \} \\ &\mathit{rep} \ z \circ (x:) \\ = &\quad \{ \text{assuming } y = \mathit{rep} \ z \} \\ &y \circ (x:) \end{aligned}$$

Avoiding the need for this kind of ad-hoc additional reasoning is precisely the benefit that we obtain by starting from condition (2) rather than (1). In particular, using $\mathit{rep} \ (f \ x \ y) = g \ x \ (\mathit{rep} \ y)$ as our specification for g makes *explicit* from the outset that we can assume the second argument to g is always of the form $\mathit{rep} \ y$.

5 Worker/wrapper for expressions

For our next example we move from the type of lists to a simple language of expressions comprising integers and addition, together with its associated *fold* operator:

$$\begin{aligned} \mathbf{data} \text{ Expr} &= \text{Val Int} \mid \text{Add Expr Expr} \\ \text{fold} &:: (a \rightarrow a \rightarrow a) \rightarrow (\text{Int} \rightarrow a) \rightarrow \text{Expr} \rightarrow a \\ \text{fold } f \ v \ (\text{Val } n) &= v \ n \\ \text{fold } f \ v \ (\text{Add } x \ y) &= f \ (\text{fold } f \ v \ x) \ (\text{fold } f \ v \ y) \end{aligned}$$

Now suppose that we wish to change the return type of a function $\text{fold } f \ v :: \text{Expr} \rightarrow a$ from the original type a to some other type b , and that we are given conversion functions $\text{rep} :: a \rightarrow b$ and $\text{abs} :: b \rightarrow a$ such that $\text{abs} \circ \text{rep} = \text{id}_a$. In this context, our general worker/wrapper theory states that any of the three conditions

$$(1) \quad g \ x \ y = \text{rep} \ (f \ (\text{abs } x) \ (\text{abs } y))$$

$$(2) \quad \text{rep} \ (f \ x \ y) = g \ (\text{rep } x) \ (\text{rep } y)$$

$$(3) \quad f \ (\text{abs } x) \ (\text{abs } y) = \text{abs} \ (g \ x \ y)$$

is sufficient to justify the following factorisation of the original *fold* that produces a result of type a into the composition of a worker *fold* that produces a result of type b , and a wrapper function that converts the result back to the original type a :

$$\text{fold } f \ v = \text{abs} \circ \text{fold } g \ (\text{rep} \circ v)$$

Now consider the problem of transforming an evaluator for expressions into continuation-passing style, the typical first step in deriving an efficient abstract machine (Hutton & Wright, 2006). Using explicit recursion, an evaluation function can be defined by

$$\begin{aligned} \text{eval} &:: \text{Expr} \rightarrow \text{Int} \\ \text{eval} \ (\text{Val } n) &= n \\ \text{eval} \ (\text{Add } x \ y) &= \text{eval } x + \text{eval } y \end{aligned}$$

or equivalently, using the *fold* operator for expressions:

$$\begin{aligned} \text{eval} &:: \text{Expr} \rightarrow \text{Int} \\ \text{eval} &= \text{fold} \ (+) \ \text{id} \end{aligned}$$

Rewriting this definition in continuation-passing style involves taking a function on integers (the continuation) as an extra argument, which using currying corresponds to changing from the original return type Int to the new return type $(\text{Int} \rightarrow \text{Int}) \rightarrow \text{Int}$. The necessary conversion functions between the two types are defined as follows:

$$\begin{aligned} \mathbf{type} \ \text{Cint} &= (\text{Int} \rightarrow \text{Int}) \rightarrow \text{Int} \\ \text{rep} &:: \text{Int} \rightarrow \text{Cint} \\ \text{rep } n &= \lambda c \rightarrow c \ n \\ \text{abs} &:: \text{Cint} \rightarrow \text{Int} \\ \text{abs } f &= f \ \text{id} \end{aligned}$$

It is easy to show that $abs \circ rep = id_{Int}$. As with fast reverse, the appropriate starting point for constructing the worker in this case is condition (2),

$$rep (x + y) = g (rep x) (rep y)$$

from which we calculate a function g satisfying this equation as follows:

$$\begin{aligned} & rep (x + y) c \\ = & \quad \{ \text{applying } rep \} \\ & c (x + y) \\ = & \quad \{ \text{abstracting over } x \} \\ & (\lambda n \rightarrow c (n + y)) x \\ = & \quad \{ \text{unapplying } rep \} \\ & rep x (\lambda n \rightarrow c (n + y)) \\ = & \quad \{ \text{abstracting over } y \} \\ & rep x (\lambda n \rightarrow (\lambda m \rightarrow c (n + m)) y) \\ = & \quad \{ \text{unapplying } rep \} \\ & rep x (\lambda n \rightarrow rep y (\lambda m \rightarrow c (n + m))) \\ = & \quad \{ \text{define } g \ a \ b = a (\lambda n \rightarrow b (\lambda m \rightarrow c (n + m))) \} \\ & g (rep x) (rep y) \end{aligned}$$

Now that we have satisfied the necessary preconditions, applying the worker/wrapper transformation for *fold* gives the following definitions

$$\begin{aligned} eval & \quad :: \ Expr \rightarrow Int \\ eval & \quad = \ abs \circ work \\ work & \quad :: \ Expr \rightarrow Cint \\ work & \quad = \ fold \ g \ (rep \circ id) \end{aligned}$$

which expand out to give the expected continuation-passing evaluator:

$$\begin{aligned} eval & \quad :: \ Expr \rightarrow Int \\ eval \ e & \quad = \ work \ e \ id \\ work & \quad :: \ Expr \rightarrow (Int \rightarrow Int) \rightarrow Int \\ work \ (Val \ n) \ c & \quad = \ c \ n \\ work \ (Add \ x \ y) \ c & \quad = \ work \ x \ (\lambda n \rightarrow work \ y \ (\lambda m \rightarrow c \ (n + m))) \end{aligned}$$

Once again, note that the derivation proceeds using simple equational reasoning and does not require induction. Moreover, in contrast to our previous derivation of such an evaluator using more elementary techniques (Hutton & Wright, 2006), using worker/wrapper condition (2) as the starting point results in a derivation whose goal is made explicit from the outset, namely to construct a function g such that $rep (x + y) = g (rep x) (rep y)$, rather than this property being implicit in the structure of the derivation itself.

6 Efficient substitution

For our final example, we consider a more challenging problem: improving the performance of monadic substitution on trees. The example is taken from (Voigtländer, 2008),

but whereas the author only sketches a proof of correctness and conjectures that a formal proof may require sophisticated techniques, we show that a simple proof is possible using our worker/wrapper technique for *fold*. We begin by generalising the type *Expr* from the previous section to the type *Tree a* of binary trees with leaves of type *a*:

```
data Tree a = Leaf a | Node (Tree a) (Tree a)
```

Now recall that in Haskell, the categorical notion of a *monad* is captured by the following class declaration, which states that a parameterised type *m* is a member of the class *Monad* of monadic types if it is equipped with *return* and *>>=* functions of the specified types:

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

The two functions must also satisfy identity and associativity properties:

```
return x >>= f = f x
e >>= return  = e
(e >>= f) >>= g = e >>= (λx → f x >>= g)
```

It is straightforward to make *Tree* into a monadic type by the following instance declaration, and to verify that the required monad laws are satisfied:

```
instance Monad Tree where
  return      :: a → Tree a
  return x    = Leaf x
  (>>=)       :: Tree a → (a → Tree b) → Tree b
  (Leaf x) >>= f = f x
  (Node l r) >>= f = Node (l >>= f) (r >>= f)
```

This declaration implements the well-known idea that substitution is monadic. In particular, if we view values of type *Tree a* as terms with variables of type *a*, then *return* converts a value into the corresponding term, and $t \gg= f$ is the term that results from applying the substitution *f* to every variable in the term *t*. For example, given the tree of characters

$$t = \text{Node} (\text{Leaf } 'a') (\text{Leaf } 'b')$$

and the substitution

```
f      :: Char → Tree Int
f 'a'  = Leaf 1
f 'b'  = Node (Leaf 2) (Leaf 3)
```


then the expression $t \gg= f$ produces the following tree of integers:

$$\text{Node} (\text{Leaf } 1) (\text{Node} (\text{Leaf } 2) (\text{Leaf } 3))$$

Now consider the following recursive function on natural numbers, which uses substitution to produce a tree of integers of a specified depth:

$$\begin{aligned} \text{fullTree} &:: \text{Int} \rightarrow \text{Tree Int} \\ \text{fullTree } 1 &= \text{return } 1 \\ \text{fullTree } (n + 1) &= \text{fullTree } n \gg\gg \lambda i \rightarrow \\ &\quad \text{Node } (\text{return } (n - i)) (\text{return } (i + 1)) \end{aligned}$$

That is, a tree of depth 1 is produced by returning a leaf, and a tree of depth $n + 1$ by recursively building a tree of depth n , and then using substitution to replace each leaf value i by a tree of depth two with leaf values $n - i$ and $i + 1$. For example, the first four trees produced by applying *fullTree* can be pictured as follows:



As we would expect from these examples, *fullTree* takes exponential time. Now consider the function *zigzag* that follows a path down a tree that alternates between moving left (*zig*) and right (*zag*), and returns the resulting leaf value:

$$\begin{aligned} \text{zigzag} &:: \text{Tree } a \rightarrow a \\ \text{zigzag} &= \text{zig} \\ &\textbf{where} \\ &\quad \text{zig } (\text{Leaf } x) = x \\ &\quad \text{zig } (\text{Node } l r) = \text{zag } l \\ &\quad \text{zag } (\text{Leaf } x) = x \\ &\quad \text{zag } (\text{Node } l r) = \text{zig } r \end{aligned}$$

In a lazy language such as Haskell, evaluating *zigzag* (*fullTree* n) only builds as much of the intermediate tree as necessary to produce the final result, which in this case is a single path. However, due to the iterative nature of *fullTree*, in which the complete tree is potentially traversed at each step in order to increase the depth by one, such an evaluation still requires quadratic time, even in a lazy language. How can this be reduced to linear time?

6.1 The codensity monad

Voigtländer's solution (2008) is based upon changing the representation of trees, using the notion of continuations. Recall that a continuation can be viewed as a function that is applied to the result of another computation. Using this idea, we can represent a value x as the function $\lambda c \rightarrow c x$ that takes a continuation c , and applies this function to x in order to produce the final result. This representation gives rise to the type $(a \rightarrow r) \rightarrow r$ of continuation computations of type a that return results of type r :

$$\textbf{type Cont } r a = (a \rightarrow r) \rightarrow r$$

It is easy to show that $Cont\ r$ is a monadic type. Moreover, we can also parameterise the declaration by another monad m to give a *monad transformer* (Liang *et al.*, 1995):

$$\mathbf{type}\ ContT\ r\ m\ a = (a \rightarrow m\ r) \rightarrow m\ r$$

For the purposes of improving the efficiency of $fullTree$, we will use the following variant, known as the *codensity* monad transformer (Jaskelioff, 2009):

$$\mathbf{type}\ CodT\ m\ a = \forall r. ((a \rightarrow m\ r) \rightarrow m\ r)$$

That is, the result type r is moved from the left-side of the declaration to the right-side, by exploiting Haskell's notion of *rank 2* types (Peyton Jones *et al.*, 2007). Moving the quantification in this manner means that whereas the continuation monad $ContT\ r\ m$ has a fixed result type r , the codensity monad $CodT\ m$ has a variable (polymorphic) result type. Making $CodT$ into a monad transformer proceeds as follows:

instance $Monad\ m \Rightarrow Monad\ (CodT\ m)$ **where**

$$return \quad :: a \rightarrow CodT\ m\ a$$

$$return\ x = \lambda c \rightarrow c\ x$$

$$(\gg\equiv) \quad :: CodT\ m\ a \rightarrow (a \rightarrow CodT\ m\ b) \rightarrow CodT\ m\ b$$

$$f\ \gg\equiv\ g = \lambda c \rightarrow f\ (\lambda x \rightarrow g\ x\ c)$$

Using the codensity monad transformer, we now define a new representation for trees, together with the necessary conversion functions between the original and new types:

$$\mathbf{type}\ Coden\ a = CodT\ Tree\ a$$

$$rep \quad \quad :: Tree\ a \rightarrow Coden\ a$$

$$rep\ t = (t\ \gg\equiv)$$


$$abs \quad \quad :: Coden\ a \rightarrow Tree\ a$$

$$abs\ c = c\ return$$

It is interesting to note the similarity to the definitions $rep\ xs = (xs++)$ and $abs\ f = f\ []$ given earlier for lists. The above definitions for trees have the same structure, except that the monoid operations $++$ and $[]$ are generalised to the monad operations $\gg\equiv$ and $return$. A simple calculation verifies the worker/wrapper assumption $abs \circ rep = id_{Tree\ a}$.

6.2 The term type

To improve the performance of $fullTree$, our aim now is to factorise this function into the composition of a more efficient worker that produces a result in the codensity monad, and a wrapper that converts the result back into the tree monad. That is, we seek to define a function $fullCoden$ that makes the following diagram commute:



Following the lead of our previous examples, we might expect to proceed by defining *fullTree* as a *fold* over the type of natural numbers, and then applying our worker/wrapper technique to derive the required worker. For this example, however, it turns out to be preferable to begin by reformulating the problem in terms of a more structured type than the natural numbers. Consider once again the definition for *fullTree*:

$$\begin{aligned} \text{fullTree} &:: \text{Int} \rightarrow \text{Tree Int} \\ \text{fullTree } 1 &= \text{return } 1 \\ \text{fullTree } (n + 1) &= \text{fullTree } n \gg\gg \lambda i \rightarrow \\ &\quad \text{Node } (\text{return } (n - i)) (\text{return } (i + 1)) \end{aligned}$$

In this definition, the resulting trees are built using three functions:

$$\begin{aligned} \text{return} &:: a \rightarrow \text{Tree } a \\ (\gg\gg) &:: \text{Tree } a \rightarrow (a \rightarrow \text{Tree } b) \rightarrow \text{Tree } b \\ \text{Node} &:: \text{Tree } a \rightarrow \text{Tree } a \rightarrow \text{Tree } a \end{aligned}$$

Based upon this observation, we can define the following type of *tree terms* whose values represent trees that are built using these functions:

data *Term a where*

$$\begin{aligned} \text{Return} &:: a \rightarrow \text{Term } a \\ \text{Bind} &:: \text{Term } a \rightarrow (a \rightarrow \text{Term } b) \rightarrow \text{Term } b \\ \text{Branch} &:: \text{Term } a \rightarrow \text{Term } a \rightarrow \text{Term } a \end{aligned}$$

Reifying functions as data in this manner is sometimes called a deep embedding. Note that because *Bind* involves terms of two different types, *Term a* is a GADT (Peyton Jones *et al.*, 2006). Categorically, defining a *fold* for such types requires moving to a functor category, in which objects are functors and arrows are natural transformations (Johann & Ghani, 2008). In Haskell, the *fold* for terms can be defined as follows:

$$\begin{aligned} \text{fold} &:: (\forall a. a \rightarrow f a) \rightarrow \\ &\quad (\forall a b. f a \rightarrow (a \rightarrow f b) \rightarrow f b) \rightarrow \\ &\quad (\forall a. f a \rightarrow f a \rightarrow f a) \rightarrow \\ &\quad (\forall a. \text{Term } a \rightarrow f a) \\ \text{fold } r \text{ b } n (\text{Return } x) &= r x \\ \text{fold } r \text{ b } n (\text{Bind } t g) &= b (\text{fold } r \text{ b } n t) (\text{fold } r \text{ b } n \circ g) \\ \text{fold } r \text{ b } n (\text{Branch } t u) &= n (\text{fold } r \text{ b } n t) (\text{fold } r \text{ b } n u) \end{aligned}$$

The use of quantifiers in the type for *fold* reflects the use of natural transformations, which in Haskell correspond to polymorphic functions. To ensure the expected universal property we also require that *Term* and *f* are functors, but we omit the details here.

Using the *fold* operator for terms, the fact that terms represent trees can now be formalised by defining an evaluation function that simply replaces the syntactic constructors on terms by the corresponding semantic operations on trees:

$$\begin{aligned} \text{eval} &:: \text{Term } a \rightarrow \text{Tree } a \\ \text{eval} &= \text{fold return } (\gg\gg) \text{Node} \end{aligned}$$


In turn, we can define a version of *fullTree* that produces a term rather than a tree, by replacing the use of the tree operations by the appropriate term constructors:

$$\begin{aligned} \mathit{fullTerm} &:: \mathit{Int} \rightarrow \mathit{Term Int} \\ \mathit{fullTerm} 1 &= \mathit{Return} 1 \\ \mathit{fullTerm} (n + 1) &= \mathit{fullTerm} n \text{ 'Bind' } \lambda i \rightarrow \\ &\quad \mathit{Branch} (\mathit{Return} (n - i)) (\mathit{Return} (i + 1)) \end{aligned}$$

A simple inductive proof shows that $\mathit{fullTree} = \mathit{eval} \circ \mathit{fullTerm}$.

6.3 Applying worker/wrapper

Having reformulated *fullTree* using an intermediate type of tree terms, we now seek to complete the following expanded version of our commuting diagram from the previous section, by defining appropriate functions *work* and *fullCoden*:



Commutativity of the upper left triangle was established in the previous section. The following definition ensures the upper right triangle also commutes, by construction:

$$\begin{aligned} \mathit{fullCoden} &:: \mathit{Int} \rightarrow \mathit{Coden Int} \\ \mathit{fullCoden} &= \mathit{work} \circ \mathit{fullTerm} \end{aligned}$$

In turn, we define the function *work* using *fold* for terms, by simply supplying the *return* and $\gg=$ operations for our codensity monad, and a suitable *node* operation:

$$\begin{aligned} \mathit{work} &:: \mathit{Term} a \rightarrow \mathit{Coden} a \\ \mathit{work} &= \mathit{fold} \mathit{return} (\gg=) \mathit{node} \\ \mathit{node} &:: \mathit{Coden} a \rightarrow \mathit{Coden} a \rightarrow \mathit{Coden} a \\ \mathit{node} f g &= \lambda c \rightarrow \mathit{Node} (f c) (g c) \end{aligned}$$

To verify that this definition makes the lower triangle in the diagram commute, i.e. $\mathit{eval} = \mathit{abs} \circ \mathit{work}$, we begin by expanding out the definitions for *eval* and *work* to give:

$$\mathit{fold} \mathit{return} (\gg=) \mathit{Node} = \mathit{abs} \circ \mathit{fold} \mathit{return} (\gg=) \mathit{node}$$

Note that *return* and $\gg=$ on the left-side of the equation are for the tree monad, and on the right-side are for the codensity monad. We then apply the worker/wrapper technique for *fold*. In particular, condition (2) for this example expands to give three equations that are

together sufficient to justify the above factorisation:

$$(2.1) \quad \text{rep} (\text{return } x) = \text{return } x$$

$$(2.2) \quad \text{rep} (t \gg\! = f) = \text{rep } t \gg\! = \text{rep} \circ f$$

$$(2.3) \quad \text{rep} (\text{Node } l \ r) = \text{node} (\text{rep } l) (\text{rep } r)$$

The first two equations state that *rep* preserves the *return* and $\gg\! =$ operations and is hence a monad morphism, while the last states that *rep* preserves the node operation. Verifying these equations is simply a matter of expanding definitions and using monad laws. We include all three proofs below to emphasise their simplicity.

Proof: (2.1)

$$\begin{aligned} & \text{rep} (\text{return } x) \ g \\ = & \quad \{ \text{applying } \text{rep} \} \\ & \text{return } x \gg\! = g \\ = & \quad \{ \text{monad law} \} \\ & g \ x \\ = & \quad \{ \text{unapplying } \text{return} \text{ for } \text{Codon} \} \\ & \text{return } x \ g \end{aligned}$$

□

Proof: (2.2)

$$\begin{aligned} & \text{rep} (t \gg\! = f) \ g \\ = & \quad \{ \text{applying } \text{rep} \} \\ & (t \gg\! = f) \gg\! = g \\ = & \quad \{ \text{monad law} \} \\ & t \gg\! = (\lambda x \rightarrow f \ x \gg\! = g) \\ = & \quad \{ \text{unapplying } \text{rep} \} \\ & t \gg\! = (\lambda x \rightarrow \text{rep} (f \ x) \ g) \\ = & \quad \{ \text{unapplying } \text{rep} \} \\ & \text{rep } t (\lambda x \rightarrow \text{rep} (f \ x) \ g) \\ = & \quad \{ \text{unapplying } \gg\! = \text{ for } \text{Codon} \} \\ & (\text{rep } t \gg\! = \text{rep} \circ f) \ g \end{aligned}$$

□

Proof: (2.3)

$$\begin{aligned} & \text{rep} (\text{Node } l \ r) \ g \\ = & \quad \{ \text{applying } \text{rep} \} \\ & \text{Node } l \ r \gg\! = g \\ = & \quad \{ \text{applying } \gg\! = \text{ for } \text{Tree} \} \\ & \text{Node} (l \gg\! = g) (r \gg\! = g) \\ = & \quad \{ \text{unapplying } \text{rep} \} \\ & \text{Node} (\text{rep } l \ g) (\text{rep } r \ g) \\ = & \quad \{ \text{unapplying } \text{node} \} \end{aligned}$$

$$\text{node } (\text{rep } l) (\text{rep } r) g$$

□

Finally, because the three internal triangles in the diagram commute, the external triangle also commutes, which verifies the desired worker/wrapper factorisation:

$$\begin{aligned} \text{fullTree} &:: \text{Int} \rightarrow \text{Tree Int} \\ \text{fullTree} &= \text{abs} \circ \text{fullCoden} \end{aligned}$$

Returning to our original problem of improving the efficiency of *zigzag* ($\text{fullTree } n$), if we now replace the original definition for *fullTree* by the new version obtained using the worker/wrapper technique, the time complexity is reduced from quadratic to linear. For example, in a simple experiment using the Glasgow Haskell Compiler the time for $n = 10000$ was reduced from around 90 seconds to 0.2 seconds. If desired, the definition $\text{fullCoden} = \text{work} \circ \text{fullTerm}$ can also be fused to eliminate the use of the intermediate term structure, further reducing the running time to under 0.1 seconds.

We conclude with a few remarks about this example. First of all, despite using a sophisticated optimisation technique in the form of the codensity monad, the proof of correctness of the efficient version of *fullTerm* still only requires simple equational reasoning. Secondly, our proof of the worker/wrapper factorisation $\text{eval} = \text{work} \circ \text{abs}$ is not specific to tree terms built using *fullTerm*, but shows how to optimise the evaluation of *any* such terms. And finally, the use of tree terms also provides an explanation for *why* the optimisation is correct, in the sense that it makes explicit the key idea of implementing the *return*, \gg , and *Node* operations on expression trees using the codensity monad.

7 Conclusion and further work

In this article we developed a general worker/wrapper theory for changing the type of recursive functions defined using fold operators, and showed how it can be used in practice as an equational reasoning technique for improving the performance of functional programs. The approach requires only basic categorical and equational reasoning principles, and using fold operators results in simpler and more structured calculations than the previous worker/wrapper theory based upon fixed point operators.

It is also interesting to recount how this work was developed. Initially we focused on the special case of *fold* for lists, and identified conditions (1) and (2) for the case of lists. However, it was not clear how these conditions were related, nor how they related to fold fusion (3), or worker/wrapper fusion (6). It was only when we generalised from lists to an arbitrary type using initial algebra semantics that it became clear that there were in fact four relevant properties, related by a simple lattice structure. Focusing on lists made it difficult to “see the wood for the trees”, and the move to a categorical approach revealed the simple underlying algebraic structure of the problem.

There are many interesting topics for further work, including mechanising the technique, other recursion operators such as *unfold*, weaker versions of the worker/wrapper assumption $\text{abs} \circ \text{rep} = \text{id}$, and other application areas. The monadic substitution example also suggests a new approach to program optimisation that we are particularly keen to

explore, based upon a deep embedding of the operations to be optimised and the use of the worker/wrapper technique to demonstrate correctness of the optimised program.

Acknowledgements

We would like to thank Ralf Hinze and the four anonymous referees for their useful comments. The last author would like to thank the Functional Programming Lab in Nottingham for funding a one week research visit during June 2009.

References

- Backhouse, Roland, Jansson, Patrik, Jeuring, Johan, & Meertens, Lambert. (1999). Generic Programming: An Introduction. *Pages 28–115 of: Swierstra, Doaitse, Henriques, Pedro, & Oliveira, Jose (eds), Advanced Functional Programming*. LNCS 1608. Springer-Verlag.
- Bird, Richard, & de Moor, Oege. (1997). *Algebra of Programming*. Prentice Hall.
- Danielsson, Nils Anders, Gibbons, Jeremy, Hughes, John, & Jansson, Patrik. (2006). Fast and Loose Reasoning is Morally Correct. *Principles of Programming Languages*. ACM Press.
- Gibbons, Jeremy. (2006). Fission for Program Comprehension. *Pages 162–179 of: Uustalu, Tarmo (ed), Mathematics of Program Construction*. Lecture Notes in Computer Science, vol. 4014. Springer-Verlag.
- Gill, Andy, & Hutton, Graham. (2009). The Worker/Wrapper Transformation. *Journal of Functional Programming*, **19**(2), 227–251.
- Hoare, Tony. (1972). Proof of Correctness of Data Representations. *Acta Informatica*, **1**(4), 271–281.
- Hughes, John. (1986). A Novel Representation of Lists and its Application to the Function Reverse. *Information Processing Letters*, **22**(3).
- Hutton, Graham. (1999). A Tutorial on the Universality and Expressiveness of Fold. *Journal of Functional Programming*, **9**(4), 355–372.
- Hutton, Graham. (2007). *Programming in Haskell*. Cambridge University Press.
- Hutton, Graham, & Wright, Joel. (2006). Calculating an Exceptional Machine. Loidl, Hans-Wolfgang (ed), *Trends in Functional Programming volume 5*. Intellect. Selected papers from the Fifth Symposium on Trends in Functional Programming, Munich, November 2004.
- Jaskelioff, Mauro. (2009). Modular Monad Transformers. *Pages 64–79 of: Proceedings of the European Symposium on Programming*. LNCS, vol. 5502. Springer.
- Johann, Patricia, & Ghani, Neil. (2008). Foundations for Structured Programming with GADTs. *Pages 297–308 of: Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press.
- Liang, Sheng, Hudak, Paul, & Jones, Mark. (1995). Monad Transformers and Modular Interpreters. *Pages 333–343 of: Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press.
- Malcolm, Grant. (1990). Algebraic Data Types and Program Transformation. *Science of Computer Programming*, **14**(2-3), 255–280.
- Meijer, Erik, Fokkinga, Maarten, & Paterson, Ross. (1991). Functional Programming with Bananas, Lenses, Envelopes and Barbed Wire. Hughes, John (ed), *Proceedings of the Conference on Functional Programming and Computer Architecture*. LNCS, no. 523. Springer-Verlag.
- Peyton Jones, Simon. (2003). *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press. Also available on the web from www.haskell.org/definition.
- Peyton Jones, Simon, & Launchbury, John. (1991). Unboxed Values as First Class Citizens in a Non-strict Functional Language. *Proceedings of the Conference on Functional Programming and Computer Architecture*. Cambridge, Massachusetts: Springer-Verlag.

- Peyton Jones, Simon, Vytiniotis, Dimitrios, Weirich, Stephanie, & Washburn, Geoffrey. (2006). Simple Unification-Based Type Inference for GADTs. *Pages 50–61 of: Proceedings of the 11th ACM SIGPLAN International Conference on Functional Programming*. ACM Press.
- Peyton Jones, Simon, Vytiniotis, Dimitrios, Weirich, Stephanie, & Shields, Mark. (2007). Practical Type Inference for Arbitrary-Rank Types. *Journal of Functional Programming*, **17**(1), 1–82.
- Voigtländer, Janis. (2008). Asymptotic Improvement of Computations over Free Monads. *Pages 388–403 of: Proceedings of the 9th International Conference on Mathematics of Program Construction*. LNCS, vol. 5133. Marseille, France: Springer-Verlag.