

The worker/wrapper transformation

ANDY GILL

University of Kansas, USA
(*e-mail: andygill@ku.edu*)

GRAHAM HUTTON

University of Nottingham, UK
(*e-mail: gmh@cs.nott.ac.uk*)

Abstract

The worker/wrapper transformation is a technique for changing the type of a computation, usually with the aim of improving its performance. It has been used by compiler writers for many years, but the technique is little known in the wider functional programming community, and has never been described precisely. In this article we explain, formalise and explore the generality of the worker/wrapper transformation. We also provide a systematic recipe for its use as an equational reasoning technique for improving the performance of programs, and illustrate the power of this recipe using a range of examples.

1 Introduction

The worker/wrapper transformation is a technique for transforming a computation of one type into a *worker* of a different type, together with a *wrapper* that acts as an impedance matcher between the original and new computations.

The technique can be used to improve the performance of functional programs by improving the choice of data structures used. For example, in the Glasgow Haskell Compiler, a specific instance of the worker/wrapper transformation is used to replace boxed data structures by unboxed structures, based upon strictness information (Peyton Jones & Launchbury, 1991). Despite its practical importance, however, the technique is little known in the wider functional programming community, and to the best of our knowledge has never been described precisely.

In this article we explain, formalise, and explore the generality of the worker/wrapper transformation. We also provide a systematic recipe for its use, and illustrate the power of this recipe using a range of programming examples. More precisely, the article makes the following contributions:

- We present the worker/wrapper transformation in a general semantic framework (Section 3), rather than via specific syntactic instances, in order to make the technique as widely applicable as possible.
- We give a correctness proof for the technique (Section 3), based upon the use of the rolling rule from fixed point calculus, which provides a range of explicit conditions under which the technique can be applied.

- We provide a systematic recipe for its use as an equational reasoning technique for improving the performance of programs (Section 3), by means of a step-by-step process for applying the transformation.
- We apply this recipe to four different approaches to program optimisation, based upon the use of accumulation (Section 4), unboxing (Section 5), memoisation (Section 6) and continuations (Section 7).

A summary of related work, and directions for further work, are provided in the concluding sections. The article is aimed at readers who are familiar with the basics of reasoning about functional programs, say to the level of Bird (1998), but no previous experience with the worker/wrapper transformation is assumed. The techniques are presented using Haskell (Peyton Jones, 2003), but can easily be adapted to other functional languages. An extended version of the article that includes all the proofs is available from the authors' web pages.

2 The basic idea

In this section we review the basic idea of the worker/wrapper transformation, in a similar manner to which it was originally described by Peyton Jones and Launchbury (1991), and implemented in the Glasgow Haskell Compiler.

Suppose that we are given a function f , defined in the form

$$f = \text{body}$$

where *body* is the right-hand side of the definition, and may include recursive calls to f . At this syntactic level, the first step in applying the worker/wrapper transformation is to define appropriate functions *wrap* and *unwrap* that allow the function f to be redefined by the equation $f = \text{wrap } (\text{unwrap } \text{body})$, which is then split into two equations by naming the intermediate result, as follows:

$$\begin{aligned} f &= \text{wrap } \text{work} \\ \text{work} &= \text{unwrap } \text{body} \end{aligned}$$

In this manner, f has been factorised into the application of a 'wrapper' function *wrap* to a 'worker' function *work*, itself defined by applying *unwrap* to the body of the original definition for f . Note that if f was originally recursive, then f and *work* are now mutually recursive. The next step in the process is to eliminate such mutual recursion by inlining the new definition for f in the definition for *work*, thereby making *work* into a recursive definition that is independent of f :

$$\begin{aligned} f &= \text{wrap } \text{work} \\ \text{work} &= \text{unwrap } (\text{body } [\text{wrap } \text{work} / f]) \end{aligned}$$

As usual, $e[e'/x]$ denotes the result of substituting e' for all free occurrences of the variable x in the expression e . The final step is then to simplify the resulting definitions for f and *work* using standard program transformation techniques.

For example, Peyton Jones and Launchbury (1991) show how this process can be used to transform a recursive definition of the factorial function into a more

efficient version that is defined in terms of a worker that only uses unboxed integers, together with a wrapper that takes care of the initial unboxing and final boxing.

The above, syntactic, description of the worker/wrapper transformation is appealingly simple, but raises a number of important questions. Is the technique actually correct? How can this be proved? Under what conditions does the proof hold? How should it be used in practice? What kind of applications is it suitable for? Addressing these questions is the aim of the remainder of this article.

3 The worker/wrapper transformation

In order to give a precise treatment of the worker/wrapper transformation, and hence prove its correctness, we move from the informal syntactic approach based upon inlining to a formal semantic approach based upon the use of fixed points. We begin by defining a fixed point operator in Haskell:

$$\begin{aligned} \text{fix} &:: (a \rightarrow a) \rightarrow a \\ \text{fix } f &= f (\text{fix } f) \end{aligned}$$

The property of this operator that we will use to formalise the worker/wrapper transformation is the *rolling rule* (Backhouse, 2002), which allows us to pull the first argument of a composition outside a fixed point, resulting in the composition swapping the order of its arguments, or ‘rolling over’:

$$\text{fix } (g \circ f) = g (\text{fix } (f \circ g))$$

Informally, this rule is valid because both sides expand to the infinite application $g (f (g (f (g (f \dots))))))$. A formal proof, together with a brief review of the fixed point approach to recursion, is provided in the Appendix.

Now consider the problem of changing the type of a recursive computation. More precisely, suppose that we are given a computation $\text{comp} :: A$, defined as the least fixed point $\text{comp} = \text{fix } \text{body}$ of some function $\text{body} :: A \rightarrow A$, and that we wish to change the underlying type from A to some other type B .

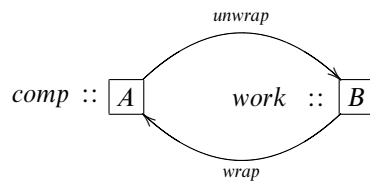
The worker/wrapper approach to this problem of changing types is based upon defining conversion functions $\text{unwrap} :: A \rightarrow B$ and $\text{wrap} :: B \rightarrow A$ such that the equation $\text{wrap} \circ \text{unwrap} = \text{id}_A$ holds, where id_A is the identity function for the type A . (To aid the reader, identity functions are subscripted with their types throughout the article.) This equation states that converting a value of the original type into the new type and then back again does not change the value, and formalises the idea that the type A can be faithfully represented by the type B .

Given such a setting, we can now calculate as follows:

$$\begin{aligned} &\text{comp} \\ = &\quad \{ \text{applying } \text{comp} \} \\ &\text{fix } \text{body} \\ = &\quad \{ \text{id is the identity for } \circ \} \\ &\text{fix } (\text{id}_A \circ \text{body}) \end{aligned}$$

$$\begin{aligned}
&= \{ \dagger \text{ assuming } \mathit{wrap} \circ \mathit{unwrap} = \mathit{id}_A \} \\
&\quad \mathit{fix} (\mathit{wrap} \circ \mathit{unwrap} \circ \mathit{body}) \\
&= \{ \text{rolling rule} \} \\
&\quad \mathit{wrap} (\mathit{fix} (\mathit{unwrap} \circ \mathit{body} \circ \mathit{wrap})) \\
&= \{ \text{define } \mathit{work} = \mathit{fix} (\mathit{unwrap} \circ \mathit{body} \circ \mathit{wrap}) \} \\
&\quad \mathit{wrap} \mathit{work}
\end{aligned}$$

That is, we have shown how an arbitrary recursive computation of type A can be factorised as a wrapper of type $B \rightarrow A$ applied to a new recursive worker computation of type B , based upon the assumption that the identity function for A can be split into the composition of two conversion functions. The following diagram summarises the primary typing relationships:



As we shall see, however, in general there is no requirement that this diagram commutes. For example, for some applications the assumption $\mathit{wrap} \circ \mathit{unwrap} = \mathit{id}_A$ may not be true in general, but only in the context (labelled by \dagger) in which it is used in the above calculation. In particular, we may sometimes require the weaker property $\mathit{wrap} \circ \mathit{unwrap} \circ \mathit{body} = \mathit{id}_A \circ \mathit{body}$, which states that the assumption is true for values produced by body , or even $\mathit{fix} (\mathit{wrap} \circ \mathit{unwrap} \circ \mathit{body}) = \mathit{fix} (\mathit{id}_A \circ \mathit{body})$, which also takes the recursive context into account. That is, we have the following hierarchy of *worker/wrapper assumptions* that support the above calculation:

$$\begin{aligned}
&\mathit{wrap} \circ \mathit{unwrap} = \mathit{id}_A \\
&\quad \Downarrow \\
&\mathit{wrap} \circ \mathit{unwrap} \circ \mathit{body} = \mathit{body} \\
&\quad \Downarrow \\
&\mathit{fix} (\mathit{wrap} \circ \mathit{unwrap} \circ \mathit{body}) = \mathit{fix} \mathit{body}
\end{aligned}$$

For any given application, we will use the strongest such assumption that is valid, i.e. the simplest to verify; often the basic equation $\mathit{wrap} \circ \mathit{unwrap} = \mathit{id}_A$ will suffice. In conclusion, the original syntactic description of the worker/wrapper transformation can now be formalised on semantic grounds as shown in Figure 1.

Once this transformation has been applied, one then attempts to simplify the worker using normal fold/unfold transformation (Burstall & Darlington, 1977). It is usually convenient to begin by rewriting the worker using explicit recursion

$$\mathit{work} = \mathit{unwrap} (\mathit{body} (\mathit{wrap} \mathit{work}))$$

If $comp :: A$ is a recursive computation defined by $comp = fix\ body$ for some $body :: A \rightarrow A$, and $wrap :: B \rightarrow A$ and $unwrap :: A \rightarrow B$ are conversion functions satisfying any of the worker/wrapper assumptions, then

$$comp = wrap\ work$$

where $work :: B$ is defined by

$$work = fix\ (unwrap \circ body \circ wrap)$$

Fig. 1. The worker/wrapper transformation.

which transformation can be verified as follows:

$$\begin{aligned} & work \\ = & \quad \{ \text{applying } work \} \\ & fix\ (unwrap \circ body \circ wrap) \\ = & \quad \{ \text{applying } fix \} \\ & (unwrap \circ body \circ wrap)\ (fix\ (unwrap \circ body \circ wrap)) \\ = & \quad \{ \text{unapplying } work \} \\ & (unwrap \circ body \circ wrap)\ work \\ = & \quad \{ \text{applying } \circ \} \\ & unwrap\ (body\ (wrap\ work)) \end{aligned}$$

Further simplification of the worker is then typically driven by the desire to fuse together instances of $unwrap$ and $wrap$, to eliminate the overhead of repeatedly converting between the two types. In general, it is not the case that $unwrap \circ wrap$ can be fused to give id_B , but the following *worker/wrapper fusion* property, in which the argument value of type B is the worker itself, is often applicable:

$$\text{If } wrap \circ unwrap = id_A, \text{ then } unwrap\ (wrap\ work) = work$$

For example, this property will be utilised in two of our programming applications (Sections 4 and 7). The fusion property itself can be verified as follows:

$$\begin{aligned} & unwrap\ (wrap\ work) \\ = & \quad \{ \text{applying } work \} \\ & unwrap\ (wrap\ (unwrap\ (body\ (wrap\ work)))) \\ = & \quad \{ \text{assuming } wrap \circ unwrap = id_A \} \\ & unwrap\ (body\ (wrap\ work)) \\ = & \quad \{ \text{unapplying } work \} \\ & work \end{aligned}$$

In summary, we have the following general recipe for applying the worker/wrapper transformation to change the type of a recursive computation:

- Express the computation as the least fixed point;
- Choose the desired new type for the computation;

- Define conversions between the original and new types;
- Check that they satisfy one of the worker/wrapper assumptions;
- Apply the worker/wrapper transformation;
- Simplify the resulting definitions.

We conclude this section by noting that the worker/wrapper transformation can also be formalised using *fixed point fusion* (Meijer *et al.*, 1991) which requires the additional property that *wrap* is strict ($wrap \perp = \perp$). Fortunately this property follows automatically from the basic worker/wrapper assumption $wrap \circ unwrap = id$, by virtue of the fact that any (monotonic) function that is right invertible must be strict. We prefer the formulation using the rolling rule because it corresponds directly to our understanding of why the transformation is valid.

4 Example: the reverse function

As a first example of our worker/wrapper recipe, we will use it to transform a simple definition for the function that reverses a list into a more efficient version that uses *accumulation*. This transformation is normally achieved using more elementary techniques (Hutton, 2007, Section 13.6), but we now show that it also fits naturally into our general worker/wrapper paradigm.

4.1 Hughes lists

We begin by reviewing the idea of representing lists using functions, which is the essential algorithmic idea underlying this example.

As observed by Hughes (1986), it is possible to represent the normal type of lists in an alternative manner using functions, which represent a list xs by the function $(xs \#)$ that prepends xs to its argument. We can implement this idea by defining a type $H\ a$ of *Hughes lists*, together with a representation function rep

$$\begin{aligned} \text{type } H\ a &= [a] \rightarrow [a] \\ rep &:: [a] \rightarrow H\ a \\ rep\ xs &= (xs \#) \end{aligned}$$

An important property of this representation is that the function rep forms a homomorphism from lists to functions, in the sense that

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

That is, rep is a homomorphism from the monoid of (finite) lists, for which the associative operator is $\#$ and the unit is $[]$, to the monoid of functions on lists, for which the operator is composition \circ and the unit is the identity function $id_{[a]}$.

In addition to the function *rep* that converts normal lists into Hughes lists, it is also natural to consider conversion in the opposite direction, which is achieved by simply applying the given function to the empty list

$$\begin{aligned} \text{abs} &:: H\ a \rightarrow [a] \\ \text{abs}\ f &= f\ [] \end{aligned}$$

This definition ensures that converting a list into Hughes form and back again returns the original list, or more formally, it ensures that $\text{abs} \circ \text{rep} = \text{id}_{[a]}$.

In summary, using the terminology of data representation (Hoare, 1972), we have shown how the ‘abstract’ type $[a]$ can be represented using the ‘concrete’ type $H\ a$. As is usual in this setting, note that the representation function *rep* is injective but not surjective, while the abstraction function *abs* is surjective but not injective.

4.2 Reverse

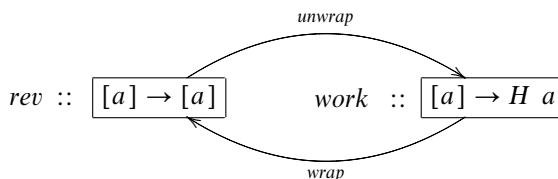
Now consider the following definition for the function that reverses a list:

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

Because of the use of $++$, this definition takes quadratic time in the length of its argument. We now show how this definition can be transformed into a more efficient *worker* that uses an extra argument to accumulate the result, together with a *wrapper* that takes care of the initial setup. The first step is to redefine *rev* as the least fixed point, by abstracting over the recursive call in the body

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

Following our worker/wrapper recipe, and utilising the idea of Hughes lists, our aim now is to transform the computation *rev* of type $[a] \rightarrow [a]$ into a worker of type $[a] \rightarrow H\ a$ (which abbreviates $[a] \rightarrow [a] \rightarrow [a]$ and hence introduces the necessary extra argument list), as illustrated by the following diagram:



Implementing the conversion functions *unwrap* and *wrap* is simply a matter of composing the conversion functions *rep* and *abs* from the previous section

$$\begin{aligned} \text{unwrap} &:: ([a] \rightarrow [a]) \rightarrow ([a] \rightarrow H\ a) \\ \text{unwrap } f &= \text{rep} \circ f \\ \text{wrap} &:: ([a] \rightarrow H\ a) \rightarrow ([a] \rightarrow [a]) \\ \text{wrap } g &= \text{abs} \circ g \end{aligned}$$

Verifying the worker/wrapper assumption, in this case $\text{wrap} \circ \text{unwrap} = \text{id}_{[a] \rightarrow [a]}$, follows from the fact that $\text{abs} \circ \text{rep} = \text{id}_{[a]}$ by straightforward calculation.

Now that we have satisfied the preconditions for the worker/wrapper transformation, applying this transformation and rewriting the worker using explicit recursion rather than the *fix* operator gives the following result:

$$\begin{aligned} \text{rev} &:: [a] \rightarrow [a] \\ \text{rev} &= \text{wrap } \text{work} \\ \text{work} &:: [a] \rightarrow H\ a \\ \text{work} &= \text{unwrap } (\text{body } (\text{wrap } \text{work})) \end{aligned}$$

Making the list argument explicit in the definition for *rev*, and expanding out *wrap*, we obtain the wrapper for the efficient version of *rev*, which simply supplies the empty list as the initial accumulator

$$\begin{aligned} \text{rev} &:: [a] \rightarrow [a] \\ \text{rev } xs &= \text{work } xs [] \end{aligned}$$

Now let us focus our attention on simplifying the definition for the worker function. First of all, we redefine *work* using pattern matching

$$\begin{aligned} \text{work } [] &= \text{rep } [] \\ \text{work } (x : xs) &= \text{rep } (\text{wrap } \text{work } xs \# [x]) \end{aligned}$$

and the transformation can be verified as follows:

$$\begin{aligned} &\text{work } xs \\ &= \quad \{ \text{applying } \text{work} \} \\ &\quad \text{unwrap } (\text{body } (\text{wrap } \text{work}))\ xs \\ &= \quad \{ \text{applying } \text{unwrap} \} \\ &\quad \text{rep } (\text{body } (\text{wrap } \text{work})\ xs) \\ &= \quad \{ \text{applying } \text{body} \} \\ &\quad \text{rep } (\mathbf{case } xs \mathbf{ of} \\ &\quad \quad [] \rightarrow [] \\ &\quad \quad (x : xs) \rightarrow \text{wrap } \text{work } xs \# [x]) \\ &= \quad \{ \text{distribution over } \mathbf{case} \} \\ &\quad \mathbf{case } xs \mathbf{ of} \\ &\quad \quad [] \rightarrow \text{rep } [] \\ &\quad \quad (x : xs) \rightarrow \text{rep } (\text{wrap } \text{work } xs \# [x]) \end{aligned}$$

The distribution step assumes that *rev* is only applied to finite lists, to ensure that the argument list *xs* in this step is not undefined (\perp).

Now we exploit the fact that $rep :: [a] \rightarrow H\ a$ is a homomorphism from lists to functions, to rewrite the new definition as follows:

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

Then we simplify the first argument of the composition in this definition by exposing the implicit use of *unwrap* and then fusing this with *wrap*

$$\begin{aligned} &rep (wrap\ work\ xs) \\ = &\quad \{ \text{unapplying } unwrap \} \\ &unwrap (wrap\ work)\ xs \\ = &\quad \{ \text{worker/wrapper fusion} \} \\ &work\ xs \end{aligned}$$

That is, we now have the following definition:

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

Finally, if we make the extra list argument in this definition explicit, and then expand out \circ and *rep*, we obtain the worker for the efficient version of reverse, which uses an extra list argument to accumulate the final result

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

We conclude this section with three observations about the above derivation. First of all, note that everything follows in a straightforward manner from the idea of using Hughes' representation of lists, in the sense that once we have made this decision, the rest of the derivation proceeds using standard fold/unfold transformation. Secondly, the derivation does not require the use of induction, other than the implicit use to verify the auxiliary result that finite lists form a monoid, which is used in proof that *rep* is a homomorphism. And finally, it is possible to generalise from this example and formulate a special case of the worker/wrapper transformation for changing the result type of a function; because this only saves a few steps in the above calculations, we prefer to use our general formulation.

5 Example: the factorial function

For our second example, we formalise an example from the original article on the worker/wrapper transformation (Peyton Jones & Launchbury, 1991), by showing how to transform a simple definition for the factorial function on integers into a more efficient version that only uses *unboxed* integers.

5.1 Boxed and unboxed integers

Consider the question of how integers should be represented in the implementation of a programming language. For strict languages, such as ML, integers can be

represented directly as binary numbers, for example as a 32-bit word. Integers represented in this manner are called *unboxed* integers, because no additional packaging is required around the raw numeric values themselves.

For non-strict languages, such as Haskell, a more refined representation is required, because evaluation is only performed on demand. In particular, we must distinguish between a computation that will return an integer value and an integer value itself. In practice, this is usually achieved using some form of tagging in the implementation. Such tagged integer representations are called *boxed* integers.

Tagging supports the use of non-strict evaluation, but carries a considerable overhead. For example, evaluating $x + y$ conceptually requires evaluating the expressions x and y , unboxing the resulting values (removing the tags), performing the actual addition and then boxing the result (reinstating the tag). It would clearly be more efficient to work with unboxed integers as much as possible.

The key to achieving this behaviour in Haskell is an idea put forward by Peyton Jones and Launchbury (1991), namely to reflect the notion of boxed and unboxed integers directly in the language itself, via the following type definition:

$$\mathbf{data} \text{ Int} = I_{\#} \text{ Int}_{\#}$$

That is, a boxed value of type *Int* can be constructed by applying the constructor tag $I_{\#}$ to a value of type $\text{Int}_{\#}$, which is the built-in type of unboxed 32-bit integers, and is unlifted in the sense that it has no \perp element. Note that, by convention in Haskell, unboxed types, values, variables and functions have a $\#$ suffix on their names, as in $\text{Int}_{\#}$, $0_{\#}$, $n_{\#}$ and $+_{\#}$. However, this symbol has no semantic meaning, and is used purely for identification purposes.

Given the above type definition, converting an unboxed integer into a boxed integer can be achieved simply by applying the constructor $I_{\#}$, and in the opposite direction by removing this constructor through pattern matching. Thus there is no need to define explicit conversion functions.

5.2 Factorial

Consider the following definition for the factorial function on non-negative integers:

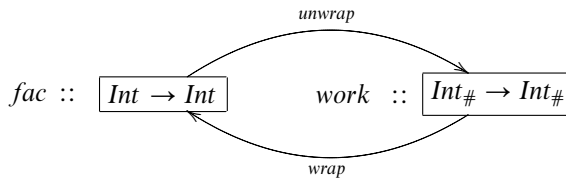
$$\begin{aligned} \text{fac} &:: \text{Int} \rightarrow \text{Int} \\ \text{fac } 0 &= 1 \\ \text{fac } n &= n * \text{fac } (n - 1) \end{aligned}$$

In a non-strict language such as Haskell, a straightforward implementation of this definition will be rather inefficient, due to repeated unboxing and boxing of integer values. However, because *fac* is strict ($\text{fac } \perp = \perp$), it is safe to replace the definition of *fac* by a more efficient *worker* that only uses unboxed integers and strict evaluation, together with a *wrapper* that takes care of the initial unboxing and final reboxing. We show how this transformation can be achieved in practice. Following our recipe,

the first step is to redefine *fac* as the least fixed point

$$\begin{aligned} \mathit{fac} &:: \mathit{Int} \rightarrow \mathit{Int} \\ \mathit{fac} &= \mathit{fix} \ \mathit{body} \\ \mathit{body} &:: (\mathit{Int} \rightarrow \mathit{Int}) \rightarrow (\mathit{Int} \rightarrow \mathit{Int}) \\ \mathit{body} \ f \ 0 &= 1 \\ \mathit{body} \ f \ n &= n * f \ (n - 1) \end{aligned}$$

Our aim now is to transform the computation *fac* of type $\mathit{Int} \rightarrow \mathit{Int}$ into a worker of type $\mathit{Int}_{\#} \rightarrow \mathit{Int}_{\#}$, as illustrated in the following diagram:



Implementing the conversion functions *wrap* and *unwrap* is simply a matter of taking care of the necessary boxing and unboxing

$$\begin{aligned} \mathit{unwrap} &:: (\mathit{Int} \rightarrow \mathit{Int}) \rightarrow (\mathit{Int}_{\#} \rightarrow \mathit{Int}_{\#}) \\ \mathit{unwrap} \ f \ x_{\#} &= \mathbf{case} \ f \ (I_{\#} \ x_{\#}) \ \mathbf{of} \\ &\quad I_{\#} \ y_{\#} \rightarrow y_{\#} \\ \mathit{wrap} &:: (\mathit{Int}_{\#} \rightarrow \mathit{Int}_{\#}) \rightarrow (\mathit{Int} \rightarrow \mathit{Int}) \\ \mathit{wrap} \ g_{\#} \ x &= \mathbf{case} \ x \ \mathbf{of} \\ &\quad I_{\#} \ x_{\#} \rightarrow \mathbf{case} \ g_{\#} \ x_{\#} \ \mathbf{of} \\ &\quad \quad y_{\#} \rightarrow I_{\#} \ y_{\#} \end{aligned}$$

Using the above definitions and properties of case expressions, it is straightforward to show that $\mathit{wrap} \circ \mathit{unwrap} = \mathit{id}_{\mathit{Int} \rightarrow \mathit{Int}}$, but only if we restrict our attention to argument functions of type $\mathit{Int} \rightarrow \mathit{Int}$ that are strict. Hence for this example, the simple assumption $\mathit{wrap} \circ \mathit{unwrap} = \mathit{id}_{\mathit{Int} \rightarrow \mathit{Int}}$ is not always valid. However, because *body f* is defined by pattern matching and is therefore strict, we do have the weaker assumption $\mathit{wrap} \circ \mathit{unwrap} \circ \mathit{body} = \mathit{body}$, which is precisely where strictness is exploited in this example. More complex functions than factorial, for which determining strictness depends upon taking account of recursion, will require our weakest worker/wrapper assumption defined using *fix*.

Now that we have satisfied the worker/wrapper preconditions, applying this transformation gives the following result:

$$\begin{aligned} \mathit{fac} &:: \mathit{Int} \rightarrow \mathit{Int} \\ \mathit{fac} &= \mathit{wrap} \ \mathit{work} \\ \mathit{work} &:: \mathit{Int}_{\#} \rightarrow \mathit{Int}_{\#} \\ \mathit{work} &= \mathit{unwrap} \ (\mathit{body} \ (\mathit{wrap} \ \mathit{work})) \end{aligned}$$

Making the integer argument explicit in the definition for *fac*, and expanding out *wrap*, we obtain the wrapper for the unboxed version of *fac*, which unboxes the integer argument, applies the worker and then boxes the result

$$\begin{aligned} \text{fac} &:: \text{Int} \rightarrow \text{Int} \\ \text{fac } n &= \mathbf{case } n \mathbf{ of} \\ &\quad I_{\#} x_{\#} \rightarrow \mathbf{case } \text{work } x_{\#} \mathbf{ of} \\ &\quad \quad y_{\#} \rightarrow I_{\#} y_{\#} \end{aligned}$$

Now let us simplify the definition for the worker function, which in this case amounts to expanding out definitions and then simplifying the result. As with the *reverse* example, the first step is to redefine *work* using pattern matching, and this transformation can be verified in a similar manner as before

$$\begin{aligned} \text{work } x_{\#} &= \mathbf{case } (\mathbf{case } (I_{\#} x_{\#}) \mathbf{ of} \\ &\quad 0 \rightarrow 1 \\ &\quad n \rightarrow \mathbf{case } (n - 1) \mathbf{ of} \\ &\quad \quad I_{\#} a_{\#} \rightarrow \mathbf{case } \text{work } a_{\#} \mathbf{ of} \\ &\quad \quad \quad b_{\#} \rightarrow n * (I_{\#} b_{\#})) \mathbf{ of} \\ &\quad I_{\#} y_{\#} \rightarrow y_{\#} \end{aligned}$$

Then if we expand out the boxed components 0, 1, *n*, *−* and *** in terms of the constructor *I_#* (for example, by replacing 0 with *I_# 0_#*), and simplify the resulting definition using properties of case, we obtain a worker with the same structure as the original factorial function, but which now operates entirely using unboxed integers, and hence avoids the need for repeated unboxing and boxing

$$\begin{aligned} \text{work} &:: \text{Int}_{\#} \rightarrow \text{Int}_{\#} \\ \text{work } x_{\#} &= \mathbf{case } x_{\#} \mathbf{ of} \\ &\quad 0_{\#} \rightarrow 1_{\#} \\ &\quad n_{\#} \rightarrow n_{\#} *_{\#} \text{work } (n_{\#} -_{\#} 1_{\#}) \end{aligned}$$

Once again, note that once we have made the initial decision regarding the change in type of the function, applying our worker/wrapper recipe is largely a matter of routine equational reasoning, and does not require the use of induction. Note also that this derivation involved changing both the argument and result type of a function, whereas our previous example only changed the result type.

6 Example: the Fibonacci function

For our next example, we show how the worker/wrapper transformation can be used to transform a simple definition for the Fibonacci function on natural numbers into a more efficient version that operates by means of a *memo table*. In contrast to our previous two examples, in which a function is represented by another function, this derivation represents a function by a data structure.

6.1 Memoisation

The term memoisation was coined by Michie (1968), and refers to the optimisation technique of tabulating all the arguments that a function is called with, together with the corresponding results returned, and reusing these results if the function is called again with any of its previous arguments. If this happens often, memoisation can lead to a dramatic improvement in performance.

Our approach to memoising the Fibonacci function is to observe that any function on natural numbers can be represented in an alternative manner as an infinite list (or stream) of results. In particular, we represent a function f as the stream $[f\ 0, f\ 1, f\ 2, \dots]$ that tabulates its behaviour for all possible argument values.

To implement this representation, let us write Nat for the subtype of Int comprising the natural numbers $0, 1, 2, 3, \dots$, and $Stream\ a$ for the subtype of $[a]$ comprising the infinite lists of elements of type a . Using these two subtypes allows us to be more precise about the types of functions that we will define, while still being able to use familiar notation and library functions for integers and lists.

For example, a function $unwrap$ that converts functions on natural numbers into streams in the manner described above can now be defined

$$\begin{aligned} unwrap &:: (Nat \rightarrow a) \rightarrow Stream\ a \\ unwrap\ f &= map\ f\ [0..] \end{aligned}$$

Dually, we can also perform conversion in the opposite direction, from a stream to a function, by simply indexing into the stream:

$$\begin{aligned} wrap &:: Stream\ a \rightarrow (Nat \rightarrow a) \\ wrap\ xs &= (xs\ !!) \end{aligned}$$

The auxiliary operator $!!$ selects the n th element of a stream

$$\begin{aligned} (!! &:: Stream\ a \rightarrow Nat \rightarrow a \\ xs\ !!\ 0 &= head\ xs \\ xs\ !!\ (n + 1) &= (tail\ xs)\ !!\ n \end{aligned}$$

Using these definitions, we can show that $wrap \circ unwrap = id_{Nat \rightarrow a}$ by making the function and numeric arguments explicit, then expanding out definitions to give $(unwrap\ f)\ !!\ n = f\ n$ and finally verifying this equation by induction on n .

The dual property, $unwrap \circ wrap = id_{Stream\ a}$, also holds. In particular, making the stream argument explicit and expanding out definitions gives $unwrap\ (xs\ !!) = xs$, which can be verified by coinduction on streams (Turner, 1995).

In conclusion, the functions $unwrap$ and $wrap$ establish an isomorphism between the types $Nat \rightarrow a$ and $Stream\ a$. This isomorphism itself is well known, but perhaps surprisingly, generalises in allowing any function type with an inductively defined argument type to be represented in a coinductive manner (Altenkirch, 2001), thereby providing a basis for memoising a large class of functions.

6.2 Fibonacci

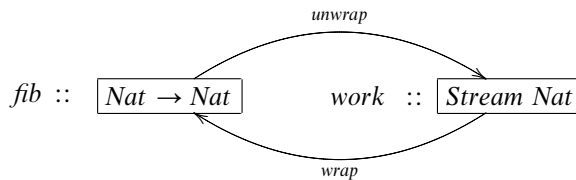
Consider the following standard definition for the Fibonacci function on natural numbers, as found in any textbook on numerical series:

$$\begin{aligned} fib &:: Nat \rightarrow Nat \\ fib\ 0 &= 0 \\ fib\ 1 &= 1 \\ fib\ (n + 2) &= fib\ n + fib\ (n + 1) \end{aligned}$$

That is, the first two Fibonacci numbers are 0 and 1, and each successive number is the sum of the previous two. While this definition directly expresses the desired behaviour of the function, it also recomputes the same results many times over, in fact requiring exponential time in the size of its argument. We now show how it can be transformed into a more efficient *worker* that constructs a memo table comprising the stream $[0, 1, 1, 2, 3, 5, \dots]$ of all Fibonacci numbers (and hence ensures that each result is computed at most once), together with a *wrapper* that selects the required element from this table. The first step is to redefine *fib* using *fix*:

$$\begin{aligned} fib &:: Nat \rightarrow Nat \\ fib &= fix\ body \\ body &:: (Nat \rightarrow Nat) \rightarrow (Nat \rightarrow Nat) \\ body\ f\ 0 &= 0 \\ body\ f\ 1 &= 1 \\ body\ f\ (n + 2) &= f\ n + f\ (n + 1) \end{aligned}$$

Our aim now is to transform the computation *fib* of type $Nat \rightarrow Nat$ into a worker of type $Stream\ Nat$, as illustrated in the following diagram:



Because we have already defined appropriate conversion functions satisfying the assumption $wrap \circ unwrap = id_{Nat \rightarrow Nat}$ in the previous section, we can now apply the worker/wrapper transformation to obtain the following result:

$$\begin{aligned} fib &:: Nat \rightarrow Nat \\ fib &= wrap\ work \\ work &:: Stream\ Nat \\ work &= unwrap\ (body\ (wrap\ work)) \end{aligned}$$

Making the numeric argument explicit in the definition for *fib*, and expanding out *wrap*, we obtain the *fib* wrapper for the memoised version of *fib*, which simply selects

the required element from the memo table

$$\begin{aligned} fib &:: Nat \rightarrow Nat \\ fib\ n &= work\ !!\ n \end{aligned}$$

Now we simplify the worker itself, which in this case just amounts to expanding out definitions, and does not require any form of fusion:

$$\begin{aligned} &work \\ = &\{ \text{applying } work \} \\ &unwrap\ (body\ (wrap\ work)) \\ = &\{ \text{applying } unwrap \} \\ &map\ (body\ (wrap\ work))\ [0..] \\ = &\{ \text{applying } body \} \\ &map\ (\lambda n \rightarrow \mathbf{case\ } n \mathbf{\ of} \\ &\quad 0 \rightarrow 0 \\ &\quad 1 \rightarrow 1 \\ &\quad (n + 2) \rightarrow wrap\ work\ n + wrap\ work\ (n + 1))\ [0..] \\ = &\{ \text{applying } wrap \} \\ &map\ (\lambda n \rightarrow \mathbf{case\ } n \mathbf{\ of} \\ &\quad 0 \rightarrow 0 \\ &\quad 1 \rightarrow 1 \\ &\quad (n + 2) \rightarrow work\ !!\ n + work\ !!\ (n + 1))\ [0..] \end{aligned}$$

That is, we have derived the following definition for the worker, which builds a memo table comprising all Fibonacci numbers, using the table itself to ensure that each number is computed at most once, and lazy evaluation to ensure that the table (which is conceptually infinite) is built on a demand-driven basis

$$\begin{aligned} work &:: Stream\ Nat \\ work &= map\ f\ [0..] \\ &\mathbf{where} \\ &\quad f\ 0 = 0 \\ &\quad f\ 1 = 1 \\ &\quad f\ (n + 2) = work\ !!\ n + work\ !!\ (n + 1) \end{aligned}$$

In conclusion, we have improved the time performance of the *fib* function from exponential to quadratic. Further improvements can readily be made in a variety of ways, such as representing the memo table in a more efficient manner (e.g. using an array with constant-time indexing rather than a stream with linear-time indexing), exploiting the structure of the computation to design a specialised version of the memo table (e.g. using a pair of numbers rather than a stream), or using more advanced programming techniques (e.g. cyclic programming). For example, using the technique of unique fixed points (Hinze, 2008), it can be shown that the above definition for the worker is equivalent to the well-known cyclic definition for the stream of Fibonacci numbers

$$\begin{aligned} fibs &:: Stream\ Nat \\ fibs &= 0 : 1 : zipWith\ (+)\ fibs\ (tail\ fibs) \end{aligned}$$

7 Example: an evaluation function

For our final example, we show how to use the worker/wrapper transformation in implementing the transformation of a function into *continuation-passing* style, an important precursor to many program analyses and optimisations.

7.1 Continuations

As observed by Reynolds (1972), it is possible to represent any type in an alternative manner using *continuations*. The notion of a continuation can be defined in many ways, but for our purposes it is simply a function that is applied to the result of another computation. Using this idea, we represent a value x as the function $\lambda c \rightarrow c\ x$ that takes another function c (a continuation) as its argument, and applies this function to x in order to produce the final result.

We can implement this representation using the type $(a \rightarrow r) \rightarrow r$ of continuation computations of type a that return results of type r (Wadler, 1992b). For our purposes, however, we only require the special case when the two types are the same, which we abbreviate by *Cont a*, and define a function *rep* that converts normal values into their representation as continuation values:

$$\begin{aligned} \mathbf{type} \text{ Cont } a &= (a \rightarrow a) \rightarrow a \\ \mathit{rep} &:: a \rightarrow \text{Cont } a \\ \mathit{rep } x &= \lambda c \rightarrow c\ x \end{aligned}$$

Dually, we can also perform conversion in the opposite direction, by supplying the identity function as the continuation:

$$\begin{aligned} \mathit{abs} &:: \text{Cont } a \rightarrow a \\ \mathit{abs } f &= f\ \mathit{id}_a \end{aligned}$$

Using these definitions, it is easy to show that $\mathit{abs} \circ \mathit{rep} = \mathit{id}_a$.

7.2 Evaluation

Consider a simple expression language comprising integers values, an addition operator, a single exceptional value called *throw* and a *catch* operator:

$$\mathbf{data} \text{ Expr} = \text{Val Int} \mid \text{Add Expr Expr} \mid \text{Throw} \mid \text{Catch Expr Expr}$$

This language provides an appropriate minimal setting in which to investigate various aspects of the semantics of exceptions (Hutton & Wright, 2004, 2006, 2007). Informally, *Throw* abandons the current computation and throws an exception, while *Catch x y* behaves as the expression x unless it throws an exception, in which case the *catch* behaves as the handler expression y .

To formalise the meaning of this language, we first recall the *Maybe* type

$$\mathbf{data} \text{ Maybe } a = \text{Nothing} \mid \text{Just } a$$

That is, a value of type *Maybe a* is either *Nothing*, which we think of as an exceptional value, or has the form *Just x* for some *x* of type *a*, which we think of as a normal value (Spivey, 1990). For the purposes of our example, however, we only require the type *Maybe Int*, which we abbreviate as *Mint*

type *Mint* = *Maybe Int*

Using this type, it is straightforward to define a function that evaluates expressions, and takes care of the necessary propagation and handling of exceptions

```

eval          :: Expr → Mint
eval (Val n)  = Just n
eval (Add x y) = case eval x of
                    Nothing → Nothing
                    Just n  → case eval y of
                                Nothing → Nothing
                                Just m  → Just (n + m)

eval (Throw)  = Nothing
eval (Catch x y) = case eval x of
                    Nothing → eval y
                    Just n  → Just n

```

Because of the repeated tagging and untagging of values of type *Mint*, this definition may be rather inefficient. We now show how it can be transformed into a *worker* that uses two continuations to make normal and exceptional control flow explicit (and hence eliminate the use of tags), together with a *wrapper* that takes care of the initial setup. Once again, the first step is to redefine *eval* using *fix*

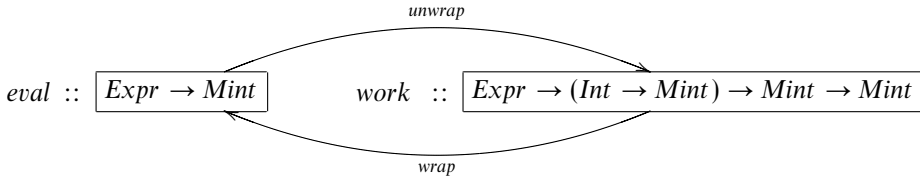
```

eval          :: Expr → Mint
eval          = fix body
body f        :: (Expr → Mint) → (Expr → Mint)
body f (Val n) = Just n
body f (Add x y) = case f x of
                    Nothing → Nothing
                    Just n  → case f y of
                                Nothing → Nothing
                                Just m  → Just (n + m)

body f (Throw)  = Nothing
body f (Catch x y) = case f x of
                    Nothing → f y
                    Just n  → Just n

```

One might expect now to aim to transform the computation *eval* of type *Expr* → *Mint* into a worker of type *Expr* → *Cont Mint*. In practice, however, the desire to have separate continuations for normal and exceptional control flow, sometimes called *double-barrelled* continuations (Thielecke, 2002), means that we start by splitting the underlying type *Mint* → *Mint* of continuations into two parts, and aim for a worker of the type illustrated in the following diagram:



Intuitively, the worker will apply a success continuation of type $Int \rightarrow Mint$ if the result of evaluating the given expression is an integer, or a failure continuation of type $Mint$ (which can be viewed as a function with no arguments) if the result is an exception. The necessary conversion functions, satisfying the worker/wrapper assumption $wrap \circ unwrap = id_{Expr \rightarrow Mint}$, can be defined as follows:

$$\begin{aligned}
 unwrap\ g\ e\ s\ f &= \mathbf{case}\ (g\ e)\ \mathbf{of} \\
 &\quad \mathit{Nothing} \rightarrow f \\
 &\quad \mathit{Just}\ n \rightarrow s\ n \\
 wrap\ h\ e &= h\ e\ \mathit{Just}\ \mathit{Nothing}
 \end{aligned}$$

Now that we have established the preconditions for the worker/wrapper transformation, applying this transformation gives

$$\begin{aligned}
 eval &:: Expr \rightarrow Mint \\
 eval &= wrap\ work \\
 work &:: Expr \rightarrow (Int \rightarrow Mint) \rightarrow Mint \rightarrow Mint \\
 work &= unwrap\ (body\ (wrap\ work))
 \end{aligned}$$

Making the expression argument explicit in the definition for $eval$, and expanding out $wrap$, we obtain the wrapper for the continuation-passing version of $eval$, which simply supplies Just and $\mathit{Nothing}$ as the initial continuations

$$\begin{aligned}
 eval &:: Expr \rightarrow Mint \\
 eval\ e &= work\ e\ \mathit{Just}\ \mathit{Nothing}
 \end{aligned}$$

Now let us focus our attention on simplifying the definition for the worker function. First of all, we redefine $work$ using pattern matching

$$\begin{aligned}
 work\ (Val\ n)\ s\ f &= s\ n \\
 work\ (Add\ x\ y)\ s\ f &= \mathbf{case}\ wrap\ work\ x\ \mathbf{of} \\
 &\quad \mathit{Nothing} \rightarrow f \\
 &\quad \mathit{Just}\ n \rightarrow \mathbf{case}\ wrap\ work\ y\ \mathbf{of} \\
 &\quad\quad \mathit{Nothing} \rightarrow f \\
 &\quad\quad \mathit{Just}\ m \rightarrow s\ (n + m) \\
 work\ (Throw)\ s\ f &= f \\
 work\ (Catch\ x\ y)\ s\ f &= \mathbf{case}\ wrap\ work\ x\ \mathbf{of} \\
 &\quad \mathit{Nothing} \rightarrow \mathbf{case}\ wrap\ work\ y\ \mathbf{of} \\
 &\quad\quad \mathit{Nothing} \rightarrow f \\
 &\quad\quad \mathit{Just}\ n \rightarrow s\ n \\
 &\quad \mathit{Just}\ n \rightarrow s\ n
 \end{aligned}$$

Now we simplify the case analyses used in the recursive definition by exposing the implicit use of *unwrap* and then fusing again with *wrap*. In particular, for an arbitrary expression x , and continuations g and s of the appropriate types

$$\begin{aligned}
 & \text{case } \textit{wrap work } x \text{ of} \\
 & \quad \textit{Nothing} \rightarrow g \\
 & \quad \textit{Just } n \rightarrow s \ n \\
 = & \quad \{ \textit{unapplying } \textit{unwrap} \} \\
 & \quad \textit{unwrap } (\textit{wrap work}) \ x \ s \ g \\
 = & \quad \{ \textit{worker/wrapper fusion} \} \\
 & \quad \textit{work } x \ s \ g
 \end{aligned}$$

Applying this result three times gives our final definition for the worker, in which control flow is now managed by explicit success and failure continuations, rather than by repeated tagging and untagging of *Maybe* values

$$\begin{aligned}
 \textit{work} & \quad \quad \quad :: \textit{Expr} \rightarrow (\textit{Int} \rightarrow \textit{Mint}) \rightarrow \textit{Mint} \rightarrow \textit{Mint} \\
 \textit{work } (\textit{Val } n) \ s \ f & \quad = s \ n \\
 \textit{work } (\textit{Add } x \ y) \ s \ f & \quad = \textit{work } x \ (\lambda n \rightarrow \textit{work } y \ (\lambda m \rightarrow s \ (n + m))) \ f \\
 \textit{work } (\textit{Throw}) \ s \ f & \quad = f \\
 \textit{work } (\textit{Catch } x \ y) \ s \ f & \quad = \textit{work } x \ s \ (\textit{work } y \ s \ f)
 \end{aligned}$$

Once again, note that everything follows in a straightforward manner once we have made the initial design decision to represent the type *Mint* using two continuations, and that the derivation does not require the use of induction. We conclude this section by noting that the resulting worker function can now readily be transformed into an abstract machine for evaluating expressions (Ager *et al.*, 2003; Hutton & Wright, 2006), resulting in an efficient machine that operates using two control stacks, one for normal evaluation and the other for handling exceptions.

8 Related work

As discussed in Section 1, the worker/wrapper transformation was first used to exploit strictness information in the Glasgow Haskell Compiler (Peyton Jones & Launchbury, 1991). However, the correctness of the resulting transformation has never been formally verified. During our work on this article we learned that the authors did in fact sketch a proof using the rolling rule, but this proof was never fully developed, nor was it published. The underlying strictness analyser in this compiler has been changed and improved several times since its initial implementation (Peyton Jones & Partain, 1993), but the use of the worker/wrapper technique to exploit this information remains exactly the same.

The worker/wrapper transformation has also been used to improve the shortcut approach to deforestation (Gill *et al.*, 1993), an optimisation technique that attempts to remove intermediate data structures from programs. In particular, the shortcut approach performs fusion when functions that produce and consume lists are written in a specific, idiomatic way, and are completely inlined. In the final chapter of his

PhD thesis (Gill, 1996), Gill describes a scheme for abstracting a function that produces a list in a way that utilised the worker/wrapper transformation to reduce the increase in code size that can result from wholesale inlining of definitions. In particular, Gill considered list producing functions of the form

$$f\ x_1 \dots x_n = \mathit{build}\ (\lambda c\ n \rightarrow \mathit{body})$$

which were then split into a wrapper and a worker

$$\begin{aligned} f\ x_1 \dots x_n &= \mathit{build}\ (\lambda c\ n \rightarrow \mathit{work}\ x_1 \dots x_n\ c\ n) \\ \mathit{work}\ x_1 \dots x_n\ c\ n &= \mathit{body} \end{aligned}$$

and finally, the new definition for f was inlined in the definition for work , thereby communicating the deforestation opportunity. However, no formal justification was given for the correctness of the worker/wrapper technique.

Chitil (2000a, 2000b) considerably extended Gill's work, by generalising to allow for multiple return lists, as directed by his type-inference based approach to deforestation. In particular, Chitil uses a static argument transformation (Santos, 1995) to turn a recursive function into a non-recursive function with local recursion, and splits the new, non-recursive function into a worker and wrapper, with detailed justification. He also explores the recursive instance of the worker/wrapper technique as applied to deforestation, acknowledging that 'whereas intuitively the semantic correctness of [this approach] is clear, a formal proof is hard', and suggests that a proof may be possible using improvement theory (Sands, 1998). However, as we have shown in this article, simpler techniques suffice.

Another application of the worker/wrapper transformation has been in changing the order of function arguments, to aid static analysis. In particular, Launchbury and Sheard (1995) use the technique to illustrate how the standard definition

$$\begin{aligned} \mathit{map} &:: (a \rightarrow b) \rightarrow [a] \rightarrow [b] \\ \mathit{map}\ f\ xs &= \mathbf{case}\ xs\ \mathbf{of} \\ &\quad [] \rightarrow [] \\ &\quad (x : xs) \rightarrow f\ x : \mathit{map}\ f\ xs \end{aligned}$$

can be translated into an equivalent version that is defined using a worker that takes the two arguments in the opposite order

$$\begin{aligned} \mathit{map} &:: (a \rightarrow b) \rightarrow [a] \rightarrow [b] \\ \mathit{map}\ f\ xs &= \mathit{work}\ xs\ f \\ \mathit{work} &:: [a] \rightarrow (a \rightarrow b) \rightarrow [b] \\ \mathit{work}\ xs\ f &= \mathbf{case}\ xs\ \mathbf{of} \\ &\quad [] \rightarrow [] \\ &\quad (x : xs) \rightarrow f\ x : \mathit{work}\ xs\ f \end{aligned}$$

It is interesting to note that the various users of the worker/wrapper transformation appear to realise its potential as a powerful tool for communicating a change of type through a recursive definition, but did not explore this generalisation.

The worker/wrapper terminology has also been used to describe a similar but distinct technique, namely when a wrapper of one type is exposed to an optimisation system for the purpose of applying a type conversion to a worker. Technically, this corresponds to the *adapter* design pattern (Gamma *et al.*, 1995). The rewrite system of the Glasgow Haskell Compiler (Peyton Jones *et al.*, 2001) uses this technique in its implementation of short-cut deforestation, as does the recent implementation of stream fusion (Coutts *et al.*, 2007). In both these systems, rather than wholesale inlining of a candidate function, an inlined wrapper expresses the opportunity for deforestation, but the actual computation is performed by the worker. In a real sense, these uses of workers and wrappers are part of the *worker/wrapper transformation*, which is the combination of *splitting* a function into a worker and wrapper in a correct and potentially useful manner, together with the *use* of these new functions to fulfill specific roles inside a rewrite system.

9 Conclusion and further work

In this article we presented the worker/wrapper transformation, and gave several examples of its use as an equational reasoning technique for improving the performance of functional programs. In particular, we showed how it can be used to derive more efficient programs based upon the use of accumulation, unboxed values, memoisation and continuations. Since these are all well-studied and widely used approaches to program optimisation, it is natural to ask what new value the worker/wrapper transformation brings? We see the following three primary benefits:

- It provides a general and systematic approach to transforming a computation of one type into an equivalent computation of another type. Such type transformations are pervasive in functional programming and reasoning, and in the efficient compilation of functional languages.
- It is straightforward to understand and apply, requiring only basic equational reasoning techniques, and often avoiding the need for induction. As such, it can readily be utilised by a wide spectrum of functional programmers as a simple but powerful technique for refactoring their programs.
- It captures many seemingly unrelated optimisation techniques inside a single unified framework, which may help reveal new connections between existing techniques, and the discovery of new techniques.

This article is by no means the end of the worker/wrapper story, but rather opens up a number of opportunities for further research. Interesting topics for further work include mechanising the technique in an equational reasoning system such as HERA (Gill, 2006), investigating how it can be automated in an optimising compiler such as GHC (Peyton Jones *et al.*, 2001), considering programs that utilise effects (Wadler, 1992a; McBride & Paterson, 2008) and more advanced applications (e.g. to programming languages themselves), exploring the use of specialised patterns of recursion (Gibbons & Jones, 1998; Hutton, 1999) and versions of the transformation (e.g. for changing the result type of a function) and further generalising the technique using category theory (Backhouse *et al.*, 1995).

Acknowledgments

We would like to thank Roland Backhouse, Olaf Chitil, Isaac Dupree, Conal Elliott, Peter Gammie, Peter Jonsson, John Matthews, Conor McBride, Simon Peyton Jones and Mark Tullsen for many useful comments and suggestions, Colin Runciman for the invitation to submit to JFP, and the three anonymous referees for their detailed reviews. Work on this article was initiated during a sabbatical visit to Galois, Inc. in Portland, Oregon by the second author during Summer 2007, for which funding from Galois is very gratefully acknowledged.

References

- Ager, Mads Sig, Biernacki, Dariusz, Danvy, Olivier & Midtgaard, Jan (2003) A functional correspondence between evaluators and abstract machines. In *Proceedings of the Fifth ACM-SIGPLAN International Conference on Principles and Practice of Declarative Programming* Uppsala, Sweden.
- Attenkirch, Thorsten (2001) Representations of first-order function types as terminal coalgebras. In *Typed Lambda Calculi and Applications*. LNCS, no. 2044. Berlin: Springer.
- Backhouse, Roland (2002) Galois connections and fixed point calculus. In *Algebraic and Coalgebraic Methods in the Mathematics of Program Construction*. LNCS Tutorial, vol. 2297. Berlin: Springer-Verlag.
- Backhouse, Roland, Bijsterveld, Marcel, van Geldrop, Rik & van der Woude, Jaap (1995) Categorical fixed point calculus. In *Proceedings of the Sixth International Conference on Category Theory and Computer Science*. Berlin: Springer-Verlag.
- Bird, Richard (1998) *Introduction to Functional Programming using Haskell*, second edition. New York: Prentice Hall.
- Burstall, Rod & Darlington, John (1977) A transformational system for developing recursive programs. *J. ACM* **24**, 44–67.
- Chitil, Olaf (2000a) *Type-Inference Based Deforestation of Functional Programs*. Ph.D. thesis, RWTH Aachen.
- Chitil, Olaf (2000b) Type-inference based short cut deforestation (nearly) without inlining. In *Proceedings of 11th International Workshop on Implementation of Functional Languages*, Chris Clack and Pieter Koopman (eds). LNCS, no. 1868. Berlin: Springer.
- Coutts, Duncan, Leshchinskiy, Roman & Stewart, Don (2007) Stream fusion: From lists to streams to nothing at all. In *Proceedings of the 2007 ACM SIGPLAN International Conference on Functional Programming*. New York: ACM Press.
- Gamma, Erich, Helm, Richard, Johnson, Ralph & Vlissides, John (1995) *Design Patterns*. Reading, Mass: Addison-Wesley Professional.
- Gibbons, Jeremy & Jones, Geraint (1998). The under-appreciated unfold. In *Proceedings of the Third ACM SIGPLAN International Conference on Functional Programming* Baltimore, Maryland.
- Gill, Andy (1996) *Cheap Deforestation for Non-Strict Functional Languages*. Ph.D. thesis, University of Glasgow.
- Gill, Andy (2006) Introducing the haskell equational reasoning assistant. In *Proceedings of the 2006 ACM SIGPLAN Workshop on Haskell*. New York: ACM Press.
- Gill, Andy, Launchbury, John & Peyton Jones, Simon (1993) A short cut to deforestation. In *Proceedings of the Conference on Functional Programming Languages and Computer Architecture*. New York: ACM Press.

- Hinze, Ralf (2008) Functional pearl: Streams and unique fixed points. In *Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming* Victoria, British Columbia.
- 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. *Info. Process. Lett.* **22**(3), pp 141–144.
- Hutton, Graham (1999) A tutorial on the universality and expressiveness of fold. *J. Funct. Prog.* **9**(4), pp 355–372.
- Hutton, Graham (2007) *Programming in Haskell*, Cambridge: Cambridge University Press.
- Hutton, Graham & Wright, Joel (2004) Compiling exceptions correctly. In *Proceedings of the Seventh International Conference on Mathematics of Program Construction*. LNCS, vol. 3125. Stirling, Scotland: Springer.
- Hutton, Graham & Wright, Joel (2006) Calculating an exceptional machine. Selected papers from the *Fifth Symposium on Trends in Functional Programming*, Munich, vol. 5, November 2004 UK: Intellect.
- Hutton, Graham & Wright, Joel (2007) What is the meaning of these constant interruptions? *J. Funct. Prog.* **17**(6), pp 777–792.
- Launchbury, John & Sheard, Tim (1995) Warm fusion: Deriving build-catas from recursive definitions. In *Proceedings of the Seventh ACM SIGPLAN/SIGARCH International Conference on Functional Programming Languages and Computer Architecture*. New York: ACM Press.
- McBride, Conor & Paterson, Ross (2008) Applicative programming with effects. *J. Funct. Prog.* **18**(1), pp 1–13.
- Meijer, Erik, Fokkinga, Maarten & Paterson, Ross (1991) Functional programming with bananas, lenses, envelopes and barbed wire. In *Proceedings of the Conference on Functional Programming and Computer Architecture*, Hughes, John (ed). LNCS, no. 523. Springer-Verlag.
- Michie, Donald (1968) Memo functions and machine learning. *Nature*, **218**, pp 19–22.
- Peyton Jones, Simon (2003) *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press. Also available on www.haskell.org/definition.
- Peyton Jones, Simon & Launchbury, John (1991) Unboxed values as first class citizens in a non-strict functional language. In *Proceedings of the Conference on Functional Programming and Computer Architecture*. Cambridge, MA: Springer-Verlag.
- Peyton Jones, Simon & Partain, Will (1993) Measuring the effectiveness of a simple strictness analyser. In *Proceedings of the 1993 Glasgow Workshop on Functional Programming*, Hammond, Kevin & O'Donnell, John (eds). Ayr, Scotland: Springer-Verlag.
- Peyton Jones, Simon, Tolmach, Andrew & Hoare, Tony (2001) Playing by the rules: Rewriting as a practical optimisation technique in GHC. In *Proceedings of the 2001 ACM SIGPLAN Workshop on Haskell*. ACM Press.
- Reynolds, John C. (1972) Definitional interpreters for higher-order programming languages. In *Proceedings of the ACM Annual Conference*. ACM Press.
- Sands, David (1998) Improvement theory and its applications. In *Higher Order Operational Techniques in Semantics*, Gordon, Andrew & Pitts, Andrew (eds). Cambridge, UK: Cambridge University Press.
- Santos, Andre (1995) *Compilation by Transformation in Non-strict Functional Languages*. Ph.D. thesis, University of Glasgow.
- Schmidt, David (1986) *Denotational Semantics: A Methodology for Language Development*. Allyn and Bacon, Inc.

- Spivey, Mike (1990) A functional theory of exceptions. *Sci. Comput. Prog.* **14**(1), pp 25–42.
- Thielecke, Hayo (2002) Comparing control constructs by double-barrelled CPS. *Higher-Order Symbol. Comput.* **15**(2/3), pp 141–160.
- Turner, David A. (1995) Elementary strong functional programming. In *Proceedings of the First International Symposium on Functional Programming Languages in Education*. LNCS, no. 1022. Springer-Verlag.
- Wadler, Philip (1992a) Monads for functional programming. In *Proceedings of the Marktoberdorf Summer School on Program Design Calculi*, Broy, Manfred (ed). Springer-Verlag.
- Wadler, Philip (1992b) The essence of functional programming. *Proc. Principles Prog. Lang.* New York: ACM Press, pp 1–14.

Appendix: fixed points and the rolling rule

For completeness, this appendix briefly reviews the necessary semantic theory of fixed points that underlies our formalisation of the worker/wrapper transformation, and presents a proof of the rolling rule in this context.

Our formalisation is based upon the domain-theoretic approach to semantics (Schmidt, 1986), in which the basic idea is that types are *complete partial orders* (cpo), that is, sets with a partial-ordering \sqsubseteq , the least element \perp , and limits of all non-empty chains. In turn, programs are *continuous functions*, that is, functions between cpo that preserve the partial-order and limit structure.

Now consider a recursive equation $x = f x$ that defines a value x in terms of itself and some continuous function f . A well-known fixed point theorem states that this equation has the least solution for x , denoted by $fix f$ and called the *least fixed point* of f , which is adopted as the semantics of the definition. Moreover, $fix f$ is constructed as the limit of the following infinite chain:

$$\perp \sqsubseteq f \perp \sqsubseteq f (f \perp) \sqsubseteq f (f (f \perp)) \sqsubseteq \dots$$

As a simple example of this approach, consider the equation $ones = 1 : ones$ that defines the infinite list $1 : 1 : 1 : \dots$. This definition can be rewritten as $ones = f ones$, where f is the function defined by $f xs = 1 : xs$. Hence, the semantics of the definition is given by $ones = fix f$, and by the fixed point theorem is constructed as the limit of the infinite chain of partial lists containing increasing numbers of ones

$$\perp \sqsubseteq 1 : \perp \sqsubseteq 1 : 1 : \perp \sqsubseteq 1 : 1 : 1 : \perp \sqsubseteq \dots$$

The fixed point approach to recursive definitions can be realised in Haskell by defining an explicit version of the function fix as follows:

$$\begin{aligned} fix &:: (a \rightarrow a) \rightarrow a \\ fix f &= f (fix f) \end{aligned}$$

For example, evaluating the expression $fix (\lambda xs \rightarrow 1 : xs)$ using this definition gives the infinite list of ones, $1 : 1 : 1 : \dots$, as expected.

In order to prove the rolling rule, which states that $fix (g \circ f) = g (fix (f \circ g))$ for any functions f and g of the appropriate types, we exploit the following two fundamental properties of fixed points (Backhouse, 2002):

$$\begin{aligned} fix f &= f (fix f) && \text{(computation)} \\ f x \sqsubseteq x &\Rightarrow fix f \sqsubseteq x && \text{(induction)} \end{aligned}$$

The first property states that $\text{fix } f$ is a fixed point of f (an expression x satisfying $f x = x$), while the second states that $\text{fix } f$ is the least *prefix* point of f (the least expression x satisfying $f x \sqsubseteq x$). One might have expected in the latter case to state that $\text{fix } f$ is the least *fixed* point rather than just the least prefixed point, but these two notions can be shown to be equivalent, and the above formulation has the advantage of being more useful for reasoning purposes. Using these two properties, the rolling rule can now be verified by mutual inclusion

$$\begin{aligned}
 & \text{fix } (g \circ f) \sqsubseteq g (\text{fix } (f \circ g)) \\
 \Leftarrow & \quad \{ \text{induction} \} \\
 & (g \circ f) (g (\text{fix } (f \circ g))) \sqsubseteq g (\text{fix } (f \circ g)) \\
 \Leftrightarrow & \quad \{ \text{applying } \circ \} \\
 & g (f (g (\text{fix } (f \circ g)))) \sqsubseteq g (\text{fix } (f \circ g)) \\
 \Leftrightarrow & \quad \{ \text{unapplying } \circ \} \\
 & g ((f \circ g) (\text{fix } (f \circ g))) \sqsubseteq g (\text{fix } (f \circ g)) \\
 \Leftrightarrow & \quad \{ \text{computation} \} \\
 & g (\text{fix } (f \circ g)) \sqsubseteq g (\text{fix } (f \circ g)) \\
 \Leftrightarrow & \quad \{ \text{reflexivity} \} \\
 & \text{True}
 \end{aligned}$$

and

$$\begin{aligned}
 & g (\text{fix } (f \circ g)) \sqsubseteq \text{fix } (g \circ f) \\
 \Leftrightarrow & \quad \{ \text{computation} \} \\
 & g (\text{fix } (f \circ g)) \sqsubseteq (g \circ f) (\text{fix } (g \circ f)) \\
 \Leftrightarrow & \quad \{ \text{applying } \circ \} \\
 & g (\text{fix } (f \circ g)) \sqsubseteq g (f (\text{fix } (g \circ f))) \\
 \Leftarrow & \quad \{ g \text{ is monotonic} \} \\
 & \text{fix } (f \circ g) \sqsubseteq f (\text{fix } (g \circ f)) \\
 \Leftrightarrow & \quad \{ \text{above result, swapping } f \text{ and } g \} \\
 & \text{True}
 \end{aligned}$$

We conclude by noting that the above proof of the rolling rule only relies on the basic notion of monotonic functions on partially ordered sets, rather than the stronger notion of continuous functions on cpos. However, working in this stronger setting automatically guarantees that the necessary fixed points always exist, by virtue of the fixed point theorem described above, which in weaker settings may become a side condition on the rolling rule.