

Fusion in Practice

Diederik van Arkel, John van Groningen, and Sjaak Smetsers

Computing Science Institute
University of Nijmegen
Toernooiveld 1, 6525 ED Nijmegen, The Netherlands
`diederik@cs.kun.nl`, `johnvg@cs.kun.nl`, `sjakie@cs.kun.nl`

Abstract. Deforestation was introduced to eliminate intermediate data structures used to connect separate parts of a functional program together. A more sophisticated technique, called fusion, tries to achieve the same goal but with better results. Fusion is based on a consumer-producer model indicating pairs of functions that can be combined into a single function. In this paper we extend this fusion algorithm by refining this model and by adding new transformation rules. The extended fusion algorithm is able to deal with standard deforestation but also with higher-order removal and dictionary elimination. We have implemented this method in the Clean 2.0 compiler.

1 Introduction

Static analysis techniques, such as typing and strictness analysis are crucial components of state-of-the-art implementations of lazy functional programming languages. These techniques are employed to determine properties of functions in a program, to be used by the programmer and also by the compiler itself. The growing complexity of functional languages like Haskell [Has92] and Clean [Cle13,Cle20] requires increasingly sophisticated methods for translating programs written in these languages into efficient executables. Often these optimization methods are implemented in an ad hoc manner: new language features seem to require new optimization techniques which are implemented simultaneously, or added later when it is noticed that the use of these features leads to inefficient code. For instance, type classes require the elimination of dictionaries, monadic programs introduce a lot of higher-order functions that have to be removed, and the intermediate data structures that are built due to function composition should be avoided.

In Clean 1.3 most of these optimizations were implemented independently. They also occurred at different phases during the compilation process making it difficult to combine them into a single optimization phase. Finally, the removal of auxiliary data structures was not implemented at all.

In the Clean 2.0 compiler one combined method is used to perform various optimizations. This method is based on Chin's *fusion algorithm* [Chin94], which in its turn was inspired by Wadler's *deforestation algorithm* [Wad88,Fer88]. We start with a few examples illustrating what types of optimizations can be performed.

1.1 Deforestation

The following example is given in both [Chin94] and [Wad88]. It elegantly shows that a programmer no longer need worry about annoying questions such as “Which of the following two expressions is more efficient?”

$$\text{append}(\text{append}(a, b), c) \quad \text{or} \quad \text{append}(a, \text{append}(b, c))$$

Experienced programmers almost automatically use the second expression, but from a more abstract point of view there seems to be no difference.

Deforestation as well as fusion will transform the first expression into `app_app(a, b, c)` and introduce a new function in which the two applications of `append` are combined:

$$\begin{aligned} \text{append}(f, t) &= \text{case } f \text{ of} \\ &\quad \text{Nil} \quad \quad \quad | t \\ &\quad \text{Cons}(x, xs) | \text{Cons}(x, \text{append}(xs, t)) \\ \\ \text{app_app}(a, b, c) &= \text{case } a \text{ of} \\ &\quad \text{Nil} \quad \quad \quad | \text{append}(b, c) \\ &\quad \text{Cons}(x, xs) | \text{Cons}(x, \text{app_app}(xs, b, c)) \end{aligned}$$

Transforming the second expression introduces essentially the same function both using Wadler’s deforestation and by Chin’s fusion. However, this saves only one evaluation step compared to the original expression at the cost of an extra function. Our fusion algorithm transforms the first expression just as deforestation or fusion would, but it leaves the second expression unchanged.

The major difficulty with both deforestation and fusion is to determine which applications can be fused (or deforested) safely. Without any precautions there are many situations in which the transformation will not terminate. Therefore it is necessary to formulate proper criteria that, on the one hand, guarantee termination, and on the other hand, do not reject too many fusion candidates. Non-termination can occur when so-called *accumulating parameters* are involved. We will have a closer look at this notion in section 4.1, but the absence of accumulating parameters is not enough. The transformation should also not introduce repeated computations, by duplicating redexes. This leads to the linearity requirement of functions: the variables in a function body should not appear more than once.

1.2 Type classes and dictionaries

Type classes are generally considered to be one of the most powerful concepts of functional languages. Firstly, type classes allow the language designer to support its implementation with a collection of libraries in which it is no longer necessary to use different names for conceptually the same operations. Secondly, it eases the definition of auxiliary operations on complex data structures. Both advantages are illustrated with the following example in which an instance of equality for lists is declared. Here we use the Clean syntax (deviating slightly from the Haskell notation).

```

instance == [a] | == a
where
    (==) [] [] = True
    (==) [x:xs] [y:ys] = x == y && xs == ys
    (==) _ _ = False

```

With type classes one can use the same name (==) for defining equality over all kinds of different types. In the body of an instance one can use the overloaded operation itself to compare substructures, i.e. it is not necessary to indicate the difference between both occurrences of == in the right-hand side of the second alternative. The translation of such instances into ‘real’ functions is easy: Each *context restriction* is converted into a *dictionary* argument containing the concrete version of the overloaded class. For the equality example this leads to

```

eqList d [] [] = True
eqList d [x:xs] [y:ys] = d.== (x,y) && eqList d xs ys
eqList d _ _ = False

```

The expression `d.==` denotes the selection of the == operation from the dictionary `d`. An application of == to two lists of integers, e.g. `[1,2] == [1,2]`, is replaced by an expression containing the list version of equality parameterized with the integer dictionary of the equality class, `eqList eqInt [1,2] [1,2]`.

Applying this simple compilation scheme introduces a lot of overhead which can be eliminated by specializing `eqList` for the `eqInt` dictionary as shown below

```

eqListeqInt [] [] = True
eqListeqInt [x:xs] [y:ys] = eqInt (x,y) && eqListeqInt xs ys
eqListeqInt _ _ = False

```

In Clean 1.3 the specialization of overloaded operations within a single module was performed immediately, i.e. dictionaries were not built at all, except for some rare, exotic cases. These exceptions are illustrated by the following type declaration (taken from [Oka98])

```

:: Seq a = Nil | Cons a (Seq [a])

```

Defining an instance of == is easy, specializing such an instance for a concrete element type cannot be done. The compiler has to recognize such situations in order to avoid an infinite specialization loop.

In Clean 2.0 specialization is performed by the fusion algorithm. The handling of the special unsolvable cases is trivial because these cases involve functions with accumulating arguments and will therefore not be accepted as fusion candidates. The linearity requirement is not necessary because dictionaries do not contain unevaluated expressions (*closures*), so copying dictionaries does not duplicate computations.

1.3 Higher-order functions

Although dictionary elimination can be seen as a form of higher-order function removal, we describe these optimizations separately. A straightforward treatment of higher-order functions introduces overhead both in time and space. E.g. measurements on large programs using a monadic style of programming show that such overhead can be large; see section 6.

In [Wad88] Wadler introduces *higher-order macros* to eliminate some of this overhead but this method has one major limitation: these macros are not considered first class citizens. Chin has extended his fusion algorithm so that it is able to deal with higher-order functions. We adopt his solution with some minor refinements.

Again accumulating parameters are a source of non-termination. Consider the following function definitions:

```
twice f x = f (f x)
```

```
f g = f (twice g)
```

The parameter of `f` is accumulating: the argument in the recursive call of `f` is 'larger' than the original argument. Trying to fuse `f` with `inc` (for some producer `inc`) in the application `f inc` will lead to an a new application of the form `f twice.inc`. Fusing this one leads to the expression `f twice_twice.inc` and so on.

Partial function application should also be treated with care. At first one might think that it is safe to fuse an application `f (g E)` in which the arity of `f` is greater than one and the subexpression `g E` is a redex. This fusion will combine `f` and `g` into a single function, say `f_g`, and replace the original expression by `f_g E`. This, however, is dangerous if the original expression was shared, as shown by the following function `h`:

```
h = z 1 + z 2
  where z = f (g E)
```

This function is not equivalent to the version in which `f` and `g` have been fused:

```
h = z 1 + z 2
  where z = f_g E
```

Here the computation encoded in the body of `g` will be performed twice, as compared to only once in the original version.

1.4 Generic functions

Generic programming allows the programmer to write a function once and use it many times even for different types. It relieves the programmer from having to define new instances of common operations each time he declares a new data type. The idea is to consider types as being built up from a small fixed set of type

constructors and to specify generic operations in terms of these constructors. In Clean 2.0, for example, one can specify all instances of equality by just a few lines of fairly obvious code:

```
generic eq a :: a a -> Bool

instance eq UNIT where
    eq x y = True
instance eq PAIR where
    eq eqx eqy (PAIR x1 y1) (PAIR x2 y2)
        = eqx x1 x2 && eqy y1 y2
instance eq EITHER where
    eq eql eqr (LEFT l1) (LEFT l2)
        = eql l1 l2
    eq eql eqr (RIGHT r1) (RIGHT r2)
        = eqr r1 r2
    eq eql eqr _ _
        = False
```

Here `UNIT`, `PAIR` and `EITHER` are the fixed generic types. With the aid of this generic specification, the compiler is able to generate instances for any algebraic data type. The idea is to convert an object of such a data type to its generic representation (this encoding follows directly from the definition of the data type), apply the generic operation to this converted object and, if necessary, convert the object back to the original data type. For a comprehensive description of how generics can be implemented, see [Ali01] or [Hin00].

Without any optimizations one obtains operations which are very inefficient. The conversions and the use of higher-order functions, e.g. the `eqx` and `eqy` arguments of the instance for `PAIR`, introduce a lot of overhead. Fortunately the combined data and higher-order fusion appears to be sufficient to get rid of all intermediate data and all higher-order calls leading to specialized operations that are as efficient as the hand coded versions. The only trivial extension of the fusion algorithm needed to achieve this is to allow constructors as producers.

2 Overview

The underlying idea for transformation algorithms like Wadler's deforestation or Chin's fusion is to combine nested applications of functions of the form $F(\dots, G\vec{E}, \dots)$ ¹ into a single application $F + G(\dots, \vec{E}, \dots)$, by performing a sequence of unfold steps of both F and G . Of course, if one of the functions involved is recursive this sequence is potentially infinite. To avoid this it is necessary that during the sequence of unfold steps an application is reached that has been encountered before. In that case one can perform the crucial fold step. But how do we know that we will certainly reach such an application?

¹ We write \vec{E} as shorthand for (E_1, \dots, E_n)

Wadler’s solution is to define the notion of *treeless form*. If the above F and G are treeless it is guaranteed that no infinite unfolding sequences will occur. However, Wadler does not distinguish between F and G . Chin recognizes that the roles of these functions in the fusion process are different. He comes up with the so called producer-consumer model: F plays the role of *consumer*, consuming data through one of its arguments, whereas G acts as a *producer*, producing data via its result. Separate safety criteria can then be applied for the different roles. In [Gil93] a restricted version of this model is used in which only a small fixed set of operators are identified as consumers and producers.

Although Chin’s criterion indicates more fusion candidates than Wadler’s, it is still too restrictive to be useful for transforming average functional programs.

To illustrate these shortcomings we first repeat Chin’s notion of safety: A function F is a safe consumer in its i^{th} argument if all recursive calls of F have either a variable or a constant on the i^{th} parameter position, otherwise it is *accumulating* in that argument. A function G is a safe producer if none of its recursive calls appears on a consuming position.

One of the drawbacks of the safety criterion for consumers is that it indicates too many consuming parameters, and consequently it limits the number of producers (since the producer property negatively depends on the consumer property). As an example, consider the following definition for `flatten`:

$$\text{flatten}(l) = \text{case } l \text{ of}$$

Nil		Nil
Cons(x, xs)		append($x, \text{flatten}(xs)$)

According to Chin, the `append` function is consuming in both of its arguments. Consequently, the `flatten` function is not a producer, for, its recursive call appears on a consuming position of `append`. Wadler will also reject `flatten` because its definition is not treeless. In our definition of consumer we will introduce an auxiliary notion, called *active arguments*, that filters out the arguments that will not lead to a fold step, like the second argument of `append`. If `append` is no longer consuming in its second argument, `flatten` becomes a decent producer.

Chin also indicates superfluous consuming arguments when we are not dealing with a single recursive function but with a set of mutually recursive functions. To illustrate this, consider the unary functions f and g being mutually recursive as follows:

$$\begin{aligned} f(x) &= g(x) \\ g(x) &= f(h(x)) \end{aligned}$$

Now f is accumulating and g is not (e.g. g ’s body contains a call to f with an accumulating argument whereas f ’s body just contains a simple call to g). Although g is a proper consumer, it will make no sense to fuse an application of g with a producer, for this producer will be passed to f but cannot be fused with f . Again no fold step will take place. Consequently, a function of which the recursive call appears as an argument of g will be unjustly rejected as a producer.

3 Syntax

We will focus on program transformations in functional programming languages by describing a ‘core language’ capturing the essential aspects such as pattern matching, sharing and higher-order functions.²

Functional expressions are built up from applications of function symbols F and data constructors C .

Pattern matching is expressed by a construction `case ... of ...`. In function definitions, one can express pattern matching with respect to one argument at a time. This means that compound patterns are decomposed into nested (‘sequentialized’) `case` expressions.

Sharing of objects is expressed by a `let` construction and higher-order applications by an `@`.

We do not allow functions that contain `case` expressions as nested subexpressions on the right-hand side. And as in [Chin94], we distinguish so called *g-type functions* (starting with a single pattern match) and *f-type functions* (with no pattern match at all).

Definition 1. (i) *The set of expressions is defined by the following grammar. Below, x ranges over variables, C over constructors and F over function symbols.*

$$\begin{aligned}
 E ::= & x \\
 & | C(E_1, \dots, E_k) \\
 & | F(E_1, \dots, E_k) \\
 & | \text{let } \vec{x} = \vec{E} \text{ in } E' \\
 & | \text{case } E \text{ of } P_1 | E_1 \dots P_n | E_n \\
 & | E @ E' \\
 P ::= & C(x_1, \dots, x_k)
 \end{aligned}$$

(ii) *The set of free variables (in the obvious sense) of E is denoted by $\text{FV}(E)$. An expression E is said to be open if $\text{FV}(E) \neq \emptyset$, otherwise E is called closed.*

(iii) *A function definition is an expression of the form*

$$F(x_1, \dots, x_k) = E$$

where all the x_i 's are disjoint and $\text{FV}(E) \subset \{x_1, \dots, x_k\}$.

4 Fusion Algorithm

4.1 Consumers

We say that an occurrence of variable x in E is active if x is either a pattern matching variable (`case x of ...`), a higher-order variable ($x @ \dots$), or x is used as an argument on an active position of a function. The idea is to define the notion of ‘active occurrence’ $\text{actocc}(x, E)$ and ‘active position’ $\text{act}(F)_i$ simultaneously as the least solution of some predicate equations.

² We leave out the treatment of dictionaries because it very much resembles the way other constructs are analyzed and transformed

Definition 2. (i) The predicates *actocc* and *act* are specified by mutual induction as follows.

$$\begin{aligned}
\text{actocc}(x, y) &= \text{true if } y = x \\
&= \text{false otherwise} \\
\text{actocc}(x, F\vec{E}) &= \text{true if for some } i \ E_i = x \text{ and } \text{act}(F)_i \\
&= \bigvee_i \text{actocc}(x, E_i) \text{ otherwise} \\
\text{actocc}(x, C\vec{E}) &= \bigvee_i \text{actocc}(x, E_i) \\
\text{actocc}(x, \text{case } E \text{ of } \dots P_i | E_i \dots) &= \text{true if } E = x \\
&= \bigvee_i \text{actocc}(x, E_i) \text{ otherwise} \\
\text{actocc}(x, \text{let } \vec{x} = \vec{E} \text{ in } E') &= \bigvee_i \text{actocc}(x, E_i) \text{ or } \text{actocc}(x, E') \\
\text{actocc}(x, E \text{ @ } E') &= \text{true if } E = x \\
&= \text{actocc}(x, E) \text{ or } \text{actocc}(x, E') \text{ otherwise}
\end{aligned}$$

Moreover, for each function F , defined by $F\vec{x} = E$

$$\text{act}(F)_i \Leftrightarrow \text{actocc}(x_i, E)$$

(ii) We say that F is active in argument i if $\text{act}(F)_i$ is true.

The notion of *accumulating parameter* is borrowed from [Chin94] and slightly modified.

Definition 3. Let F_1, \dots, F_n be a set of mutually recursive functions with respective right-hand sides E_1, \dots, E_n . The function F_j is accumulating in its i^{th} parameter if either

- (1) there exists a right-hand side E_k containing an application $F_j(\dots, E'_i, \dots)$ in which E'_i is open and not just an argument or a pattern variable, or
- (2) the right-hand side of F_j , E_j , contains an application $F_k(\dots, E'_l, \dots)$ such that $E'_l = x_i$ and F_k is accumulating in l .

The first requirement corresponds to Chin's notion of accumulating parameter. The second requirement will prevent functions that recursively depend on other accumulating functions from being regarded as non-accumulating. Combining active and accumulating leads to the notion of consuming.

Definition 4. A function F is consuming in its i^{th} parameter if it is both non-accumulating and active in i .

4.2 Producers

The notion of *producer* is also taken from [Chin94] again with a minor adjustment to deal with constructors.

Definition 5. Let F_1, \dots, F_n be a set of mutually recursive functions. These functions are called *producers* if none of their recursive calls (in the right-hand sides of their definitions) occur on a consuming position. Constructors and partial applications (of functions and constructors) are always considered to be producers.

4.3 Linearity

The notion of *linearity* is identical to that used by Chin.

Definition 6. Let F be a function with definition $F(x_1, \dots, x_n) = E$. The function F is linear in its i^{th} parameter if

- (1) F is an f -type function and x_i occurs at most once in E , or
- (2) F is a g -type function and x_i occurs at most once in each of the branches of the top-level case .

4.4 Transformation rules

Our adjusted version of the fusion algorithm consists of one general transformation rule, dealing with all expressions, and three auxiliary rules for function applications, higher-order application, and for case expressions. These are essentially the same as the rules presented in [Chin94]. The idea is that during fusion both consumer and producer have to be unfolded and combined. This combination forms the body of the newly generated function. Sometimes however, it appears to be more convenient if the unfold step of the consumer could be undone, in particular if the consumer and the producer are both g -type functions. For this reason we supply some of the case expressions with the function symbol to which it corresponds (case_F). Note that this correspondence is always unique because g -type functions contain exactly one case on their right-hand side.

We use $F+_iS$ as a name for the function that results from fusing F that is consuming in its i^{th} argument, with producer S . Suppose F is defined as $F\vec{x} = E$. Then the resulting function is defined, distinguishing three cases

1. S is a fully applied function and F is a g -type function: The resulting function consists of the unfoldings of F and S . The top-level case is annotated.
2. S is a partially applied function or F is a f -type function: The resulting function consists of the unfolding of the F with the application of S substituted for formal argument i .
3. S is a constructor (either fully or partially applied): The resulting function consists of the unfolding of the F with the application of S substituted for formal argument i .

Note that we generate each $F+_iS$ only once.

We use F^+ as a name for the function that results from raising the arity of F by one.

Definition 7. *Transformation rules for first and higher-order expressions*

$$\begin{aligned}
T[x] &= x \\
T[C\vec{E}] &= CT[\vec{E}] \\
T[F\vec{E}] &= \mathcal{F}[\mathcal{F}T[\vec{E}]] \\
T[\text{case } E \text{ of } \dots P_i | E_i \dots] &= C[\text{case } T[E] \text{ of } \dots P_i | E_i \dots] \\
T[\text{let } \vec{x} = \vec{E} \text{ in } E'] &= \text{let } \vec{x} = T[\vec{E}] \text{ in } T[E'] \\
T[E \textcircled{E} E'] &= \mathcal{H}[T[E] \textcircled{E} T[E']] \\
\\
\mathcal{F}[F(E_1, \dots, E_i, \dots, E_m)] &= \mathcal{F}[F_{+i}S(E_1, \dots, E_{i-1}, E'_1, \dots, E'_n, E_{i+1}, \dots, E_m)] \\
&\quad \text{if, for some } i, F \text{ is consuming,} \\
&\quad E_i = S(E'_1, \dots, E'_n) \text{ with } S \text{ a producer and} \\
&\quad F \text{ is linear in } i \text{ if } S \text{ is a function} \\
&= F(E_1, \dots, E_i, \dots, E_m) \text{ otherwise} \\
\\
\mathcal{H}[C(E_1, \dots, E_k) \textcircled{E}] &= C(E_1, \dots, E_k, E) \\
\mathcal{H}[F(E_1, \dots, E_k) \textcircled{E}] &= \mathcal{F}[F(E_1, \dots, E_k, E)] \\
&\quad \text{if } \text{arity}(F) > k \\
&= \mathcal{F}[F^+(E_1, \dots, E_k, E)] \text{ otherwise} \\
\\
C[\text{case}_F G(E_1, \dots, E_n) \text{ of } \dots] &= \mathcal{F}[F(x_1, \dots, x_{i-1}, G(E_1, \dots, E_n), x_{i+1}, \dots, x_n)] \\
&\quad \text{where} \\
&\quad F(x_1, \dots, x_n) = \text{case } x_i \text{ of } \dots \\
C[\text{case}_F C_i(E_1, \dots, E_n) \text{ of } \dots C_i(x_1, \dots, x_n) | E'_i \dots] &= T[E'_i[E_1/x_1, \dots, E_n/x_n]] \\
C[\text{case}_F(\text{case } E \text{ of } \dots P_i | E'_i \dots) \text{ of } \dots] &= \text{case } T[E] \text{ of } \dots P_i | E''_i \dots \\
&\quad \text{where } E''_i = \mathcal{F}[F(x_1, \dots, x_{i-1}, T[E'_i], x_{i+1}, \dots, x_n)]^3 \\
&\quad \text{and} \\
&\quad F(x_1, \dots, x_n) = \text{case } x_i \text{ of } \dots \\
C[\text{case}_{Fx} \text{ of } \dots P_i | E_i \dots] &= \text{case}_{Fx} \text{ of } \dots P_i | T[E_i] \dots
\end{aligned}$$

³ A minor improvement here is obtained by examining the expression E'_i . If this expression starts with a constructor, it is better to perform the pattern match instead of undoing the unfold step of F

Definition 8. *Rules for introducing new functions.*

$$\begin{aligned}
& F+_iG(x_1, \dots, x_{i-1}, y_1, \dots, y_m, x_{i+1}, \dots, x_n) \\
& \quad = \mathcal{T}[[E[E'/x_i]], \quad \text{if } \text{arity}(G) = m \text{ and } F \text{ is a } g\text{-type function} \\
& \quad \quad \quad \text{where } G\vec{y} = E' \\
& \quad = \mathcal{T}[[E[G(y_1, \dots, y_m)/x_i]], \quad \text{otherwise} \\
& F+_iC(x_1, \dots, x_{i-1}, y_1, \dots, y_m, x_{i+1}, \dots, x_n) \\
& \quad = \mathcal{T}[[E[C(y_1, \dots, y_m)/x_i]] \\
\\
& F^+(x_1, \dots, x_n, x) \\
& \quad = \mathcal{T}[[\mathcal{R}_x[E]]] \\
\\
& \mathcal{R}_y[[x]] \\
& \quad = x \textcircled{y} \\
& \mathcal{R}_y[[F(E_1, \dots, E_n)]] \\
& \quad = F(E_1, \dots, E_n, y), \quad \text{if } \text{arity}(F) > n \\
& \quad = F(E_1, \dots, E_n) \textcircled{y} \quad \text{otherwise} \\
& \mathcal{R}_y[[\text{let } \vec{x} = \vec{E} \text{ in } E']] \\
& \quad = \text{let } \vec{x} = \vec{E} \text{ in } \mathcal{R}_y[[E']] \\
& \mathcal{R}_y[[\text{case } E \text{ of } \dots P_i | E_i \dots]] \\
& \quad = \text{case } E \text{ of } \dots P_i | \mathcal{R}_y[[E_i]] \dots \\
& \mathcal{R}_y[[E \textcircled{E'}]] \\
& \quad = (E \textcircled{E'}) \textcircled{y}
\end{aligned}$$

5 Examples

We now present a few examples of fusion using the adjusted transformation rules.

5.1 Deforestation

We start with a rather trivial example involving the functions **Append**, **Flatten** and **Reverse**. The first two functions have been defined earlier. We will use three different versions of **Reverse** to illustrate the effect of alternative but equivalent function definitions on the fusion result. The applications being fused are the right-hand sides of the functions called **test1**, **test2** and **test3**.

The first definition of **Reverse** uses a helper function with an explicit accumulator:

$$\begin{aligned}
\text{Reverse1}(l) &= \text{Rev}(l, \text{Nil}) \\
\\
\text{Rev}(l, a) &= \text{case } l \text{ of} \\
& \quad \text{Nil} \quad \quad \quad | a \\
& \quad \text{Cons}(x, xs) \quad | \text{Rev}(xs, \text{Cons}(x, a)) \\
\\
\text{test1}(l) &= \text{Reverse1}(\text{Flatten}(l))
\end{aligned}$$

The result of applying the transformation rules to the function **test1** is shown below. To make the description less baroque we have left out the + and subscript

in the names of the combined functions.

$$\begin{aligned}
\text{test1}(l) &= \text{Reverse1Flatten}(l) \\
\text{Reverse1Flatten}(l) &= \text{RevFlatten}(l, \text{Nil}) \\
\text{RevFlatten}(l, r) &= \text{case } l \text{ of} \\
&\quad \text{Nil} \quad \quad \quad | r \\
&\quad \text{Cons}(x, xs) | \text{RevAppendFlatten}(x, xs, r) \\
\text{RevAppendFlatten}(xs, l, r) &= \text{case } xs \text{ of} \\
&\quad \text{Nil} \quad \quad \quad | \text{RevFlatten}(l, r) \\
&\quad \text{Cons}(x, xs) | \text{RevAppendFlatten}(xs, l, \text{Cons}(x, r))
\end{aligned}$$

The transformation steps that were needed to achieve part of this result are:

$$\begin{aligned}
&\mathcal{T}[\text{Reverse1}(\text{Flatten}(l))] \\
&= \mathcal{F}[\text{Reverse1}(\text{Flatten}(l))] \\
&= \text{Reverse1Flatten}(l) \\
&\text{Reverse1Flatten}(l) \\
&= \mathcal{T}[\text{Rev}(\text{Flatten}(l), \text{Nil})] \\
&= \text{RevFlatten}(l, \text{Nil}) \\
&\text{RevFlatten}(l, r) \\
&= \mathcal{C}[\text{case } (\text{case } l \text{ of} \\
&\quad \quad \quad \text{Nil} \quad \quad \quad | \text{Nil} \\
&\quad \quad \quad \text{Cons}(x, xs) | \text{Append}(x, \text{Flatten}(xs))) \\
&\quad \quad \quad \text{Nil} \quad \quad \quad | r \\
&\quad \quad \quad \text{Cons}(x, xs) | \text{Rev}(xs, \text{Cons}(x, r))] \\
&= \text{case } l \text{ of} \\
&\quad \text{Nil} \quad \quad \quad | r^4 \\
&\quad \text{Cons}(x, xs) | \mathcal{F}[\text{Rev}(\text{Append}(x, \text{Flatten}(xs)), r)] \\
&= \text{case } l \text{ of} \\
&\quad \text{Nil} \quad \quad \quad | r \\
&\quad \text{Cons}(x, xs) | \mathcal{F}[\text{RevAppend}(x, \text{Flatten}(xs), r)] \\
&= \text{case } l \text{ of} \\
&\quad \text{Nil} \quad \quad \quad | r \\
&\quad \text{Cons}(x, xs) | \mathcal{F}[\text{RevAppendFlatten}(x, xs, r)]
\end{aligned}$$

The generation of `RevAppendFlatten` out of `RevAppend` and `Flatten` proceeds in the same way.

⁴ Here we have applied the optimization step that was suggested in the footnote of definition 7

The second definition of `Reverse` uses the standard higher-order *foldl* function:

```
Reverse2(l) = Foldl(Snoc, Nil, l)

Foldl(f, r, l) = case l of
  Nil          | r
  Cons(x, xs)  | Foldl(f, (f @ r) @ x, xs)

Snoc(xs, x) = Cons(x, xs)

test2(l)     = Reverse2(Flatten(l))
```

Transforming `test2` result in two mutually recursive functions that are, except for the the order of the parameters, identical to the auxiliary functions `RevFlatten` and `RevAppendFlatten` generated by the transformation of `test1`.

Fusion appears to be much less successful if the following direct definition of `Reverse` is used:

```
Reverse3(l) = case l of
  Nil          | Nil
  Cons(x, xs)  | Append(Reverse3(xs), Cons(x, Nil))

test3(l)     = Reverse3(Flatten(l))
```

When transforming `test3` again `Reverse3` will be combined with `Flatten` but the fact that `Reverse3` itself is not a producer (the recursive occurrence of this function appears on a consuming position), will prevent the combined right-hand side from being deforested completely.

In general combinations of standard list functions (except for *accumulating* functions, such as `Reverse`), e.g.

```
Sum(Map Inc(Map Inc(Take(n, Repeat(1))))))
```

are transformed into a single function that generates no list at all and that does not contain any higher-order function applications. This shows that there is no longer a need to avoid higher-order functions like *foldl* or *map* for efficiency reasons.

5.2 Generics

Generic functions are translated into real functions in three steps. The first step, taking place before type checking, converts the generic specifications into class definitions and instances of those classes.⁵ Furthermore, conversion functions (so-called *bidirectional mappings*) are generated to convert the function arguments and result to and from the generic representation.

⁵ Currently the programmer has to specify for which types an instance is wanted. We could also generate instances lazily, i.e. when the compiler detects that a certain instance is required. For technical reasons this appears to be difficult.

The second translation step takes place during type checking: all overloading is resolved and replaced by appropriate dictionaries.

Finally, the overhead introduced by dictionaries, higher-order functions and by using the generic representations is removed by the fusion algorithm. For an in-depth description we refer to [Ali01]. Here, we briefly illustrate fusion with the aid of the example mentioned in the introduction. Suppose we have used the instance of equality for lists of integers, called `eqListInt`. This function can be directly expressed in terms of the library function `eqInt`, for comparing two integers, and the function `eqList`, for comparing lists. The latter is generated according to the generic specification of equality. The result is shown below:

```
eqListInt l1 l2 = eqList eqInt l1 l2
```

```
eqList eqa x y = fromBM (bimapEq isoList) (eqListG eqa) x y
```

The bidirectional mapping `bimapEq isoList` converts a function of type `[a] -> [a] -> Bool` to and from a function of type `GList a -> GList a -> Bool`, where `GList` denotes the generic list type. In this example we need the ‘from’ version selected by `fromBM`. The function `eqListG` is defined directly in terms of the user defined instances of equality for the generic type constructors. The structure of the list type itself determines how these user defined instances are combined.

```
eqListG eqa = eqEITHER eqUNIT (eqPAIR eqa (eqList eqa))
```

The complete program contains 10 to 20 auxiliary operations that are used to define `bimapEq` and `isoList`. There is no need to say that, without any further optimizations, operations defined in this way are very inefficient. It is pointless to show all fusion steps that were performed in order to obtain the optimized result: it appears that no less than 70 intermediate functions were generated. However, most of these become obsolete after the fusion process. After removing these superfluous function we end up with the following result where the names of the generated functions give some information about the functions that were combined.

```
eqListeqInt x y =
  eqEITHEReqUNITeqPAIReqListeqInteqIntmapListTomapListTo x y

eqEITHEReqUNITeqPAIReqListeqInteqIntmapListTomapListTo l1 l2 =
  case l1 of
    Nil
      = eqEITHERLEFTeqUNITUNITmapListTo l2
    Cons x xs
      = eqEITHERRIGHTeqPAIReqListeqInteqIntPAIRmapListTo x xs l2

eqEITHERLEFTeqUNITUNITmapListTo l =
  case l of
    Nil
      = True
    Cons x xs
```

```

    = False

eqEITHERRIGHTeqPAIReqListeqInteqIntPAIRmapListTo y ys l =
  case l of
    Nil
      = False
    Cons x xs
      = eqPAIR2eqListeqInteqIntPAIR y ys x xs

eqPAIR2eqListeqInteqIntPAIR y ys x xs =
  And (Equal x y) (eqListeqInt xs ys)

```

This result is essentially the same as any efficient hand written program.

6 Self-application

As a practical test of the effectiveness of the adjusted fusion algorithm we compared the memory allocation and speed of a Clean 2.0 compiler generated with fusion enabled with that of the same compiler generated without fusion. In both cases specialization of overloaded functions was enabled.

The fused compiler compiled its own code approximately 5 percent faster, and 9 percent less memory was allocated relative to the non-fused compiler. These improvements were almost entirely caused by better optimization of a few modules that use a monadic style, and not by removal of intermediate data structures using deforestation.

In one of these modules all curried function applications introduced by the monadic style were eliminated. Due to these optimizations the module executed 27 percent faster and allocated 52 percent less memory.

However in the most expensive module using a monadic style only 68 percent of the curried function applications were eliminated. This improved the execution speed 33 percent and the memory allocation 51 percent.

It should however be possible to remove nearly all these curried function applications. The current algorithm is not able to do this, because some functions that are called with a higher order function as argument are not optimized, because the argument is accumulating.

This happens also in the following example:

```

add :: Int Int -> Int
add a b = a+b

f :: [Int] Int -> Int
f [e:l] s = g (add e) l s
f [] s = s

g :: (Int -> Int) [Int] Int -> Int
g h l s = f l (h s)

```

The argument `h` of the function `g` is accumulating because `g` is called with `(add e)` as argument, therefore `g` is not fused with `(add e)`. In this case it would be safe to fuse.

This limitation prevents the compiler from removing nearly all the remaining curried function applications from the above mentioned module. However, if a call `f (g x)`, for some functions `f` and `g`, appears in a program, the argument of the function `f` does not always have to be treated as an accumulating argument. This is the case when the argument of `g` is always the same or does not grow in the recursion. By recognizing such cases we hope to optimize most of the remaining curried function applications. Or instead, we could fuse a limited number of times in these cases, to make the fusion algorithm terminate.

Another example of curried applications in this module that cannot be optimized are `foldl` calls that yield a higher order function. Such a higher order function occurs at an accumulating argument position in the `foldl` call, and can therefore not be fused.

7 Conclusion

The adjusted fusion algorithm combines deforestation together with dictionary elimination and higher-order removal. Practical results seem to indicate that the main benefits are achieved for specialised features such as type classes, generics, and monads rather than in ‘ordinary’ code. Further work remains to be done in the handling of accumulating parameters. Finally a formal proof of termination would be reassuring to have.

References

- [Ali01] A. Alimarine and R. Plasmeijer. *A Generic Programming Extension for Clean*. In: Arts, Th., Mohnen M., eds. Proceedings of the 13th International Workshop on the Implementation of Functional Languages, IFL 2001, Selected Papers, Älvsjö, Sweden, September 24–26, 2001, Springer-Verlag, LNCS 2312, pages 168–185.
- [Chin94] W.-N. Chin. *Safe fusion of functional expressions II: Further improvements* Journal of Functional Programming, Volume 6, Part 4, pp 515–557, 1994.
- [Cle13] M. Plasmeijer, M. van Eekelen. *Language Report Concurrent Clean. Version 1.3*. Technical Report CSI R9816, Faculty of mathematics and Informatics, Catholic University of Nijmegen, June 1998. Also available at www.cs.kun.nl/~clean/contents/contents.html
- [Cle20] M. Plasmeijer, M. van Eekelen. *Language Report Concurrent Clean. Version 2.0. DRAFT!* Faculty of mathematics and Informatics, Catholic University of Nijmegen, December 2001. Also available at www.cs.kun.nl/~clean/contents/contents.html
- [Fer88] A. Ferguson and P. Wadler. *When will Deforestation Stop*. In Proc. of 1988 Glasgow Workshop on Functional Programming, pp 39–56, Rothasay, Isle of Bute, August 1988.

- [Gil93] A. Gill, S. Peyton Jones, J. Launchbury. *A Short Cut to Deforestation*, Proc. Functional Programming Languages and Computer Architecture (FPCA'93), Copenhagen, June 1993, pp223–232.
- [Has92] P. Hudak, S. Peyton Jones, P. Wadler, B. Boutel, J. Fairbairn, J. Fasel, K. Hammond, J. Hughes, Th. Johnsson, D. Kieburtz, R. Nikhil, W. Partain, J. Peterson. *Report on the programming language Haskell*, In ACM SigPlan Notices, 27 (5): 1–164. May 1992.
- [Hin00] R. Hinze and S. Peyton Jones. *Derivable Type Classes*. In Graham Hutton, editor, Proceedings of the Fourth Haskell Workshop, Montreal, Canada, September 17, 2000
- [Oka98] C. Okasaki. *Purely Functional Data Structures* Cambridge University Press, ISBN 0-521-63124-6, 1998
- [Wad88] P. Wadler. *Deforestation: Transforming Programs to Eliminate Trees* Proceedings of the 2nd European Symposium on Programming, Nancy, France, March 1988. Lecture Notes in Computer Science 300.