

Fusion with Stacks and Accumulating Parameters ^{*}

Susumu Nishimura

Department of Mathematics, Faculty of Science, Kyoto University, Kyoto 606–8502, Japan
susumu@math.kyoto-u.ac.jp

ABSTRACT

We propose a new algorithm for fusion transformation that allows both stacks and accumulating parameters. The new algorithm can fuse programs that cannot be handled by existing fusion techniques, *e.g.*, XML transformations. The algorithm is formulated in a modular, type-directed style where the transformation process is comprised of several transformation steps that change types but preserve the observational behavior of programs. We identify a class of functions to which our new fusion method successfully applies and show that a closure property holds for that class.

Categories and Subject Descriptors

D.1.2 [Programming Techniques]: Automatic Programming—*program transformation*; D.1.1 [Programming Techniques]: Functional Programming; D.3.3 [Programming Languages]: Language Constructs and Features—*data types and structures*

General Terms

Algorithms, Languages

Keywords

Stacks, accumulating parameters, higher-order removal, shortcut fusion, type-directed program transformation.

1. INTRODUCTION

Modular programming, where a program is written as a collection of smaller parts that solve subproblems, makes programs significantly easier to understand and implement. However, finer modularity often goes against efficiency, because of the overhead of interfacing between subprograms.

^{*}Part of preliminary result of this paper was informally presented in Asian Workshop on Programming Languages and Systems 2002 (APLAS'02).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PEPM'04, August 24–26, 2004, Verona, Italy.
Copyright 2004 ACM 1-58113-835-0/04/0008 ...\$5.00.

Functional programs, in which interfacing is explicit as function composition, are well-suited for studying this issue, and the research within the last two decades has developed a family of novel optimization techniques called *fusion* (or *deforestation*) [27, 4, 8, 20, 14, 22, 11]. Fusion is a program transformation which unifies two recursive functions into one. A remarkable improvement on efficiency is achieved when the two functions are interfaced by a large sized intermediate recursively defined data structures, such as lists.

In this paper, we propose a new fusion algorithm that can deal with *stacks* and *accumulating parameters*, which are difficult or even impossible to handle in the existing fusion methods. The new algorithm is intended to encompass more realistic applications. For example, consider a program for XML document transformation. Such a program would be implemented in the following modular style.

$$\text{XML} \xrightarrow{\text{parse}} \text{XML tree} \xrightarrow{f} \text{XML tree} \xrightarrow{\text{unparse}} \text{XML}$$

The input XML text is parsed by the function *parse* to produce an internal tree representation. Then, a function *f* for document transformation is applied to produce another intermediate tree representation for the output. The resulting tree is translated back into the text representation by *unparse* for the final output. In this program, stacks and accumulating parameters naturally arise. Stacks are needed in order to parse against non-regular languages [10]. Accumulating parameters are also required for propagating the context information during parsing, unparsing, and any other computation of a reasonable complexity. (The discussion above holds for many software systems that operate on non-regular languages, *e.g.*, compilers.)

The technical contributions of this paper are as follows.

- Our algorithm is formulated in a highly modular manner in the sense that the algorithm is comprised of a series of *type-directed* transformation steps, where each step represents a *type conversion* on the program. This type-directed formulation gives a logically clean separation of the transformation process into smaller pieces that are easier to handle. Higher-order functions are intensively used for interfacing between transformation steps. This differentiates our work from existing program composition techniques. (See the discussion in Section 6.)
- Our fusion algorithm always succeeds if the source program meets a *linearity condition*. It is also shown that a *closure property* holds, *i.e.*, if the source program meets a (stronger) linearity condition, the transforma-

tion result is also linear. This closure property guarantees that a program comprising of more than two functions can be fused by repeated applications of our algorithm.

In this paper, we discuss a fusion algorithm for two functions, each of which recurses over a single input data structure. (We do not consider zip-like functions that recurse over two or more inputs simultaneously.) We assume a call-by-name semantics for a functional language we consider. When we discuss fusion for a function composition $g \circ f$, we call f the *producer function* and g the *consumer function*, where the latter consumes the intermediate result produced by the former. Each function may have accumulating parameters, while only the producer function can have stacks in its accumulating parameters and/or return results for managing data elements in the LIFO order. Parts of final output may be held in accumulating parameters and stacks, whereas no inspection into these data elements (*e.g.*, case distinction) is allowed.

The rest of this paper is organized as follows. Section 2 gives our running example for which we demonstrate our fusion algorithm. In Section 3, we explain our fusion transformation in a simplified setting. Section 4 presents the full transformation algorithm, which extends the aforementioned algorithm with stacks and multiple accumulating parameters and return results. In Section 5, we provide a condition for successful transformation and show a closure property. Section 6 mentions related work and finally Section 7 concludes with a brief summary and some issues for future work.

For the lack of space, we omit proofs and some program codes from this paper. The omitted materials are available on-line from the author's homepage <http://www.math.kyoto-u.ac.jp/~susumu/>.

2. TRANSFORMATION EXAMPLE

As an example of fusion transformation, let us consider a program that operates on sequences of matching parentheses. We chose this simple example for explanatory purposes only. We can consider a more involved example, say, XML transformation systems [16, 18].

Throughout the paper, we present programs by borrowing the syntax from Haskell [9] and extending it with records and overloaded functions. (See Section 4 for the definition of overloaded functions.) We write $M[x_1 \setminus N_1, \dots, x_n \setminus N_n]$ to represent a *substitution* to expression M , where the free occurrences of variables x_1, \dots, x_n in M are simultaneously replaced with N_1, \dots, N_n , respectively. We also use a special symbol Ω to represent an arbitrary well-typed expression. Every occurrence of Ω represents an unreachable point in the program and hence can be replaced with any (appropriately typed) expression without affecting the program's behavior. In practice, we can encode $\Omega = \text{error "unreachable"}$.

We model the sequence of (possibly unmatching) parentheses by the following datatype

```
data L = Begin L | End L | Nil
```

where the constructor **Begin** stands for the left parenthesis, **End** for the right parenthesis, and **Nil** for the end of sequence.

A sequence of matching parentheses is internally represented by a binary tree as defined below.

```
parse :: L -> (se :: Stk T) -> (l :: T, sp :: Stk T)
parse (Begin t) x0 =
  let x1 = parse t (se=Stk Leaf (#se x0))
      Stk a as = #sp x1
  in (l=Node (#l x1) a, sp=as)

parse (End t) x0 =
  let x1 = parse t (se=as)
      Stk a as = #se x0
  in (l=a, sp=Stk (#l x1) (#sp x1))

parse Nil x0 = (l=Leaf, sp=empty)

cascade :: T -> (h :: T) -> (l :: T)
cascade (Node t1 t2) x0 =
  let x1 = cascade t1 (h=Node (#l x2) Leaf)
      x2 = cascade t2 (h=#h x0)
  in (l=#l x1)

cascade Leaf x0 = (l=#h x0)


unparse :: T -> (k :: L) -> (j :: L)
unparse (Node t1 t2) u0 =
  let u1 = unparse t1 (k=End (#j u2))
      u2 = unparse t2 (k=#k u0)
  in (j=Begin (#j u1))

unparse Leaf u0 = (j=#k u0)
```

Figure 1: Source program

```
data T = Node T T | Leaf
```

Each binary node represents a matching pair of parentheses, where its left subtree corresponds to the subsequence embraced by that pair, while its right subtree corresponds to the rest of the sequence to follow. For example, a matching sequence $()(())$ is translated into a binary tree depicted below.



We consider a fusion for a system that is composed of three functions given in Fig. 1, where **parse** is the parsing function, **cascade** is the transformation function on the parsed binary tree, and **unparse** is the unpadding function.

The parsing function utilizes stacks, which we implement as follows.

```
data Stk a = Stk a (Stk a)
empty = error "empty stack"
```

We represent a stack with n elements a_1, \dots, a_n ($n \geq 0$) by a list-like structure terminated by **empty**, *i.e.*, $\text{Stk } a_1 (\text{Stk } a_2 (\dots (\text{Stk } a_n \text{ empty}) \dots))$. In Fig. 1, we use pattern matching $\text{Stk } a \text{ } as = S$ to bind the topmost element of a stack S to the variable a and the rest to the variable as . This corresponds to the “pop” operation on stacks.

The functions in Fig. 1 are all defined as a recursive function that takes a set of accumulating parameters and returns a set of computation results, where each set is expressed by a *record*.¹ A record, written $(l_1 = M_1, \dots, l_n = M_n)$

¹The record syntax is available as **Trex** extension in Haskell interpreter Hugs (<http://www.haskell.org/hugs>).

($n \geq 1$), is a tuple whose each element M_i (often called a *record field*) is indexed by a unique *record label* l_i . We write $(l_1 :: \tau_1, \dots, l_n :: \tau_n)$ for the type of records consisting of n fields where each record field labeled by l_i has type τ_i . (The order of labels is insignificant both in expressions and types.) To access a record field labeled by l in a record R , we use *field selection operator*, written $\#l R$.

We use record labels as a convenient means for marking certain set of expressions by syntactical objects. In the final transformation step of our algorithm (Section 3.4 and 4.6), it is crucial to identify some expressions according to their origin. Record labels are convenient for this, as the names of labels last unchanged throughout the transformation process. (cf. the names of program variables, which are subject to change due to α -conversion.) Ordinary (unlabeled) tuples may be used instead, but records can coexist more smoothly with our transformation process, which includes flattening and concatenation operations on tuples.

The explanation of the functions given in Fig. 1 follows below. The parsing function `parse` takes one accumulating parameter labeled by `se` and returns two results labeled by `sp` and `l`, where the `se` field holds a stack that remembers the number of parentheses yet to match and the `sp` field represents the parsing stack. The result of parsing is returned in the field `l`. When a malformed input sequence is fed to the `parse` function with `empty` being the initial value of the `se` field, an error will be signaled due to a “pop” operation on an empty stack.

The `cascade` function takes one accumulating parameter labeled by `h` and returns one result labeled by `l`. Given `Leaf` as the initial accumulating parameter, `cascade` recursively replaces the leftmost leaf node in every binary tree with the right subtree of the root node of that tree. This results in a binary tree where all the input binary nodes are reordered onto the spine of left branches in the output. For instance, a binary tree corresponding to the sequence $()((()))$ is transformed into the one corresponding to $((()))$.

The `unparse` function translates back the binary tree representation into the sequence representation. It takes one accumulating parameter labeled by `k`, whose initial value is `Nil`, and returns one result labeled by `j`.

In the following sections, we will derive a program that fuses these three functions. The result of fusion is given in Fig. 5 of Section 4. The resulting program does not produce any intermediate binary tree representation; it is defined as a first-order function that recurses over the input sequence with a new set of accumulating parameters and return results, which may circularly refers to each other.

3. FUSION WITH ACCUMULATING PARAMETERS

In this section we demonstrate our transformation algorithm by applying it for fusing the producer function `cascade` and the consumer function `unparse`.

Fig. 2 illustrates the outline of the transformation process.² Our transformation algorithm is *type-directed*, in a sense that each transformation step corresponds to a type conversion. (In Fig. 2, types that are subject to conversion

²For the sake of readability, the figure shows the case where each function takes one accumulating parameter and returns one result, and stacks are only in the accumulating parameter of the producer function.

are underlined.) It first applies shortcut fusion [8] for unifying the producer function `prod` and the consumer function `cons`. The result is a higher-order function `prod'`, which is turned into a first-order function in the remainder of transformation steps. This section demonstrates a simplified version of our transformation algorithm and does not deal with transformation steps (II) and (IV), which concern stack types. We will discuss the general case, allowing stacks and multiple accumulating parameters and return results, in the following section.

We will argue the correctness of each transformation step up to observational equivalence. Let us write $M_1 = M_2$ to mean observational equivalence, *i.e.*, M_1 and M_2 show the same behavior under any evaluation context. Under a call-by-name evaluation semantics, $\beta\eta$ -equality induces an observational congruence.³ As for records, the following relation holds.

$$\#l (l = M, \dots) = M \quad (\text{REC-CANCEL})$$

We will freely use $\beta\eta$ -equality and also (REC-CANCEL) for simplifying programs.

In addition, assuming M is a non-diverging record expression, we will use the following to reason about the correctness of transformation. (In this paper, the type annotation syntax $(M :: \tau)$ represents a typing constraint that M has type τ .)

$$(M :: (l_1 :: \tau_1, \dots, l_n :: \tau_n)) = (l_1 = \#l_1 M, \dots, l_n = \#l_n M) \quad (\text{REC-ETA})$$

This relation can be safely applied for reasoning about correctness anywhere in this paper, since we assume a particular syntactic form (discussed in Section 5), in which every record is created by the explicit record syntax $(l_1 = M_1, \dots, l_n = M_n)$ that never diverge.

In this paper, we carry out the discussion on the correctness in a rather conventional way, in order not to include too much theoretical details. (The argument on the correctness in this paper should be formalized in the framework of Pitts’ syntactic logical relations [19].) Suppose that we have a program transformation \mathcal{P} , whose corresponding type conversion is \mathcal{T} . To verify the correctness of transformation, we check if $\mathcal{P}[[M_1]] = \mathcal{P}[[M_2]]$ for every observationally equivalent pair of expressions M_1 and M_2 , and the interesting proof case is that for expressions of type τ such that the outermost type constructors in τ and $\mathcal{T}[[\tau]]$ differ. In that case, we show that, for any constructor C and destructor D_i for the type τ , observational equivalence $D_i (C M_1 \dots M_n) = M_i$ is preserved by the transformation, *i.e.*, $\mathcal{P}[[D_i (C M_1 \dots M_n)]] = \mathcal{P}[[M_i]]$ holds. As for a record type, this corresponds to verifying that (REC-CANCEL) is preserved by the transformation.

3.1 Shortcut fusion

We apply shortcut fusion to fuse the following function application.

```
unpcas :: T -> L
unpcas x =
  #j (unparse (#l (cascade x (h=Leaf))) (k=Nil))
```

³We do not consider Haskell’s *seq* operator, which breaks this property.




Figure 2: Transformation steps

For this, we rewrite `cascade` and `unparse` in terms of `build` and `fold`, whose definitions on the datatype `T` of binary trees are given below.

```

build :: (∀b. (b → b → b) → b → b) → T
build g = g Node Leaf

```

```

fold :: (a → a → a) → a → T → a
fold n l (Node t1 t2) = n (fold n l t1) (fold n l t2)
fold n l Leaf = l

```

The producer function is expressed as `build g`, where `g` is a function that abstracts over all occurrences of constructors for intermediate data (namely, `Node` and `Leaf`) and `build` reinstatiates the abstractions with concrete constructors. The consumer function is expressed as `fold c z`, using the generic recursion operator `fold` that replaces every `Node` constructor with `c` and `Leaf` with `z`.

The composition of the producer and consumer functions are fused by the following `fold/build` rule [8].

$$\text{fold } c \ z \ (\text{build } g) = g \ c \ z$$

Intuitively, this equation means that, instead of applying `fold c z` to the intermediate result produced by `build g`, we can obtain the same effect by instantiating abstracted constructors in `g` with `c` and `z`.

The correctness of shortcut fusion follows from free theorems [26, 12], the parametericity property induced from the polymorphic typing constraint on `build`.

The shortcut fusion applies to the present example as follows. (For general methods for translating normal recursive programs into `fold` and `build` forms, see [14, 5].) The consumer function `unparse` is defined with `fold` as follows.

```

unparse :: T → (k :: L) → (j :: L)
unparse = fold c z
  where c t1 t2 =
        \u0 -> let u1 = t1 (k= End (#j u2))
                u2 = t2 (k= #k u0)
                in (j= Begin (#j u1))
        z = \u0 -> (j= #k u0)

```

The producer function, as in the form used in the defini-

tion of the `unpcas` function above, is expressed as follows:

```

#l (cascade x (h=Leaf))
= #l (build x
      (\n l->cascade' t n l (h=fold n l Leaf)))

```

where `cascade'` is defined as follows.

```

cascade' n l (Node t1 t2) = \x0 ->
  let x1 = cascade' n l t1 (h= n (#l x2, l))
      x2 = cascade' n l t2 (h= #h x0)
  in (l= #l x1)
cascade' n l Leaf = \x0 -> (l= #h x0)

```

We note that we need to abstract over constructors not only in the function body of `cascade` but also its accumulating parameters by applying `fold n l`, so that `cascade'` meets the above mentioned polymorphic typing constraint [7].

Application of the `fold/build` rule, followed by some simplifications, gives the following result.

```

cascade' :: T → (h :: (k :: L) → (j :: L))
           → (l :: (k :: L) → (j :: L))
cascade' (Node t1 t2) = \x0 ->
  let x1 = cascade' t1
      (h= \u -> (j=Begin (#j
                        ((#l x2) (k=End (#k u))))))
      x2 = cascade' t2 (h= #h x0)
  in (l= #l x1)
cascade' Leaf = \x0 -> (l= #h x0)

```

```

unpcas :: T → L
unpcas x =
  #j ((#l (cascade' x (h= \x0 -> (j= #k x0))))
      (k=Nil))

```

3.2 Function type lifting

The transformation step that follows shortcut fusion is called *function type lifting*, which corresponds to the following type conversion

$$\begin{array}{ll}
\text{from} & (p :: \tau_1 \rightarrow \tau_2) \\
\text{into} & (p :: \tau_1) \rightarrow (p :: \tau_2),
\end{array}$$

where `p` is either `h` or `l`.

This type conversion is realized by applying a transformation function called \mathcal{L} to every defining equation in the program. \mathcal{L} traverses the program structure and applies the following transformation rules wherever possible.

$$\begin{aligned}\mathcal{L}[(p = (M :: \tau_1 \rightarrow \tau_2))] &= \backslash y \rightarrow (p = \mathcal{L}[[M]] (\#p y)) \\ \mathcal{L}[\#p (M :: (p :: \tau_1 \rightarrow \tau_2))] &= \backslash y \rightarrow \#p (\mathcal{L}[[M]] (p = y))\end{aligned}$$

(In both rules, y denotes a fresh variable.)

We omitted trivial transformation rules that are needed just for traversing the program structure. The same convention applies to the rest of the paper.

According to the discussion in the beginning of this section, we can verify the correctness of the present transformation by checking (REC-CANCEL) is preserved by the transformation.

$$\begin{aligned}\mathcal{L}[\#p (p = M)] \\ = \backslash y \rightarrow \#p ((\backslash y \rightarrow (p = \mathcal{L}[[M]] (\#p y))) (p = y)) \\ = \backslash y \rightarrow \mathcal{L}[[M]] y = \mathcal{L}[[M]] \quad \text{by } \eta\text{-equality}\end{aligned}$$

The result of function type lifting transformation is given as follows.

$$\begin{aligned}\text{cascade}' &:: T \rightarrow (((h :: (k :: L)) \rightarrow (h :: (j :: L))) \\ &\quad \rightarrow ((l :: (k :: L)) \rightarrow (l :: (j :: L)))) \\ \text{cascade}' (\text{Node } t1 \ t2) &= \backslash x0 \rightarrow \\ &\quad \text{let } x1 = \text{cascade}' \ t1 \\ &\quad \quad (\backslash y1 \rightarrow (h=(j=\text{Begin } (\#j (\#1 \ x2 \\ &\quad \quad \quad (l=(k=\text{End } (\#k (\#h \ y1)))))))))) \\ &\quad \quad x2 = \text{cascade}' \ t2 \\ &\quad \quad \quad (\backslash y2 \rightarrow (h= \#h \ (x0 \ (h= \#h \ y2)))) \\ &\quad \text{in } \backslash y0 \rightarrow (l= \#1 \ (x1 \ (l= \#1 \ y0))) \\ \text{cascade}' \text{ Leaf} &= \backslash x0 \rightarrow \\ &\quad \backslash y0 \rightarrow (l= \#h \ (x0 \ (h= \#1 \ y0))) \\ \text{unpcas} &:: T \rightarrow L \\ \text{unpcas } x &= \\ &\quad \#j (\#1 (\text{cascade}' \ x \ (\backslash y \rightarrow (h=(j= \#k \ (\#h \ y)))) \\ &\quad \quad (l=(k=\text{Nil}))))\end{aligned}$$

3.3 Record type flattening

This transformation step intends a type conversion

$$\begin{aligned}\text{from} & \quad (p :: (q :: L)) \\ \text{into} & \quad (pq :: L),\end{aligned}$$

where p is either h or l , q is either k or j , and pq is a new label obtained by concatenating the two.

This type conversion represents a transformation that merges any pair of nested record labels p and q into a single one pq , i.e., $(p = (q = M))$ into $(pq = M)$ and $\#q (\#p M)$ into $\#pq M$. Formal transformation rules are given below.

$$\begin{aligned}\mathcal{F}[(p = (M :: (q :: L)))] &= (pq = \#q \mathcal{F}[[M]]) \\ \mathcal{F}[\#p (M :: (p :: (q :: L)))] &= (q = \#pq \mathcal{F}[[M]])\end{aligned}$$

We can check this transformation preserves (REC-CANCEL).

$$\begin{aligned}\mathcal{F}[\#p (p = (M :: (q :: L)))] \\ = (q = \#pq (pq = \#q \mathcal{F}[[M]]) \\ = (q = \#q \mathcal{F}[[M]]) = \mathcal{F}[[M]] \quad \text{by (REC-ETA)}\end{aligned}$$

The result of record type flattening follows below.

$$\begin{aligned}\text{cascade}' &:: T \rightarrow (((hk :: L) \rightarrow (hj :: L)) \\ &\quad \rightarrow ((lk :: L) \rightarrow (lj :: L))) \\ \text{cascade}' (\text{Node } t1 \ t2) &= \backslash x0 \rightarrow \\ &\quad \text{let } x1 = \text{cascade}' \ t1 \\ &\quad \quad (\backslash y1 \rightarrow (hj=\text{Begin } (\#lj \ (x2 \\ &\quad \quad \quad (lk=\text{End } (\#hk \ y1)))))) \\ &\quad \quad x2 = \text{cascade}' \ t2 \\ &\quad \quad \quad (\backslash y2 \rightarrow (hj= \#hj \ (x0 \ (hk= \#hk \ y2)))) \\ &\quad \text{in } \backslash y0 \rightarrow (lj= \#lj \ (x1 \ (lk= \#lk \ y0))) \\ \text{cascade}' \text{ Leaf} &= \backslash x0 \rightarrow \\ &\quad \backslash y0 \rightarrow (lj= \#hj \ (x0 \ (hk= \#lk \ y0))) \\ \text{unpcas} &:: T \rightarrow L \\ \text{unpcas } x &= \\ &\quad \#lj (\text{cascade}' \ x \ (\backslash y \rightarrow (hj= \#hk \ y)) \ (lk=\text{Nil}))\end{aligned}$$

3.4 Higher-order removal

Now we are ready for removing higher-order functions. At the type level, the higher-order removal transformation conducts a type conversion

$$\begin{aligned}\text{from} & \quad ((hk :: L) \rightarrow (hj :: L)) \rightarrow ((lk :: L) \rightarrow (lj :: L)) \\ \text{into} & \quad (hj :: L, lk :: L) \rightarrow (hk :: L, lj :: L).\end{aligned}$$

We remark that this type conversion preserves polarity of record fields. That is, record fields that occur positively in the higher-order function type (namely, hk and lj fields) also appear in positive positions of the resulting first-order function type; so do record fields in negative positions (namely, hj and lk fields).

In a previous work [17], the author has proposed a transformation algorithm that applies to a class of higher-order functions whose each defining equation has the following specific syntactic form

$$\begin{aligned}f &:: T \rightarrow ((hk::L) \rightarrow (hj::L)) \rightarrow ((lk::L) \rightarrow (lj::L)) \\ f \ (C \ t_1 \ \dots \ t_c) &= \backslash x_0 \rightarrow \\ &\quad \text{let } x_1 = f \ t_{j(1)} \ (\backslash y_1 \rightarrow (hj=N_1)) \\ &\quad \quad \dots \\ &\quad \quad x_r = f \ t_{j(r)} \ (\backslash y_r \rightarrow (hj=N_r)) \\ &\quad \text{in } \backslash y_0 \rightarrow (lj=N_0)\end{aligned} \tag{3.1}$$

where $r, c \geq 0$, $1 \leq j(i) \leq c$ for every i ($1 \leq i \leq r$), N_i 's belong to a set of first-order expressions defined by the following grammar

$$\begin{aligned}\mathcal{M} &::= \#hj \ (x_0 \ (hk = \mathcal{M})) \mid \#lj \ (x_i \ (lk = \mathcal{M})) \\ &\quad \mid \#lk \ y_0 \mid \#hk \ y_i \mid C \ \mathcal{M} \ \dots \ \mathcal{M} \quad (1 \leq i \leq r),\end{aligned}$$

and furthermore the uses of variables x_0, \dots, x_r are *linear*. It follows from the linearity that, for every i and appropriate record labels p and q , there is one and only one expression M such that M occurs in the form $\#p (x_i (q = M))$. Let us write M_{x_0} for the unique expression M in the form $\#hj (x_0 (hk = M))$ and M_{x_i} for the unique expression M in the form $\#lj (x_i (lk = M))$ ($1 \leq i \leq r$).

Each defining equation of the intermediate transformation result in the previous subsection has the above syntactic form. This holds for any program that is derived from a source program that satisfies certain syntactic condition (involving linearity on the uses of variables), which we will discuss in Section 5.

In what follows, let us explain our higher-order removal transformation method in a rather informal way, since a formal justification requires an intricate discussion on the observational equivalence of programs, which is out of the scope of the present paper. Interested readers are deferred to [17].

In the light of the intended conversion on type, we can guess that the transformation would result in a set of defining equations of the following form.

$$\begin{aligned}
f' &:: T \rightarrow (hj :: L, lk :: L) \rightarrow (hk :: L, lj :: L) \\
f' (C t_1 \dots t_c) x_0 &= \\
\text{let } x_1 &= f' t_{j(1)} (hj = M_1^{hj}, lk = M_1^{lk}) \\
&\dots \\
x_r &= f' t_{j(r)} (hj = M_r^{hj}, lk = M_r^{lk}) \\
\text{in } (lj &= M_0^{lj}, hk = M_0^{hk})
\end{aligned}$$

We need to find appropriate expressions to fill in M_0^{lj} , M_0^{hk} , M_i^{hj} 's, and M_i^{lk} 's in the above program. First, we can observe that expression N_i ($1 \leq i \leq r$) in (3.1) is each appropriate for M_i^{hj} , as it represents the one and only one hj value computed by the recursive call $f t_{j(i)} \dots$. Similarly, N_0 is appropriate for M_0^{lj} . However, we notice that any occurrence of y_i 's in these expressions becomes void in the target program, since there are no binding constructs for them. To fix this ill-definedness, we rewrite every subexpression $\#lk y_0$ ($\#hk y_i$, *resp.*) into $\#lk x_0$ ($\#hk x_i$, *resp.*), expecting that bindings for x_0, \dots, x_r in the target program would give the appropriate lk (hk , *resp.*) value. We note that this rewrite gives rise to a circular definition, where some record fields returned by a function application are referenced from the argument positions of that function call itself. (This *circular tupling technique* has been demonstrated by Bird [1] and has also been exploited in [24] to deforest in accumulating parameters.)

For the above circular binding to be valid, each $\#hk y_i$ ($1 \leq i \leq r$) in the source program (function f) should denote the same value as $\#hk x_i$ does in the target program (function f'). Since the function call $f t_{j(i)} (\backslash y_i \rightarrow \dots)$ in f binds the variable x_0 to the function $\backslash y_i \rightarrow \dots$ and uses it in the form $\#hj (x_0 (hk = M_{x_0}))$, y_i is bound to $(hk = M_{x_0})$ eventually. On the other hand, in f' , x_i is bound to the result of function call $f' t_{j(i)} \dots$, which returns the record $(hk = M_0^{hk}, \dots)$. Hence, M_0^{hk} must be M_{x_0} .

We also require that $\#lk y_0$ in f denotes the same value as $\#lk x_0$ in f' . (The initial call sites of both functions are also assumed to meet this condition, as indicated by Eq. (3.3) that we will see later in this section.) It follows from this requirement that, by a similar discussion as above, M_i^{lk} must be M_{x_i} for each i ($1 \leq i \leq r$).

To complete the transformation, we rewrite $\#hj (x_0 (hk = M))$ into $\#hj x_0$, where the hj value returned by an application of x_0 in the source program is given in the target program through the hj field in the record bound by x_0 . Similarly, we rewrite $\#lj (x_i (lk = M))$ into $\#lj x_i$ for every i ($1 \leq i \leq r$).

By the discussion so far, we transform the function definition in the form (3.1) into the following first-order program

$$\begin{aligned}
f' &:: T \rightarrow (hj :: L, lk :: L) \rightarrow (hk :: L, lj :: L) \\
f' (C t_1 \dots t_c) x_0 &= \\
\text{let } x_1 &= f' t_{j(1)} (hj = \mathcal{H}[\llbracket N_1 \rrbracket], lk = \mathcal{H}[\llbracket M_{x_1} \rrbracket]) \\
&\dots \\
x_r &= f' t_{j(r)} (hj = \mathcal{H}[\llbracket N_r \rrbracket], lk = \mathcal{H}[\llbracket M_{x_r} \rrbracket]) \\
\text{in } (lj &= \mathcal{H}[\llbracket N_0 \rrbracket], hk = \mathcal{H}[\llbracket M_{x_0} \rrbracket])
\end{aligned}$$

```

cascade' :: T → (hj::L, lk::L) → (hk::L, lj::L)
cascade' (Node t1 t2) x0 =
  let x1 = cascade' t1
      (lk= #lk x0, hj= Begin (#lj x2))
      x2 = cascade' t2
      (lk= End (#hk x1), hj= #hj x0)
  in (hk= #hk x2, lj= #lj x1)

cascade' Leaf x0 = (hk= #lk x0, lj= #hj x0)

```

```

unpcas :: T → L
unpcas x = let r = cascade' x (hj= #hk r, lk=Nil)
           in #lj r

```

Figure 3: Fusion result on unpcas

where \mathcal{H} is the transformation function defined by the following set of rules.

$$\mathcal{H}[\llbracket x_i M \rrbracket] = x_i \quad (0 \leq i \leq r) \quad (3.2a)$$

$$\mathcal{H}[\llbracket y_i \rrbracket] = x_i \quad (0 \leq i \leq r) \quad (3.2b)$$

The following equation, taken from [17], characterizes the relation between the higher-order function f and the resulting first-order function f' :

$$\begin{aligned}
&\#lj (f L G (lk = L')) \\
&= \#lj (f' L (hj = \#hj (G \\
&\quad (hk = \#hk (f' L (hj = \Omega, lk = L'))), \\
&\quad lk = L'))),
\end{aligned}$$

where L is an expression of type T , L' is of type L , and G is a function of type $(hk :: L) \rightarrow (hj :: L)$.

The two different calls to f' in the RHS of the equation can be unified into a single call with a circular definition, as follows.

$$\begin{aligned}
&\#lj (f L G (lk = L')) \\
&= \text{let } r = f' L (hj = \#hj (G (hk = \#hk r)), lk = L') \\
&\text{in } \#lj r \quad (3.3)
\end{aligned}$$

It is easy to see the former equation follows from the latter, by unwinding the circular let and exchanging Ω with an appropriate expression. (Remember that Ω is a special symbol that represents arbitrary expression.)

Applying the higher-order removal transformation to `cascade'` and rewriting the definition of `unpcas` according to the equation (3.3), we obtain the fusion result given in Fig. 3.

4. FUSION WITH STACKS

In this section, we carry out the remaining half of transformation. That is, we apply our fusion algorithm to the following program

```

unpcaspar :: L → L
unpcaspar x = unpcas (#l (parse x (se=empty)))

```

which expands to

```

unpcaspar x =
  let r = cascade' (#l (parse x (se=empty)))
      (hj= #hk r, lk=Nil)
  in #lj r

```

This section applies our full transformation algorithm, which refines and augments the one given in the previous section in two points. First, two new transformation steps concerning stack types ((II) and (IV) in Fig. 2) are added. Second, the transformation rules are generalized so that they allow multiple accumulating parameters and return values.

For this generalization, we extend our language with *first-class overloaded functions*, which are expressed by the following syntax.

$$\{\backslash(x_1 :: \tau_1) \rightarrow M_1; \dots; \backslash(x_n :: \tau_n) \rightarrow M_n\} \quad (n \geq 1)$$

An overloaded function is a collection of functions, from which one that has a matching argument type is selected and evaluated per each application of that overloaded function. (A normal λ -abstraction can be recognized as a special instance of an overloaded function.) The type of an overloaded function is written as a collection of function types

$$(\tau_1 \rightarrow \tau'_1; \dots; \tau_n \rightarrow \tau'_n) \quad (\text{or } \sum_{i=1}^n \tau_i \rightarrow \tau'_i \text{ for short}),$$

where τ'_1, \dots, τ'_n can be different types. In this paper, every different function that is a member of an overloaded function has a different argument type ($p :: \tau$) distinguished by the record label p . Therefore no ambiguity arises in the selection of function and the return type of each application of an overloaded function is statically determined.

In the rest of this section, let us write \vec{p} for a non-empty set of record labels $\{p_1, \dots, p_n\}$ ($n \geq 1$). We also write $(\vec{p} :: \tau)$ to abbreviate $(p_1 :: \tau, \dots, p_n :: \tau)$ and also write $\vec{\tau}$ to denote either τ or $\text{Stk } \tau$.

We write σ_1, σ_2 , and σ_3 for the types of the input, intermediate, and output data structures, respectively. We also write $\vec{h}, \vec{l}, \vec{k}, \vec{j}$ for disjoint sets of record labels, where \vec{h} are used for labeling accumulating parameters of the producer function and \vec{l} for return results; \vec{k} are used for labeling accumulating parameters of the consumer function and \vec{j} for return results.

In this section, we apply the transformation algorithm to the pair of the producer function `parse` and the consumer function `cascade`, where $\sigma_1 = L, \sigma_2 = T, \sigma_3 = L, \vec{h} = \{\text{se}\}, \vec{l} = \{1, \text{sp}\}, \vec{k} = \{\text{hj}, \text{lk}\}$, and $\vec{j} = \{\text{hk}, \text{lj}\}$. We omit some of the intermediate transformation results, in order not to clutter pages with lengthy codes.

4.1 Shortcut fusion

As we did in the previous section, we first apply shortcut fusion. The result is displayed below.

```

parse' ::
  L → (se :: Stk((hj :: L, lk :: L) → (hk :: L, lj :: L)))
    → (sp :: Stk((hj :: L, lk :: L) → (lj :: L, hk :: L)),
       l :: (hj :: L, lk :: L) → (hk :: L, lj :: L))
parse' (Begin t) = \x0 ->
  let x1 = parse' t
      (se=Stk (\u0 -> (hk= #lk u0, lj= #hj u0))
        (#se x0))
  in (l= \u0 ->
      let u1 = (#l x1) (lk= #lk u0,
                      hj=Begin (#lj u2))
          u2 = (hd (#sp x1))
              (lk=End (#hk u1), hj= #hj u0)
          in (hk= #hk u2, lj= #lj u1),
      sp=t1 (#sp x1))

```

```

parse' (End t) = \x0 ->
  let x1 = parse' t (se=t1 (#se x0))
  in (l=hd (#se x0), sp=Stk (#l x1) (#sp x1))

```

```

parse' Nil = \x0 ->
  (l= \u0 -> (hk= #lk u0, lj= #hj u0), sp=empty)

```

```

unpcaspar :: L → L
unpcaspar x =
  let r = #l (parse' x (se= empty))
      (hj= #hk r, lk=Nil)
  in #lj r

```

This is obtained by a similar derivation in Section 3.1. The only difference is that we must abstract over constructors in the accumulating parameter of `parse` by applying `mapS` (`fold n l`), where n and l are abstraction variables and `mapS` is a map function on stack elements defined by:

$$\text{mapS } f \text{ (Stk } a \text{ as)} = \text{Stk } (f \ a) \text{ (mapS } f \text{ as)}$$

For the sake of brevity, we have replaced every reference to pattern matching variables a and as bound by a local definition $\text{Stk } a \text{ as} = S$ with `hdS` and `t1S`, respectively, where `hd` and `t1` are functions satisfying the following equations.

$$\text{hd (Stk } A \ S) = A \quad (\text{STK-HD})$$

$$\text{t1 (Stk } A \ S) = S \quad (\text{STK-TL})$$

4.2 Function type lifting for stacks

We introduce a new transformation step that represents the following type conversion:

$$\begin{array}{ll} \text{from} & \text{Stk}(\tau_1 \rightarrow \tau_2) \\ \text{into} & \text{Stk } \tau_1 \rightarrow \text{Stk } \tau_2. \end{array}$$

This conversion makes function arrows exposed outside of the enclosing stack type constructors, so that the transformation result is ready for the next transformation step. This is realized by the following set of transformation rules

$$\begin{aligned} \mathcal{S}[\text{Stk } A \ M :: \text{Stk}(\tau_1 \rightarrow \tau_2)] &= \\ & \backslash w \rightarrow \text{Stk}(\mathcal{S}[A](\text{hd } w))(\mathcal{S}[M](\text{t1 } w)) \\ \mathcal{S}[\text{hd } M] &= \backslash u \rightarrow \text{hd } (\mathcal{S}[M] \text{ (Stk } u \ \Omega)) \\ \mathcal{S}[\text{t1 } M] &= \backslash w \rightarrow \text{t1 } (\mathcal{S}[M] \text{ (Stk } \Omega \ w)) \end{aligned}$$

where w and u are fresh variables.

We remark that the above embedding into the target type $\text{Stk } \tau_1 \rightarrow \text{Stk } \tau_2$ is not full: the resulting function always maps a stack into another stack of the same length. Furthermore, the function operates on stacks elementwisely. That is, each element in the output stack is determined only by the element at the same depth in the input stack; the other elements in the input stack are never referenced and are represented by Ω in the rules above.

By a similar discussion in Section 3, the correctness of this transformation is verified by checking if the transformation preserves (STK-HD) and (STK-TL). We show the former case only; the verification on the latter is similar.

$$\begin{aligned} & \mathcal{S}[\text{hd (Stk } A \ M)] \\ &= \backslash w \rightarrow \text{hd } ((\backslash w' \rightarrow \text{Stk } (\mathcal{S}[A](\text{hd } w')) (\mathcal{S}[M](\text{t1 } w')))) \\ & \quad (\text{Stk } w \ \Omega)) \\ &= \backslash w \rightarrow \mathcal{S}[A] \ w = \mathcal{S}[A] \quad \text{by (STK-HD) and } \eta\text{-equality} \end{aligned}$$

4.3 Function type lifting for records

This transformation step is a generalization of the one given in Section 3.2 and corresponds to the following type conversion:

$$\begin{array}{l} \text{from} \quad (p_1 :: \tau_1 \rightarrow \tau'_1, \dots, p_n :: \tau_n \rightarrow \tau'_n) \\ \text{into} \quad ((p_1 :: \tau_1) \rightarrow (p_1 :: \tau'_1); \dots; (p_n :: \tau_n) \rightarrow (p_n :: \tau'_n)) \end{array}$$

where $n \geq 1$ and $\vec{p} = \vec{h}$ or $\vec{p} = \vec{l}$.

This conversion is realized by the following set of transformation rules

$$\begin{aligned} \mathcal{L}[(p_1 = (M_1 :: \tau_1 \rightarrow \tau'_1), \dots, p_n = (M_n :: \tau_n \rightarrow \tau'_n))] &= \\ \{ \backslash (y :: (p_1 :: \tau_1)) \rightarrow (p_1 = \mathcal{L}[[M_1]](\#p_1 y)); & \\ \dots; & \\ \backslash (y :: (p_n :: \tau_n)) \rightarrow (p_n = \mathcal{L}[[M_n]](\#p_n y)) \} & \\ \mathcal{L}[\#p (M :: (p_1 :: \tau_1 \rightarrow \tau_1, \dots, p_n :: \tau_n \rightarrow \tau'_n))] &= \\ \backslash v \rightarrow \#p(\mathcal{L}[[M]](p = v)) & \end{aligned}$$

where y and v are fresh variables.

In this transformation step, the set of functions in a record $R = (p_1 = F_1, \dots, p_n = F_n)$ is translated into an overloaded function $F = \{F'_1; \dots; F'_n\}$, where each function F'_i has type $(p_i :: \tau_i) \rightarrow (p_i :: \tau'_i)$. An application of the function F_i in the form $(\#p_i R)M$ in the source program is interpreted in the target program as $\#p_i (F (p_i = M))$, where the corresponding function F'_i in the overloaded function is selected up to the record label p_i . A similar translation may be possible by using a function of type $\langle p_1 : \tau_1, \dots, p_n : \tau_n \rangle \rightarrow \langle p_1 : \tau'_1, \dots, p_n : \tau'_n \rangle$ that operates on disjoint sum types.⁴ However, this not only adds extra syntactic verbosity but also gives an imprecision in typing: it blurs the fact that the function maps every input with a record (or injection) label into an output with the same label.

The correctness of this transformation is verified by a derivation similar to that in Section 3.2.

4.4 Record type lifting

Here we need another new transformation step that operates on stack types:

$$\begin{array}{l} \text{from} \quad \text{Stk}(\vec{q} :: \tau) \\ \text{into} \quad (\vec{q} :: \text{Stk } \tau), \end{array}$$

where $\vec{q} = \vec{j}$ or $\vec{q} = \vec{k}$.

The following transformation rules realize such type conversion.

$$\begin{aligned} \mathcal{R}[(\text{Stk } A \ M :: \text{Stk}(\vec{q} :: \tau))] &= \\ (q_1 = \text{Stk}(\#q_1 \mathcal{R}[[A]])(\#q_1 \mathcal{R}[[M]]), \dots, & \\ q_m = \text{Stk}(\#q_m \mathcal{R}[[A]])(\#q_m \mathcal{R}[[M]]) & \\ \mathcal{R}[\text{hd} (M :: \text{Stk}(\vec{q} :: \tau))] &= \\ (q_1 = \text{hd}(\#q_1 \mathcal{R}[[M]]), \dots, q_m = \text{hd}(\#q_m \mathcal{R}[[M]]) & \\ \mathcal{R}[\text{tl} (M :: \text{Stk}(\vec{q} :: \tau))] &= \\ (q_1 = \text{tl}(\#q_1 \mathcal{R}[[M]]), \dots, q_m = \text{tl}(\#q_m \mathcal{R}[[M]]) & \end{aligned}$$

A similar remark in Section 4.2 applies to this transformation. That is, the embedding into the target type is not full

⁴ $\langle p_1 : \tau_1, \dots, p_n : \tau_n \rangle$ represents a disjoint sum type with n distinct injection labels into the sum type.

in the sense that every field in the resulting record holds a stack of the same length.

The following derivation verifies that the transformation preserves (STK-HD).

$$\begin{aligned} \mathcal{R}[\text{hd} (\text{Stk } A \ M :: \text{Stk}(\vec{q} :: \tau))] & \\ = (q_1 = \#q_1 \mathcal{R}[[A]], \dots, q_m = \#q_m \mathcal{R}[[A]]) & \text{ by (REC-CANCEL)} \\ = \mathcal{R}[[A]] & \text{ by (REC-ETA)} \end{aligned}$$

The derivation for (STK-TL) is similar.

4.5 Record type flattening

At this transformation step, we flatten nested record types. We generalize the transformation given in Section 3.3 so that it can handle records with multiple fields.

Let us write $p\vec{q}$ for the set of labels $\{pq \mid q \in \vec{q}\}$. The generalized transformation represents a type conversion

$$\begin{array}{l} \text{from} \quad (p :: (\vec{q} :: \vec{\tau})) \\ \text{into} \quad (p\vec{q} :: \vec{\tau}), \end{array}$$

where $p \in \vec{h} \cup \vec{l}$ and either $\vec{q} = \vec{k}$ or $\vec{q} = \vec{j}$.

The transformation rules are given below.

$$\begin{aligned} \mathcal{F}[(p = (M :: (\vec{q} :: \tau)))] &= \\ (pq_1 = \#q_1 \mathcal{F}[[M]], \dots, pq_m = \#q_m \mathcal{F}[[M]]) & \\ \mathcal{F}[\#p (M :: (p :: (\vec{q} :: \tau)))] &= \\ (q_1 = \#pq_1 \mathcal{F}[[M]], \dots, q_m = \#pq_m \mathcal{F}[[M]]) & \end{aligned}$$

The correctness of this transformation is verified by a derivation similar to that in Section 3.3.

The result of this transformation step is given in Fig. 4, where we saved space by writing **let** U **in** $(p_1=M_1, \dots, p_n=M_n)$ instead of $(p_1=\text{let } U \text{ in } M_1, \dots, p_n=\text{let } U \text{ in } M_n)$.

4.6 Higher-order removal

This section presents a higher-order removal transformation method that generalizes the one given in Section 3.4. In the following, let us write $p\vec{q}$ for the set of labels $\{pq \mid p \in \vec{p}, q \in \vec{q}\}$. We write $\prod_{p \in \vec{p}} (p = M_p)$ to abbreviate $(p_1 = M_{p_1}, \dots, p_n = M_{p_n})$ and also write $\prod_{p \in \vec{p}} (p = M_p) \otimes \prod_{q \in \vec{q}} (q = N_q)$ to denote the record expression $(p_1 = M_{p_1}, \dots, p_n = M_{p_n}, q_1 = N_{q_1}, \dots, q_m = N_{q_m})$.

The present transformation can be explained as a type conversion

$$\begin{array}{l} \text{from} \quad \sum_{i=1}^n (h_i \vec{k} :: \vec{\sigma}_3) \rightarrow (h_i \vec{j} :: \vec{\sigma}_3) \\ \rightarrow \sum_{i=1}^m (l_i \vec{k} :: \vec{\sigma}_3) \rightarrow (l_i \vec{j} :: \vec{\sigma}_3) \\ \text{into} \quad (\vec{h} \vec{j} :: \vec{\sigma}_3, \vec{l} \vec{k} :: \vec{\sigma}_3) \rightarrow (\vec{h} \vec{k} :: \vec{\sigma}_3, \vec{l} \vec{j} :: \vec{\sigma}_3). \end{array}$$

The transformation steps so far would result in a higher-order function f , whose each defining equation has the following form.

$$\begin{aligned} f :: \sigma_1 \rightarrow \sum_{i=1}^n (h_i \vec{k} :: \vec{\sigma}_3) \rightarrow (h_i \vec{j} :: \vec{\sigma}_3) & \\ \rightarrow \sum_{i=1}^m (l_i \vec{k} :: \vec{\sigma}_3) \rightarrow (l_i \vec{j} :: \vec{\sigma}_3) & \\ f (C \ t_1 \ \dots \ t_c) = \backslash x_0 \rightarrow \text{let } x_1 = f \ t_{j(1)} \ F_1 & \quad (4.1) \\ \dots & \\ x_r = f \ t_{j(r)} \ F_r & \\ \text{in } F_0 & \end{aligned}$$


```

parse' :: L → ((selk::Stk L, sehj::Stk L) → (sehk::Stk L, selj::Stk L))
          → ((lhj::L, llk::L) → (lhk::L, llj::L); (sphj::Stk L, splk::Stk L) → (splj::Stk L, sphk::Stk L))
parse' (Begin t) = \x0 ->
  let x1 = parse' t (\y1 -> (sehk=Stk (hd (#selk y1)) (#sehk (x0 (selk=t1 (#selk y1), sehj=t1 (#sehj y1)))),
                          selj=Stk (hd (#sehj y1), (#selj (x0 (selk=t1 (#selk y1), sehj=t1 (#sehj y1)))))))
  in {\(y0 :: (lhj::L, llk::L)) ->
    let u1 = (hk= #lhk (x1 (llk= #llk y0, lhj=Begin (#lj u2))), lj= #llj (x1 (llk= #llk y0, lhj=Begin (#lj u2))))
      u2 = (hk= hd (#sphk (x1 (splk=Stk (End (#hk u1)) Ω, sphj=Stk (#lhj y0) Ω))),
          lj= hd (#splj (x1 (splk=Stk (End (#hk u1)) Ω, sphj=Stk (#lhj y0) Ω))))
    in (lhk= #hk u2, llj= #lj u1) ;
    \y0 :: (sphj::Stk L, splk::Stk L) ->
      (sphk=t1 (#sphk (x1 (splk=Stk Ω (#splk y0), sphj=Stk Ω (#sphj y0)))),
       splj=t1 (#splj (x1 (splk=Stk Ω (#splk y0), sphj=Stk Ω (#sphj y0)))) ) }

parse' Nil = \x0 -> {\(y0 :: (lhj::L, llk::L)) -> (lhk= #lhk y0, llj= #lhj y0);
                    \y0 :: (sphj::Stk L, splk::Stk L) -> (sphk=empty, splj=empty) }

uncaspar :: L → L
uncaspar x = let r = parse' x (\w->(sehk=empty, selj=empty)) (lhj= #lhk r, llk=Nil) in #llj r

```

Figure 4: Result of record type flattening (The End case is omitted.)

where $c, r \geq 0$, $1 \leq j(i) \leq c$ for every i ($1 \leq i \leq r$), and

$$\begin{aligned}
F_0 &= \{y_0 :: (l_1 \vec{k} :: \vec{\sigma}_3) \rightarrow \prod_{l_1 j \in l_1 \vec{j}} (l_1 j = N_0^{l_1 j}); \\
&\quad \dots; \\
&\quad \{y_0 :: (l_m \vec{k} :: \vec{\sigma}_3) \rightarrow \prod_{l_m j \in l_m \vec{j}} (l_m j = N_0^{l_m j})\} \\
F_i &= \{y_i :: (h_1 \vec{k} :: \vec{\sigma}_3) \rightarrow \prod_{h_1 j \in h_1 \vec{j}} (h_1 j = N_i^{h_1 j}); \\
&\quad \dots; \\
&\quad \{y_i :: (h_n \vec{k} :: \vec{\sigma}_3) \rightarrow \prod_{h_n j \in h_n \vec{j}} (h_n j = N_i^{h_n j})\}
\end{aligned}$$

In the form above, each N_i^{pq} belongs to a class of expressions specified by the following grammar

$$\begin{aligned}
\mathcal{M} ::= & \#hj (x_0 (\vec{h} \vec{k} = \mathcal{M})) \mid \#lj (x_i (\vec{l} \vec{k} = \mathcal{M})) \\
& \mid \#lk y_0 \mid \#hk y_i \mid C \mathcal{M} \cdots \mathcal{M} \\
& \mid \text{Stk } \mathcal{M} \mathcal{M} \mid \text{hd } \mathcal{M} \mid \text{t1 } \mathcal{M} \mid \Omega \\
& \mid \text{let } u_1 = \mathcal{M} \cdots u_v = \mathcal{M} \text{ in } \mathcal{M}
\end{aligned}$$

where $1 \leq i \leq r$, $h \in \vec{h}$, $k \in \vec{k}$, $l \in \vec{l}$, $j \in \vec{j}$.

The higher-order removal transformation method in this section extends the one given in Section 3.4 so that multiple record fields are properly handled. For this, for each different overloaded function g occurring in f and each different record label p , we need to identify a unique expression M such that ($p = M$) is the argument to g . In Section 3.4, such identification was possible, resorting to the linearity on a set of variables. However, in the present generalized setting, the linearity does not necessarily holds for the following reasons. First, the rewrite into **hd** and **t1** notations in Section 4.1 and the transformation rules in Section 4.4 and 4.5 may duplicate expressions, thus breaking the linearity. Second, the transformation rules in Section 4.2 apply the same function $S[M]$ to different arguments, namely, $\text{Stk } u \Omega$ and $\text{Stk } \Omega w$. Finally, a function in the target program may not have any record fields for some labels, when the corresponding fields are not referenced at all in the source program.

We remedy this problem by adopting a relaxed linearity condition described below. We allow the same function to be applied in two or more application sites. If a function is applied to different arguments which have an identical expression for the same field, we recognize the identical expression as the unique expression for that field. Suppose otherwise, *i.e.*, different expressions are assigned for the same field.

We can observe that such a set of incompatible expressions is produced only by the transformation rules in Section 4.2, which transform the same expression into a pair of expressions where Ω appears in complementary positions. This indicates that we can always make the set S of incompatible expressions turned into an identical expression, written $\text{lub } S$, by replacing the occurrences of Ω with appropriate expressions.

The formal definition of $\text{lub } S$ is as follows: let us write $M \preceq N$ if N is obtained by replacing some of the occurrences of Ω in M with arbitrary expressions. Then the binary relation \preceq induces a partial order with Ω being the least element. For any finite set S of expressions, we write $\text{lub } S$ to denote the least upper bound of S w.r.t. \preceq , if any. In particular, $\text{lub } \emptyset = \Omega$.

As an example of the relaxed linearity condition, consider the program given in Fig. 4. In the defining equation of the **Begin** case, **x0** has two application sites, in both of which the **selk** field of the argument is assigned **t1 (#selk y1)** and the **sehj** field is assigned **t1 (#sehj y1)**. Hence the expressions for the **selk** field and the **sehj** field are uniquely determined. In the same defining equation, **x1** is applied to different arguments that have incompatible expressions for the same field: for example, there are two different expressions for the **splk** field, namely, $M_1 = \text{Stk } \Omega (\#\text{splk } y_0)$ and $M_2 = \text{Stk } (\text{End } (\#\text{hk } u_1)) \Omega$. These expressions has the least compatible expression $\text{lub } \{M_1, M_2\} = \text{Stk } (\text{End } (\#\text{hk } u_1)) (\#\text{splk } y_0)$, which we use as the unique expression for the **splk** field. Similarly we have $\text{Stk } (\#\text{lhj } y_0) (\#\text{sphj } y_0)$ for the **sphj** field.

In case there is no corresponding application site of a function for some record field, we have $\text{lub } \emptyset = \Omega$ and consider Ω as the unique expression for that field. This means we consider as if there were a virtual argument expression, which is a record whose every field is set Ω . For example, in the defining equation for the case of **Nil** in Fig. 4, there is no application site of **x0**. So we consider as if there were an invisible application site **x0 (selk= Ω , sehj= Ω)**. We claim this is valid, since these **selk** and **sehj** fields are never referenced in the source program and hence they would represent unreachable points in the target program too.

The formal definition of the relaxed condition follows below.

```

parse' :: L → (lhj::L, llk::L, sehk::Stk L, selj::Stk L, sphj::Stk L, splk::Stk L)
      → (sehj::Stk L, selk::Stk L, llj::L, lhk::L, splj::Stk L, sphk::Stk L)
parse' (Begin t) x0 =
  let x1 = parse' t (sehk= Stk (hd (#selk x1)) (#sehk x0), selj= Stk (hd (#sehj x1)) (#selj x0), llk= #llk x0,
      lhj= Begin (hd (#splj x1)), splk= Stk (End (#lhk x1)) (#splk x0), sphj= Stk (#lhj x0) (#sphj x0))
  in (lhk= hd (#sphk x1), llj= #llj x1, sphk= tl (#sphk x1),
      splj= tl (#splj x1), selk= tl (#selk x1), sehj= tl (#sehj x1))

parse' (End t) x0 =
  let x1 = parse' t (sehk= tl (#sehk x0), selj= tl (#selj x0), llk= hd (#splk x0),
      lhj= hd (#sphj x0), splk= tl (#splk x0), sphj= tl (#sphj x0))
  in (lhk= hd (#sehk x0), llj= hd (#selj x0), sphk= Stk (#lhk x1) (#sphk x1),
      splj= Stk (#llj x1) (#splj x1), selk= Stk (#llk x0) (#selk x1), sehj= Stk (#lhj x0) (#sehj x1))

parse' Nil x0 = (lhk= #llk x0, llj= #lhj x0, sphk=empty, splj=empty, selk=Ω, sehj=Ω)

unpcaspar :: L → L
unpcaspar x = let r = parse' x (sehk=empty, selj=empty, lhj= #lhk r, llk=Nil, sphj=Ω, splk=Ω) in #llj r

```

Figure 5: Final transformation result

DEFINITION 1. A defining equation of the form (4.1) is called linear if

$M_i^{pq} = \text{lub} \{M \mid M \text{ appears in the form } x_i (pq = M, \dots)\}$
is well-defined for every i and pq such that $i = 0$ and $pq \in \vec{hk}$ or $1 \leq i \leq r$ and $pq \in \vec{lk}$.

This relaxed linearity condition is always satisfied when the original source program meets a certain linearity condition, as we will discuss in the following section.

The higher-order removal transformation in Section 3.4 can be generalized (to allow records with multiple field labels, $\vec{h}j$, \vec{lk} , \vec{hk} , and $\vec{l}j$) as follows. Every defining equation (4.1) of the higher-order function f is transformed into the following defining equation of the first-order function f' .

$$\begin{aligned}
f' &:: \sigma_1 \rightarrow (\vec{h}j :: \vec{\sigma}_3, \vec{l}k :: \vec{\sigma}_3) \rightarrow (\vec{hk} :: \vec{\sigma}_3, \vec{l}j :: \vec{\sigma}_3) \\
f' (C t_1 \dots t_c) x_0 &= \text{let } x_1 = f' t_{j(1)} F'_1 \\
&\quad \dots \\
&\quad x_r = f' t_{j(r)} F'_r \\
&\text{in } F'_0
\end{aligned} \tag{4.2}$$

where F'_0 and F'_i ($0 \leq i \leq r$) are defined by

$$\begin{aligned}
F'_0 &= \prod_{lj \in \vec{l}j} (lj = \mathcal{H}[\llbracket N_0^{lj} \rrbracket]) \otimes \prod_{hk \in \vec{hk}} (hk = \mathcal{H}[\llbracket M_0^{hk} \rrbracket]) \\
F'_i &= \prod_{hj \in \vec{h}j} (hj = \mathcal{H}[\llbracket N_i^{hj} \rrbracket]) \otimes \prod_{lk \in \vec{lk}} (lk = \mathcal{H}[\llbracket M_i^{lk} \rrbracket])
\end{aligned}$$

with \mathcal{H} being the transformation specified by the following set of rules.

$$\begin{aligned}
\mathcal{H}[\llbracket \text{let } u_1 = M_1 \dots u_v = M_v \text{ in } M \rrbracket] &= \mathcal{H}[\llbracket M[u_1 \setminus M_1, \dots, u_v \setminus M_v] \rrbracket]
\end{aligned} \tag{4.3a}$$

$$\mathcal{H}[\llbracket x_i M \rrbracket] = x_i \quad (0 \leq i \leq r) \tag{4.3b}$$

$$\mathcal{H}[\llbracket y_i \rrbracket] = x_i \quad (0 \leq i \leq r) \tag{4.3c}$$

We note that the transformation rule (4.3a) never causes infinite unwinding, since we can prove that, in every mutually recursive local definition $\text{let } u_1 = M_1 \dots u_v = M_v \text{ in } M$ (appearing in the intermediate transformation result), every variable u_k ($1 \leq k \leq v$) occurs in one of the expressions M_1, \dots, M_v behind some x_i , i.e., in the form $x_i (\dots u_k \dots)$. This, together with the rule (4.3b), prevents nested unwinding.

The equation (3.3) in Section 3.4 is generalized as follows.

For every $l_0 \in \vec{l}$ and $j \in \vec{j}$,

$$\begin{aligned}
&\#l_0j (f L G (\prod_{lk \in \vec{lk}} (lk = L'_{lk}))) \\
&= \text{let } r = f' L (\prod_{hj \in \vec{h}j} = \#hj (G \\
&\quad (\prod_{hk \in \vec{hk}} (hk = \#hk r) \otimes \prod_{lk \in \vec{lk}} (lk = L''_{lk}))))), \\
&\text{in } \#l_0j r
\end{aligned}$$

where $L''_{lk} = L'_{l_0k}$ if $l = l_0$; otherwise, $L''_{lk} = \Omega$.

Applying this higher-order removal transformation to parse' , we obtain the final transformation result given in Fig. 5.

5. CONDITION FOR SUCCESSFUL TRANSFORMATION

This section discusses the sufficient condition on the source program (i.e., the pair of a producer and a consumer function) for our transformation algorithm to succeed.

Let us assume that each defining equations of the producer and consumer functions are given in the following form

$$\begin{aligned}
\mathbf{f} (C t_1 \dots t_c) x_0 &= \\
\text{let } x_1 &= \mathbf{f} t_{j(1)} (p_1 = M_1^{p_1}, \dots, p_n = M_1^{p_n}) \\
&\quad \dots \\
x_r &= \mathbf{f} t_{j(r)} (p_1 = M_r^{p_1}, \dots, p_n = M_r^{p_n}) \\
\text{Stk } a_1 \text{ } a_{s_1} &= N_1 \\
&\quad \dots \\
\text{Stk } a_m \text{ } a_{s_m} &= N_m \\
\text{in } (q_1 = M_0^{q_1}, \dots, q_{n'} = M_0^{q_{n'}})
\end{aligned} \tag{5.1}$$

where $c, r, m \geq 0$, $1 \leq j(i) \leq c$ for every i ($1 \leq i \leq r$). We assume $M_i^{p_i}$'s, $M_0^{q_i}$'s, and N_k 's belong to a subset of well-typed first-order expressions, defined by the following grammar.

$$\begin{aligned}
M &::= \#p x_0 \mid \#q x_i \mid a_s \mid a_{s_s} \mid C M \dots M \\
&\quad \mid \text{Stk } M M \mid \text{empty} \mid \Omega
\end{aligned}$$

where $p \in \vec{p}$, $q \in \vec{q}$, $1 \leq i \leq r$, $1 \leq s \leq m$.

DEFINITION 2 (LINEARITY CONDITION). A function is called stack linear if its every defining equation in the form (5.1) has at most one occurrence of expressions a_i and a_{s_i} for each i ($1 \leq i \leq m$).

A function is called context linear if it is stack linear and there is at most one occurrence of $\#p x_0$ for each $p \in \vec{p}$.

A function is called *recursion linear* if it is *stack linear*, $j(1), \dots, j(r)$ are *pairwise distinct*, and there is at most one occurrence of $\#q x_i$ for each i ($0 \leq i \leq r$) and $q \in \bar{q}$.

Intuitively, a context linear function is constrained to have at most one reference to each accumulating parameter (indexed by a distinct record label); a recursion linear function is constrained to visit every substructure of a node of the input datatype at most once and also every field in the returned record is referenced at most once.

We can show, subject to the following condition on the source program, that the linearity condition required for the higher-order transformation in Section 4.6 is satisfied and hence the transformation succeeds.

THEOREM 1 (SUCCESSFUL TRANSFORMATION). *The transformation in Section 4 succeeds, if the producer is a context linear function and the consumer is a recursion linear function.*

We can also show that a closure property holds under a stronger condition.

THEOREM 2 (CLOSURE). *If the transformation in Section 4 is applied to a pair of a producer function and a consumer function that are both context and recursion linear, then the result is also a context and recursion linear function.*

In the previous sections, we have implicitly exploited this property. The functions we have considered so far are all both context and recursion linear and hence we were able to apply our transformation algorithm repeatedly to derive a single function that fuses three functions.

6. RELATED WORK

To the best of author’s knowledge, there has been no proposal of a fusion transformation for functions that allows both stacks and accumulating parameters. Recent studies on fusion has seen significant advances in dealing with accumulating parameters [13, 21, 25], but none of them deal with stacks.

Our transformation algorithm, as for fusion of functions with accumulating parameters, has a competitive transformation power to Voigtländer’s *lazy composition* [25] — both of them allow the pair of a context linear producer function and a recursion linear consumer function. Also, they make use of the circular tupling technique. Interestingly, despite of these similarities, they are based on quite different transformation principles. Lazy composition is an algorithmic instance of unfold/fold-technique [3], while ours are driven by a series of type-directed transformations. We leave more precise comparison to future investigation.

Apart from the researches in functional programming, similar program composition techniques have been studied for attribute grammars as well. Ganzinger and Giegerich [6] developed an algorithm, called *descriptive composition*, where accumulating parameters are allowed as so-called *inherited attributes*. Recently, Nakano [15] extended this composition method to allow stacks.

The present work draws some ideas from their work — use of record labels and flattening of nested records. However, we emphasize that the present work provides an alternative, modular presentation of the same transformation,

whereas the attribute grammar composition method has a ‘big’ transformation step that is hardly modularized. The modularity of our transformation algorithm is achieved by representing intermediate transformation results as higher-order programs, which cannot be expressed in attribute grammars. We also note that our algorithm allows non-recursion linear producer functions (Theorem 1), which cannot be expressed within the limited recursion scheme of attribute grammars.

7. CONCLUSION AND FUTURE WORK

We have presented a new program transformation algorithm for fusion that allows both stacks and accumulating parameters. The proposed algorithm is partly comprised of some of the transformation techniques that have been developed in the communities of functional programming [8, 1, 17] and attribute grammars [6, 15], but reformulates and refines them into a series of type-directed transformations, leading to a logically clean account for the transformation process. We have also shown that the transformation algorithm, under a linearity condition, satisfies a closure property, which allows repeated applications of fusion.

Our algorithm eliminates all the intermediate data structures created by the producer function and yields a first-order program that uses accumulating parameters and stacks, on which parts of the final output are held. However, the elimination of intermediate data structures alone does not necessarily imply a better efficiency, due to the following factors that may degrade the overall run-time efficiency. (i) The record type flattening in Section 4.5 duplicates every record field of the producer function by the number of record labels of the consumer function. This indicates that the execution cost of each recursive call is increased by a constant factor. (ii) Tupling/untupling operations on records arise per each recursive call. In other words, records in the resulting program could be seen as another intermediate structure. (iii) The resulting function includes a circular let, whose execution could be costly.

These factors of deterioration would be compensated by a post-optimization on the resulting program, however. (The concern of this paper is on the derivation of the fusion algorithm and the empirical study on the effect of optimizations discussed below is left to future investigation.) Voigtländer [25] proposed a post-processing method for reducing “ballasts” in his fused programs (tupling/untupling, circular let, etc.) Van Groningen [23] proposed a compiler optimization technique for reducing the overhead raised by allocations of tuples. These techniques would be effectively applied to our fusion results, too. Another prospective direction would be application-specific optimizations. In [16, 18], Nakano and the author developed an algorithm that derives memory efficient XML stream transformers from the result of attribute grammar composition. Refining this method so that it applies to a class of programs involving stacks would be an interesting topic for future research.

In this paper, we consider stacks in producer functions only and defined stacks as a datatype that has only one data constructor. These are not a fundamental restriction of the algorithm, however. It is easy to allow stacks in consumer functions, in which case the resulting program would involve “stacks of stacks”. We may also refine our algorithm so that stacks are defined as a datatype with an explicit empty stack constructor and with additional stack constructors as well.

(This would be useful, say, for parsing a language that comprises of two or more kinds of braces.) Putting it more general, we may even allow any polynomial datatypes in place of stacks. Nevertheless, to deal with datatypes that consist of two or more constructors generally, we would need to allow case distinction over these constructors. This would complicate the transformation process to a considerable degree, as so does the extension of the attribute grammar composition with conditionals [2]. A complete solution to this issue would require further elaboration and is left to future work.

Acknowledgment

I am grateful to Masahito Hasegawa and Jacques Garrigue for their valuable comments on an earlier draft. I also thank anonymous reviewers, whose comments were very helpful for improving the paper.

8. REFERENCES

- [1] R. S. Bird. Using circular programs to eliminate multiple traversals of data. *Acta Informatica*, 21:239–250, 1984.
- [2] J. Boyland and S. L. Graham. Composing tree attributions. In *Proc. of the 21th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 375–388. ACM Press, 1994.
- [3] R. M. Burstall and J. Darlington. A transformation system for developing recursive programs. *Journal of ACM*, 24(1):44–67, 1977.
- [4] W.-N. Chin. Safe fusion of functional expressions II: Further improvements. *Journal of Functional Programming*, 4(4):515–555, 1994.
- [5] O. Chitil. Type inference builds a short cut to deforestation. In *Proc. of International Conference on Functional Programming (ICFP’99)*, pages 249–260. ACM Press, 1999.
- [6] H. Ganzinger and R. Giegerich. Attribute coupled grammars. In *Proceedings of the ACM SIGPLAN ’84 Symposium on Compiler Construction*, volume 19(6) of *SIGPLAN Notices*, pages 157–170, June 1984.
- [7] A. Gill. *Cheap Deforestation for Non-strict Functional Languages*. PhD thesis, University of Glasgow, 1996.
- [8] A. Gill, J. Launchbury, and S. Peyton Jones. A short cut to deforestation. In *Proc. of the Conference on Functional Programming Languages and Computer Architecture*, pages 223–232. ACM Press, June 1993.
- [9] The Haskell home page. <http://www.haskell.org/>.
- [10] J. E. Hopcroft, R. Motwani, and J. D. Ullman. *Introduction to automata theory, languages, and computation*. Addison-Wesley, 2nd edition, 2001.
- [11] Z. Hu, H. Iwasaki, and M. Takeichi. Deriving structural hylomorphisms from recursive definitions. In *Proc. of International Conference on Functional Programming*, pages 73–82, Philadelphia, Pennsylvania, May 1996.
- [12] P. Johann. Short cut fusion is correct. *Journal of Functional Programming*, 13(4):797–814, 2003.
- [13] K. Kakehi, R. Glück, and Y. Futamura. An extension of shortcut deforestation for accumulative list folding. *IEICE Transactions on Information and Systems*, E85-D(9):1372–1383, 2002.
- [14] J. Launchbury and T. Sheard. Warm fusion: Deriving build-catas from recursive definitions. In *International Conference on Functional Programming Languages and Computer Architecture, (FPCA’95)*, pages 314–323. ACM Press, 1995.
- [15] K. Nakano. Composing stack-attributed tree transducers. Technical Report METR-2004-01, Department of Mathematical Informatics, University of Tokyo, Japan, 2004.
- [16] K. Nakano and S. Nishimura. Deriving event-based document transformers from tree-based specifications. In *LDTA’2001 Workshop on Language Descriptions, Tools and Applications*, volume 44 of *Electronic Notes in Theoretical Computer Science*. Elsevier Science, 2001.
- [17] S. Nishimura. Correctness of a higher-order removal transformation through a relational reasoning. In *Programming Language and Systems, First Asian Symposium, APLAS 2003 Proceedings*, volume 2895 of *LNCS*, pages 358–375. Springer Verlag, 2003. A full version is available as a preprint Kyoto-Math 2003-06 from <http://www.math.kyoto-u.ac.jp/~susumu/papers/aplas03-long.ps.gz>.
- [18] S. Nishimura and K. Nakano. XML stream transformer generation through program composition and dependency analysis. submitted.
- [19] A. Pitts. Parametric polymorphism and operational equivalence. *Mathematical Structures in Computer Science*, 10(3):321–359, 2000.
- [20] T. Sheard and L. Fegaras. A fold for all seasons. In *International Conference on Functional Programming Languages and Computer (FPCA’93)*, pages 233–242. ACM Press, 1993.
- [21] J. Svenningsson. Shortcut fusion for accumulating parameters & zip-like functions. In *Proc. of the 2002 International Conference on Functional Programming*, 2002.
- [22] A. Takano and E. Meijer. Shortcut deforestation in calculational form. In *Proc. of International Conference on Functional Programming Languages and Computer Architecture (FPCA’95)*, pages 306–313. ACM Press, 1995.
- [23] J. H. van Groningen. Optimising recursive functions yielding multiple results in tuples in a lazy functional language. In *Implementation of Functional Languages, 11th International Workshop, IFL’99*, volume 1868 of *Lecture Notes in Computer Science*, pages 59–76. Springer Verlag, 1999.
- [24] J. Voigtländer. Using circular programs to deforest in accumulating parameters. In *ASIAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, pages 126–137. ACM Press, 2002.
- [25] J. Voigtländer. Using circular programs to deforest in accumulating parameters. *Higher-Order and Symbolic Computation*, 17(1), 2004.
- [26] P. Wadler. Theorems for free! In *International Conference on Functional Programming and Computer Architecture (FPCA’89)*, pages 347–359. Addison-Wesley, 1989.
- [27] P. Wadler. Deforestation: transforming programs to eliminate trees. *Theoretical Computer Science*, 73(2):231–248, June 1990.