



Warm fusion in Stratego: A case study in generation of program transformation systems

Patricia Johann^a and Eelco Visser^b

^a *Department of Mathematics, Bates College, Lewiston, Maine 04240, USA*

E-mail: pjohann@bates.edu

^b *Institute of Information and Computing Sciences, Universiteit Utrecht, P.O. Box 80089, 3508 TB Utrecht, The Netherlands*

E-mail: visser@acm.org

Stratego is a domain-specific language for the specification of program transformation systems. The design of Stratego is based on the paradigm of rewriting strategies: user-definable programs in a little language of strategy operators determine where and in what order transformation rules are (automatically) applied to a program. The separation of rules and strategies supports modularity of specifications. Stratego also provides generic features for specification of program traversals.

In this paper we present a case study of Stratego as applied to a non-trivial problem in program transformation. We demonstrate the use of Stratego in eliminating intermediate data structures from (also known as *deforesting*) functional programs via the *warm fusion* algorithm of Launchbury and Sheard. This algorithm has been specified in Stratego and embedded in a fully automatic transformation system for kernel Haskell. The entire system consists of about 2600 lines of specification code, which breaks down into 1850 lines for a general framework for Haskell transformation and 750 lines devoted to a highly modular, easily extensible specification of the warm fusion transformer itself. Its successful design and construction provides further evidence that programs generated from Stratego specifications are suitable for integration into real systems, and that rewriting strategies are a good paradigm for the implementation of such systems.

1. Introduction

Automatic program transformation is applied in many branches of software engineering – including application generation and compiler construction – to translate high-level, but inefficient, specification code to lower-level and more efficient implementation code. It plays a particularly important role in compilers for functional programming languages [1,6,7,23,26].

1.1. Transforming programs with rewriting strategies

An important paradigm for the description of program transformation systems is that of rewrite rules. Ad-hoc implementation of transformation systems based on rewrite rules can be difficult, however, because the rules must be embedded in algorithms that

determine strategies for applying them. Stratego [15,31,34,35] is a domain-specific language for the specification of program transformation systems. Its design is based on the paradigm of rewriting strategies. Rewriting strategies combine user-definable rewriting-based programs with a little language of independent strategy operators that can be used to specify where and in what order transformation rules are applied to a program.

Stratego's separation of rewrite rules from the strategies which control their application facilitates modular specification of program transformations: transformation rules are specified independently of the application strategy and can be reused in more than one strategy. Stratego also offers both fine and coarse grain control over the application of transformation rules. This control makes it possible to specify the exact forms that programs can assume at various stages of processing. It also allows the programmer to govern the interactions between individual transformation rules. The Stratego compiler translates specifications to C programs that transform abstract syntax trees to abstract syntax trees.

In [35] it is shown how rewriting strategies can be used to modularly specify and implement optimizers for functional programs. A set of transformation rules is combined into a code simplification algorithm by means of a strategy that traverses programs and applies rules where appropriate. The emphasis in [35] is on rules that are independently applicable. As demonstrated there, it is particularly easy to combine transformation rules into different simplification strategies by adding or omitting rules. But in many settings the construction of interrelated transformation rules from several more primitive rules is necessary.

1.2. Applying strategies in deforesting functional programs

In this paper we present a case study illustrating the use of rewriting strategies to eliminate intermediate data structures from (deforest) functional programs. Deforestation algorithms typically perform a number of smaller transformations before determining whether or not the deforestation is considered successful. Combining primitive rules into complex program transformations often requires the exchange of more information between their rules than is contained in the individual program fragments they transform. The parameterization of strategies supported by Stratego provides a means of specifying and implementing rules which pass such information between them. In the case study presented here information exchanged between transformations takes the form, for example, of assumptions about bindings, dynamic rewrite rules that recognize recursive function calls, and terms generated by splitting functions to facilitate program transformation. Parameterized strategies have not been used extensively in previous Stratego specifications.

We have specified the warm fusion algorithm of Launchbury and Sheard [14] in Stratego. This technique combines the cheap deforestation based on foldr-build fusion of Gill et al. [7,8] with the fold promotion of Sheard and Fegaras [24] and a generalization of the technique of Peyton Jones and Launchbury [22] for splitting a function into a worker and a wrapper. The foldr-build fusion, which has been implemented in

the Glasgow Haskell Compiler (GHC), requires the manual transformation of functions to build-foldr form and is only defined for lists. The warm fusion algorithm generalizes cheap deforestation to arbitrary regular data types and automatically derives more general build-cata forms.

As a case study, the warm fusion algorithm is an interesting example of a non-trivial program transformation, and its specification provides evidence of the feasibility of implementing such transformations in Stratego. The case study supplies experience with the design and implementation of a complete transformation system, including interfaces with a parsing and typechecking front-end and a pretty-printing back-end for Haskell. The application to Haskell provides an environment in which to assess the effectiveness of warm fusion for deforesting more realistic programs than would otherwise be possible. The case study also demonstrates Stratego's support for the construction of transformation rules that combine basic transformation steps in various ways, the description and checking of intermediate representation formats, language independent definition of substitution and the renaming of bound variables, and the discovery of new programming idioms resulting from the strategy-induced shift away from a purely functional implementation style.

Warm fusion is also an interesting problem in its own right. The first fully automatic implementation of warm fusion was hand-coded in Haskell in 1997 [11]. The algorithm had previously been implemented only as "a toolbox of operations" [14]. This is perhaps because the description of warm fusion in [14] elides much of the detail required to turn the theory into practice. The type-driven nature of the algorithm, in particular, is fundamental to its automation, as well as to its extension to non-list data structures. The critical dependence of warm fusion on type information is reflected in its Stratego specification.

The product of our case study is a fully automatic implementation of the warm fusion algorithm. This implementation could be an important step toward the use of warm fusion in compilers or as a preprocessor for (library) programs. It can also serve as a basis for further experimentation with extensions of cheap deforestation; Stratego makes it easy, for example, to modify the set of program transformation rules and to experiment with a variety of application orders. Experience with a working system often gives rise to a deeper understanding of its underlying algorithm. It was such experience that led, for instance to our "double splitting" wrapper-worker technique for recognizing certain variables as static parameters of programs undergoing warm fusion. (This step happens "automagically" in [14].) This technique has since been incorporated into the Haskell implementation of warm fusion detailed in [11].

1.3. Outline

In the next section we briefly review some background on deforestation, discuss the principles of cata-build fusion, and illustrate the warm fusion transformation technique by means of an example. In section 3 we give an overview of the operators of System S, a calculus for the definition of tree transformations, as well as of the syntactic

abstractions built on System S that form Stratego. In section 4 we present the overall architecture of the warm fusion transformation tool built with Stratego. In sections 5–7 we discuss several highlights from the specification, focusing particularly on some of the new programming idioms that have emerged during the process of specifying the warm fusion algorithm in Stratego. The full text of the specification can be found in [13].

2. Warm fusion

Modularity in functional programming is achieved by dividing programs into small, generally applicable functions that communicate via data structures. Such functions are commonly defined as recursive operations that construct and deconstruct data structures. The definitions in figure 1 are common examples of such functions; `sum` and `foldr` consume lists, `upto` produces lists, and `map` does both. Using these functions we can, for instance, define the sum of the squares of the numbers `lo` to `hi` as

```
sos :: Int -> Int -> Int
sos = \lo hi -> sum(map(square)(upto lo hi))
```

where the function `square` is defined as

```
square :: Int -> Int
square = \x -> (x * x)
```

```
data Bool = True | False;
data List a = Nil | Cons a (List a);
map :: (a -> b) -> List a -> List b;
map = \f l ->
  case l of {
    Nil      -> Nil;
    Cons x xs -> Cons(f x)(map f xs)};
foldr :: b -> (a -> b -> b) -> List a -> b;
foldr = \n c xs ->
  case xs of {
    Nil      -> n;
    Cons y ys -> c y (foldr n c ys)};
upto :: Int -> Int -> List Int;
upto = \low high ->
  case low > high of {
    True  -> Nil;
    False -> Cons low (upto(low + 1)(high))};
sum :: List Int -> Int;
sum = foldr 0 (+);
sos :: Int -> Int -> Int;
sos = \lo hi -> sum(map(square)(upto lo hi))
```

Figure 1. Recursive functions on lists.

This implementation of the sum-of-squares function is straightforward and modular. Its disadvantage is that it constructs, traverses, and deconstructs two intermediate lists – even though both the input and output of the computation are integers. This is computationally expensive, both slowing execution time and increasing heap space requirements.

It is often possible to avoid manipulating intermediate data structures by using a more elaborate style of programming in which parts from component functions are intermingled. In this monolithic style of programming the sum-of-squares function is defined as

```
sos' :: Int -> Int -> Int
sos' = \lo hi ->
  let {sos'' :: Int -> Int;
      sos'' = \i -> case i > hi of {
                    True  -> 0;
                    False -> square(i) + sos''(i + 1)}}
  in sos''(lo)
```

Note that no intermediate data structures at all are processed by `sos'`. In this case, eliminating the manipulation of intermediate lists results in an order of magnitude gain in program performance.

Experienced programmers writing a square summing function would instinctively produce `sos'` rather than `sos`; small functions like `sos` are easily optimized at the keyboard. But when programs are either very large or very complex, even experienced programmers may find that eliminating intermediate data structures by hand is not a very attractive alternative to the modular style of programming. In such situations a tool for automatically eliminating them is needed.

2.1. Deforestation

Automatic elimination of intermediate data structures by transformation combines the clarity and maintainability of the modular style of programming with the efficiency of the monolithic style. The process of eliminating intermediate data structures from programs is often called *deforestation* after an early transformation technique of Wadler [36] which removes tree-like data structures from first-order programs.

In Wadler's deforestation, compositions of treeless expressions (a syntactic restriction of normal expressions that allows no intermediate data structures) are transformed into new treeless expressions. The technique uses function unfolding to expose consumption of constructors by case selections. Subsequent folding creates new recursive functions. To prevent non-termination of unfolding, global program patterns must be monitored. Because this is computationally expensive, Wadler's deforestation has not been incorporated into functional language compilers.

Gill et al. [7,8] introduce a less general, but cheaper, variant of deforestation for list-producing and -consuming functions. The key observation underlying their *short*

cut to deforestation is that many list-manipulating functions can be written in terms of the uniform list-consuming function `foldr` and the uniform list-producing function `build`. Since `foldr` is essentially another name for the standard catamorphism for lists, we denote it by `cata-list` in this paper. And since the `build` function of Gill et al. is the instantiation to lists of a more general `build` function applying to arbitrary regular data types, we denote it by `build-list` below.

Operationally, `cata-list` takes as input types `a` and `b`, a replacement function `f1 :: a -> b -> b` for `Cons[a]`, a replacement function `f2 :: b` for `Nil[a]`, and a list `ls` of type `List a`. (The list constructors `Cons` and `Nil` have the polymorphic types `forall a. a -> List a -> List a` and `forall a. List a`, respectively, and so must be instantiated for each particular list type; the notation `e[t]` instantiates the polymorphic expression `e` to type `t`.) It replaces by `f1` and `f2`, respectively, all occurrences of `Cons[a]` and `Nil[a]` in `ls` which actually contribute to the result of the computation. The result is a value of type `b`. The function `build-list`, on the other hand, takes as input a function `g` providing a type-independent template for constructing lists and instantiates its “abstract” list constructors with appropriate instances of the “concrete” list constructors `Cons` and `Nil`. In other words, if `g` is any function with polymorphic type `forall b . b -> (a -> b -> b) -> b`, then

$$\text{build-list}[a](g) = g[\text{List } a] (\text{Nil}[a]) (\text{Cons}[a])$$

Compositions of list-consuming and -producing functions defined in terms of `cata-list` and `build-list` can be simplified (deforested) by means of the short cut fusion rule for lists:

$$\text{cata-list}[a][t](f1, f2)(\text{build-list}[a](g)) = g[t] f1 f2$$

The short cut describes one precise way in which compilers can take advantage of uniformity in the production and consumption of lists to optimize programs which manipulate them. It makes sense intuitively: the result of a computation is the same regardless of whether the function `g` is first applied to `Cons` and `Nil` and occurrences of `Cons` and `Nil` in the resulting list are then replaced by `f1` and `f2`, respectively, or the abstract constructors in `g` are replaced by `f1` and `f2`, respectively, directly. The fact that `g` is polymorphic in its result type `t` ensures the correctness of this fusion rule [12].

2.2. An example of cata-build fusion

Figure 2 shows the `build-cata` forms of the functions in figure 1. The notation `/\a -> e` denotes the abstraction of type variable `a` from the expression `e`. Such an expression has type `forall a . t`, where `t` is the type of `e`. Type abstraction is normally implicit in definitions in Haskell because it only occurs at the top of a definition, i.e., a Haskell definition `f = \x -> e` that is polymorphic in type variable `a` abbreviates the definition `f = /\a -> \x -> e`.

The deforested function `sos'` can be derived from `sos` by inlining the definitions in figure 2 and applying the short cut in conjunction with the standard program simpli-

```

map :: (a -> b) -> List a -> List b;
map = \f l ->
    build[List b](/\t -> \(n :: t) (c :: (b -> t -> t)) ->
        cata[List a][t](n, \(y :: b) -> c(f y)) l);

foldr :: b -> (a -> b -> b) -> List a -> b;
foldr = \n c -> cata[List a][b](n, c);

upto :: Int -> Int -> List Int;
upto = \lo hi ->
    build[List Int]
        (/\t -> \(n :: t) (c :: int -> t -> t) ->
            let {upto' :: Int -> t;
                upto' = \i -> case i > hi of {
                    True -> n;
                    False -> c(i)(upto'(i + 1))}}
            in upto'(lo));

sum :: List Int -> Int;
sum = cata[List Int][Int](0, (+))

```

Figure 2. Functions in build-cata form.

fication rules in section 7. Inlining the (type-instantiated) function definitions for `sum`, `map` and `square` gives

```

sos = \lo hi -> sum(map(square)(upto lo hi))
    = \lo hi ->
        cata[List Int][Int](0, (+))
        ((\f l -> build[List Int]
            (/\t -> \(n :: t) (c :: Int -> t -> t) ->
                cata[List Int][t](n, \(y :: Int) -> c(f y)) l))
        (\x -> x * x) (upto lo hi))

```

Simplifying the application of `map` to `square` and `upto lo hi` produces

```

= \lo hi ->
    cata[List Int][Int](0, (+))
    (build[List Int]
        (/\t -> \(n :: t) (c :: Int -> t -> t) ->
            cata[List Int][t](n, \(y :: Int) -> c(y*y))(upto lo hi))

```

Applying the short cut rule to the `cata-build` pair and simplifying yields

```

= \lo hi ->
    cata[List Int][Int](0, \(y :: Int) -> (+)(y*y))(upto lo hi)

```

Inlining the definition for `upto` gives

```

= \lo hi ->
  cata[List Int][Int](0, \y :: Int -> (+)(y*y))
  (build[List Int]
   (\t -> \ (n :: t) (c :: Int -> t -> t) ->
    let {upto' :: Int -> t;
        upto' = \i -> case i > hi of {
                        True -> n;
                        False -> c(i)(upto'(i+1))}}
    in upto'(lo))

```

Using the short cut and simplifying once more gives

```

sos = \lo hi ->
  let {upto' :: Int -> Int;
      upto' = \i -> case i > hi of {
                  True -> 0;
                  False -> (i*i) + (upto'(i+1))}}
  in upto'(lo)

```

Up to renaming and inlining of square in the local function, this is precisely the definition of `sos'`.

2.3. Warm fusion: Automatically deriving cata-build forms

The short cut fusion rule calculates program improvement based on a program's explicit local structure. To do this, it requires that functions be written in the highly stylized build-cata form, rather than using explicit recursion. But this is often not the most natural way to develop programs. Moreover, because `build` does not have a Hindley–Milner type – and so can only be used in certain well-defined ways – providing it for programmers' direct use is problematic. The warm fusion algorithm of Launchbury and Sheard [14] was designed to automate the safe introduction of `build` into recursive list-processing functions, as well as the transformation of the resulting functions into equivalent ones in build-cata form.

The existence of a catamorphism and a `build` function for each regular data type makes it possible to generalize the warm fusion method to arbitrary regular data types. If F is a functor defining a regular data type, then the catamorphism $\text{cata}[F\ a_1 \dots a_n][t](f_1, \dots, f_n)$ replaces the constructors of a data structure of type $F\ a_1 \dots a_n$ with the functions f_i . The result of the catamorphism has type t . The data structure-producing function $\text{build}[F\ a_1 \dots a_n]$, on the other hand, takes as input a polymorphic function g which constructs the kind of data structures associated with the functor F . It replaces the abstract data constructors of g by the concrete data constructors c_i to produce the data structure of type $F\ a$ whose description g embodies. That is,

$$\text{build}[F\ a_1 \dots a_n](g) = g[F\ a_1 \dots a_n]\ c_1 \dots c_n$$

Note that `cata-list[a][t]` is just `cata[List a][t]` and `build-list[a]` is precisely `build[List a]`, where `List` is the functor associated with the list data type. The short cut fusion rule for `cata-list` and `build-list` generalizes to:

$$\text{cata}[F\ a1\dots an][t](f1, \dots, fn)(\text{build}[F\ a1\dots an](g)) = g[t]\ f1\dots fn$$

2.4. Warm fusion by example

To illustrate the process of warm fusion we will examine the transformation of the consumer-producer `map`. In the following examples we will omit the type declarations for variables and constructors when these are clear from the context or from previous declarations.

2.4.1. Abstracting from constructors

The goal of the preprocessing step of warm fusion is to transform a recursive definition into a definition in `build-cata` form:

$$f = /\!a1 \dots an \rightarrow \!x \dots \rightarrow \\ \text{build}[F\ a1\dots an](/\!t \rightarrow \!c1\dots cn \rightarrow \\ \text{cata}[F\ a1\dots an][t](h1, \dots, hm)\ x)$$

The functional argument of `build` is a catamorphism that consumes the input data structure `x` and builds up a structure that is constructed with the abstract constructors `ci`. This transformation shifts the recursion boundary of the function from the site of construction of the result data structure to the site of consumption of the input data structure. All recursion in `build-cata` forms is expressed via catamorphisms.

The first phase of the transformation abstracts away from the concrete constructors in the body of the function. This cannot be done simply by replacing all constructors in the body by variables, however, because not all occurrences of constructors necessarily contribute to the result of the computation. By applying `cata[F a1...an][t](c1, ..., cn)` to the body of the function, the result-producing constructors are transformed into the corresponding abstract constructors `ci`.

The identity

$$x = \text{build}[F\ b1\dots bn](/\!t \rightarrow \!c1 \dots cn \rightarrow \\ \text{cata}[F\ b1\dots bn][t](c1, \dots, cn)\ x)$$

is used to introduce this catamorphism to the body. For `map` this becomes

$$\text{map} = /\!a\ b \rightarrow \!f\ l \rightarrow \\ \text{build}[List\ b](/\!t \rightarrow \!(n :: t)\ (c :: b \rightarrow t \rightarrow t) \rightarrow \\ \text{cata}[List\ b][t](n, c)(\\ \text{case } l \text{ of } \{ \\ \text{Nil} \rightarrow \text{Nil}; \\ \text{Cons } x\ xs \rightarrow \text{Cons}(f\ x)(\text{map}[a][b]\ f\ xs))\})$$

Distribution of the catamorphism over the case expression gives

```
map = /\a b -> \f l -> build[List b](/\t -> \n c ->
case l of {
  Nil -> cata[List b][t](n, c) Nil;
  Cons x xs -> cata[List b][t](n, c)(Cons(f x)(map[a][b] f xs)))}
```

Specialization of the catamorphism to the constructors that it is applied to produces:

```
map = /\a b -> \f l -> build[List b](/\t -> \n c ->
case l of {
  Nil -> n;
  Cons x xs -> c(f x)(cata[List b][t](n, c)(map[a][b] f xs)))}
```

Note that the catamorphism is applied to the recursive second argument of the abstract replacement function for Cons.

2.4.2. Splitting off the recursive consumer

We have now abstracted away from the result-producing constructors of map and written it in the form of an abstracted call to build. Next we derive a catamorphism to replace the case analysis in map's body. This is accomplished according to the steps outlined in the remainder of this section.

First the function body is split into two new definitions. For map we get the “wrapper” map and the “worker” map# (a generally applicable idea first presented in [22]):

```
map = /\a b -> \f l ->
      build[List b](/\t -> \n c -> map# l [t] n c)
map# = \l -> /\t -> \n c ->
case l of {
  Nil -> n;
  Cons x xs -> c(f x)(cata[List b][t](n, c)(map[a][b] f xs))}
```

The splitting has the effect of isolating a recursive definition not involving build.

Note that the function f and the type variables a and b are not passed to map#. From the definition of map before splitting it is clear that these arguments are passed unchanged to the recursive call of map. That is, they are *static parameters* of map. Since we do not abstract over them, the static parameters of a function remain free in the definition of its worker. This means that f , a , and b remain free in map#. When, at the end of the transformation, the transformed version of the function's worker is folded back into the definition of its wrapper, its static parameters will become bound again.

By unfolding the wrapper in the worker we obtain a recursive definition of the worker. For map we get

```
map# = \l -> /\t -> \n c ->
case l of {
  Nil -> n;
  Cons x xs ->
    c(f x)(cata[List b][t](n, c)
      ((/\a' b' -> \f' l' ->
```

```
build[List b](/\t' -> \n' c' -> map# l' [t'] n' c'))
[a][b] f xs))}
```

Beta-reduction and short cut fusion reduces this to

```
map# = \l -> /\t -> \n c ->
      case l of {
        Nil -> n;
        Cons x xs -> c(f x)(map# xs [t] n c)}
```

Observe now that all arguments except for `l` are static parameters of `map#`. By repeating the splitting and unfolding procedure once more we get

```
map# = \l -> /\t -> \n c -> map## l
map## = \l -> case l of {
          Nil -> n;
          Cons x xs -> c(f x)(map## xs)}
```

The parameters `t`, `n`, and `c` of `map#` are now also recognized as static in `map##`. The free variable `f` in `map##` is inherited from `map`. In [14], mechanical recognition of the abstracted constructors as static parameters (when they are), happens magically.

2.4.3. Recursion to catamorphism

Finally, the recursive definition of `map##` is turned into a catamorphism by means of fold promotion. Fold promotion is based on a generic promotion theorem introduced by Malcolm [16]. The promotion theorem, which has its origins in a categorical description of programming [9], describes conditions under which the composition of an arbitrary (strict) function and a catamorphism over a regular data type may be fused to arrive at a new catamorphism equivalent to the original composition. For `map##` the promotion theorem takes the form

```
map## Nil = h1,
map##(Cons y1 y2) = h2 y1 (map## y2)
-----
map##(cata[List a][List a](Nil, Cons) xs)
  = cata[List a][t](h1, h2) xs
```

This means that we can find `h1` and `h2` by applying `map##` to `Nil` and `Cons y1 y2`, respectively, and abstracting from the recursive call to `map##`. For `Nil` this simply produces the abstracted constructor `n`. For `Cons` we get

```
h2 = \y1 y2 -> (\l -> case l of {
                Nil -> n;
                Cons x xs -> c(f x)(map## xs)})
      (Cons z1 z2)
```

where the `zi` are special constants. This reduces to

$$\backslash y1\ y2 \rightarrow c(f\ z1)(map\#\# z2)$$

Now we use special rewrite rules generated from the type of the constructor to rewrite the dummy variables z_i to the real variables y_i . This makes it possible to discover the recursive invocation of the `map##` function and replace it by the induction variable. For the `Cons` constructor the rewrite rules $z1 \rightarrow y1$ and $map##\ z2 \rightarrow y2$ are generated. The first corresponds to an occurrence of the type parameter `a` and the second to a recursive occurrence of the type `List a`.

By application of the rewrite rules $z1 \rightarrow y1$ and $map##\ z2 \rightarrow y2$ the recursive call is recognized and we get

$$h2 = \backslash y1\ y2 \rightarrow c(f\ y1)(y2)$$

Putting this together gives the non-recursive definition

$$map\#\# = \backslash l \rightarrow cata[List\ a][t](n, \backslash y1\ y2 \rightarrow c(f\ y1)(y2))\ l$$

2.4.4. Folding

By unfolding the worker functions `map##` and `map#` back into their subsequent wrappers we obtain the build-cata form of `map`:

$$map = \backslash a\ b \rightarrow \backslash f\ l \rightarrow build[List\ b](\backslash t \rightarrow \backslash n\ c \rightarrow cata[List\ a][t](n, \backslash y1\ y2 \rightarrow c(f\ y1)(y2))\ l)$$

2.4.5. Transforming programs

The transformation procedure illustrated above is attempted (it may fail) for all functions. Compositions of functions in build-cata form can be deforested by unfolding their definitions and applying short cut fusion as part of standard simplification (see section 7). The unfolding can be done without risk of non-termination because the functions are not explicitly recursive.

The build-cata forms in figure 2 are all obtained using this transformation. Note that not all of these functions do both produce and consume a list; `foldr` only consumes a list and `upto` only produces a list. Their cata-and-or-build forms are obtained using variants of the transformation process described above. These variants are discussed in section 7 below.

We have specified the warm fusion transformation algorithm in Stratego. In the remainder of this paper we will give an overview of the specification. In particular, we will discuss the basic steps of the transformation such as splitting, unfolding, folding and deriving a catamorphism and how these can be used in various combinations and orders to obtain different results. First we give an overview of Stratego itself.

3. Stratego

In this section we briefly introduce System S, a calculus for the definition of tree transformations, and Stratego, a specification language providing syntactic abstractions

for System S expressions. For a detailed description of Stratego, its operational semantics, and additional examples of its use we refer the reader to [15,31,33–35,37]. Figure 3 shows a Stratego module defining several generic transformation operators. Other example specifications that use these operators will be discussed in the rest of the paper.

3.1. System S

System S is a hierarchy of operators for expressing term transformations. The first level provides control constructs for sequential non-deterministic programming, the second level introduces combinators for term traversal and the third level defines operators for binding variables and for matching and building terms.

Transformations in System S are applied to first-order terms, which are expressions over the grammar

$$t := x \mid C(t_1, \dots, t_n) \mid [t_1, \dots, t_n] \mid (t_1, \dots, t_n)$$

where x ranges over variables and C over constructors. The notation $[t_1, \dots, t_n]$ abbreviates the list $\text{Cons}(t_1, \dots, \text{Cons}(t_n, \text{Nil}))$. In addition, the notation $[t_1, \dots, t_n \mid t]$ denotes $\text{Cons}(t_1, \dots, \text{Cons}(t_n, t))$.

Level 1: Sequential non-deterministic programming. Strategies are programs that attempt to transform terms into terms, at which they may succeed or fail. In case of success the result of such an attempt is a transformed term. In case of failure there is no result of the transformation. Strategies can be combined into new strategies by means of the following operators. The *identity* strategy id leaves the subject term unchanged and always succeeds. The *failure* strategy fail always fails. The *sequential composition* $s_1 ; s_2$ of strategies s_1 and s_2 first attempts to apply s_1 to the subject term and, if that succeeds, applies s_2 to the result. The *non-deterministic choice* $s_1 + s_2$ of strategies s_1 and s_2 attempts to apply either s_1 or s_2 . It succeeds if either succeeds and it fails if both fail; the order in which s_1 and s_2 are tried is unspecified. The *deterministic choice* $s_1 <+ s_2$ of strategies s_1 and s_2 attempts to apply either s_1 or s_2 , in that order. The *recursive closure* $\text{rec } x(s)$ of the strategy s attempts to apply s to the entire subject term and the strategy $\text{rec } x(s)$ to each occurrence of the variable x in s . The *test* strategy $\text{test}(s)$ tries to apply the strategy s . It succeeds if s succeeds, and reverts the subject term to the original term. It also fails if s fails. The *negation* $\text{not}(s)$ succeeds (with the identity transformation) if s fails and fails if s succeeds. Two examples of strategies that can be defined with these operators are the try and repeat strategies in figure 3.

Level 2: Term traversal. The combinators discussed above combine strategies that apply transformations to the root of a term. In order to apply transformations throughout a term it is necessary to traverse it. For this purpose, System S provides a *congruence* operator $C(s_1, \dots, s_n)$ for each n -ary constructor C . It applies to terms of the form $C(t_1, \dots, t_n)$ and applies s_i to t_i . An example of the use of congruences is the operator $\text{map}(s)$ defined in figure 3 that applies a strategy s to each element of a list.

```

module traversals
imports lists
strategies
  try(s)          = s <+ id
  repeat(s)       = rec x(try(s; x))
  map(s)          = rec x(Nil + Cons(s, x))
  filter(s)       = rec x(Nil + Cons(s, x) <+ Tl; x)
  topdown(s)      = rec x(s; all(x))
  bottomup(s)     = rec x(all(x); s)
  downup(s)       = rec x(s; all(x); s)
  downup2(s1, s2) = rec x(s1; all(x); s2)
  alltd(s)        = rec x(s <+ all(x))
  oncetd(s)       = rec x(s <+ one(x))
  sometd(s)       = rec x(s <+ sone(x))
  manytd(s)       = rec x(s; all(try(x)) <+ some(x))
  onebu(s)        = rec x(one(x) <+ s)
  somebu(s)       = rec x(some(x) <+ s)

```

Figure 3. Specification of several generic term traversal strategies.

Congruences can be used to define traversals over specific data structures. Specification of generic traversals (e.g., pre- or post-order over arbitrary structures) requires more generic operators. The operator `all(s)` applies `s` to all children of a constructor application $C(t_1, \dots, t_n)$. In particular, `all(s)` is the identity on constants (constructor applications without children). The strategy `one(s)` applies `s` to one child of a constructor application $C(t_1, \dots, t_n)$; it is precisely the failure strategy on constants. The strategy `some(s)` applies `s` to some of the children of a constructor application $C(t_1, \dots, t_n)$, i.e., to at least one and as many as possible. Like `one(s)`, `some(s)` fails on constants.

Figure 3 defines various traversals based on these operators. For instance, `oncetd(s)` tries to find *one* application of `s` somewhere in the term starting at the root working its way down; `s <+ one(x)` first attempts to apply `s`, if that fails an application of `s` is (recursively) attempted at one of the children of the subject term. If no application is found the traversal fails. Compare this to the traversal `alltd(s)`, which finds *all* outermost applications of `s` and never fails.

Level 3: Match, build and variable binding. The operators we have introduced thus far are useful for repeatedly applying transformation rules throughout a term. Actual transformation rules are constructed by means of pattern matching and building of pattern instantiations.

A `match ?t` succeeds if the subject term matches with the term `t`. As a side-effect, any variables in `t` are bound to the corresponding subterms of the subject term. If a variable was already bound before the match, then the binding only succeeds if the terms are the same. This enables non-linear pattern matching, so that a match such as

$?F(x, x)$ succeeds only if the two arguments of F in the subject term are equal. This non-linear behavior can also arise across other operations. For example, the two consecutive matches $?F(x, y); ?F(y, x)$ succeed exactly when the two arguments of F are equal. Once a variable is bound it cannot be unbound.

A *build* $!t$ replaces the subject term with the instantiation of the pattern t using the current bindings of terms to variables in t . A scope $\{x_1, \dots, x_n: s\}$ makes the variables x_i local to the strategy s . This means that bindings to these variables outside the scope are undone when entering the scope and are restored after leaving it. The operation $where(s)$ applies the strategy s to the subject term. If successful, it restores the original subject term, keeping only the newly obtained bindings to variables.

Built-in data types. There are two predefined sorts with an infinite number of constructors: integers and strings. Several operators provide standard operations on these data types. Of particular importance for our purposes is the operator *new* that builds a new string that does not occur anywhere in the term being transformed.

3.2. Specifications

The specification language Stratego provides syntactic abstractions for System S expressions. A specification consists of a collection of modules that define signatures, transformation rules, and strategy definitions.

A signature declares the sorts and operations (constructors) that make up the structure of the language(s) being transformed. An example signature is shown in figure 4. A strategy definition $f(x_1, \dots, x_n) = s$ introduces a new strategy operator f parameterized with strategies x_1 through x_n and with body s . Such definitions cannot be recursive, i.e., they cannot refer (directly or indirectly) to the operator being defined. All recursion must be expressed explicitly by means of the recursion operator *rec*. Labeled transformation rules are abbreviations of a particular form of strategy definitions. A conditional rule $L : l \rightarrow r \text{ where } s$ with label L , left-hand side l , right-hand side r , and condition s denotes a strategy definition $L = \{x_1, \dots, x_n: ?l; \text{where}(s); !r\}$. Here, the body of the rule first matches the left-hand side l against the subject term, and then attempts to satisfy the condition s . If that succeeds, it builds the right-hand side r . The rule is enclosed in a scope that makes all term variables x_i occurring freely in l , s and r local to the rule. If more than one definition is provided with the same name, e.g., $f(xs) = s_1$ and $f(xs) = s_2$, this is equivalent to a single definition with the sum of the original bodies as body, i.e., $f(xs) = s_1 + s_2$.

Strategy operators can only have strategies as arguments. Data can be passed to strategy operators by wrapping them in *build* expressions. For instance, the strategy $\text{map}(!A)$ will replace every element of a list by the constant term A . Parameterized strategies have not often been used in previous Stratego specifications. They are nevertheless critical in specifying the warm fusion transformer and in other situations in which information must be passed between strategies.

```

module AHaskell
signature
  sorts Decl Constr Type Exp Alt
operations
  Program      : List(Decl)                -> Program

  Data         : Type * List(Constr) * Deriving -> Decl
  ConstrDecl  : Option(Forall) * Option(Context)
                * String * List(Type)       -> Constr

  SignDecl     : Vars * Type                -> Decl
  Valdef       : Exp * Exp                  -> Decl

  TCon         : String                      -> Type
  TVar         : String                      -> Type
  TApp         : Type * List(Type)           -> Type
  TFun         : List(Type) * Type           -> Type
  Forall       : List(String) * Type        -> Type

  Typed       : Exp * Type                  -> Exp
  Var         : String                      -> Exp
  Constr      : String                      -> Exp
  Lit         : Literal                     -> Exp
  Abs         : List(Exp) * Option(Type) * Exp -> Exp
  App         : Exp * List(Exp)              -> Exp
  Let         : List(Decl) * Exp             -> Exp
  Case       : Exp * List(Alt)              -> Exp
  Alt         : Exp * Option(Type) * Exp     -> Alt
  TAbs       : List(String) * Exp           -> Exp
  TInst      : Exp * List(Type)             -> Exp

  Build       : Type * Exp                  -> Exp
  Cata        : Type * Type * List(Exp)     -> Exp

```

Figure 4. Signature for kernel Haskell.

The following definitions provide a useful shorthand. The notation $\langle s \rangle t$ denotes $!t; s$, i.e., the strategy which builds the term t and then applies s to it. The notation $s \Rightarrow t$ denotes $s; ?t$, i.e., the strategy which applies s to the current subject term and then matches the result against t . The combined notation $\langle s \rangle t \Rightarrow t'$ thus denotes $!t; s; ?t'$. The $\langle s \rangle t$ notation can also be used inside a term in a build expression. For example, the strategy expression $!F(\langle s \rangle t, t')$ corresponds to $\{x: \langle s \rangle t \Rightarrow x; !F(x, t')\}$, where x is a new variable.

3.3. Derived idioms

Stratego's syntactic abstractions give rise to a number of useful programming idioms. Foremost among these are *recursive patterns* and *distributed patterns*.

Recursive patterns are strategy expressions that describe term formats by means of congruences and recursion. Nested congruences in Stratego are similar to pattern matching in functional languages, and Stratego’s recursive patterns involving nested patterns are akin to recursive functions which verify the structure of terms. Like pattern matching in functional languages, Stratego’s recursive patterns are completely general. For example, the following recursive pattern describes the subset of Haskell expressions that corresponds to untyped λ -calculus terms:

```
lambda-exp =
  rec x(Var(id) + App(x, x) + Abs([Var(id)], None, x))
```

Their use is further demonstrated in the term format checking in section 5. They can also be used to characterize more complicated formats such as normal forms or expressions in a core language. More generally, recursive patterns can be used whenever expressions in a sublanguage of a larger representation language must be recognized or manipulated.

Distributed patterns combine the pattern matching of recursive patterns with the traversal capabilities of strategy operators. They serve as “pattern templates” that can be used to match against expressions containing specified subexpressions at variable depths within them. For example, the warm fusion transformer uses the distributed pattern `underabstr` to determine whether or not a term in the expression language of figure 4 contains an application whose argument term is an abstraction in which the variable (determined by the strategy) `s` appears:

```
underabstr(s) = oncetd(App(id, Abs(id, id, oncetd(Var(s))))))
```

Note that the argument term to the abstraction need not actually *be* the variable determined by `s`; all that is required is that the variable appear somewhere within the argument term. More general distributed patterns are constructed with the same ease.

3.4. Implementation

The Stratego compiler translates a specification to a C program that reads a term and applies the specified transformation to it. The compiler first translates a specification to a System S expression, which is then translated to a list of abstract machine instructions. The instructions are implemented in C. The run-time system is based on the ATerm library [28], which supports complete sharing of subterms (hash-consing). ATerms are also used for exchange of data between components of a transformation systems. The compiler is bootstrapped, i.e., implemented in Stratego itself. The Stratego library [32] provides a large of number generic, language independent rules and strategies.

4. Architecture

The architecture of the warm fusion program transformation system is depicted in figure 5. The system consists of four main components: a parser, typechecker, the actual warm fusion transformer, and a pretty-printer. The system could have been defined as

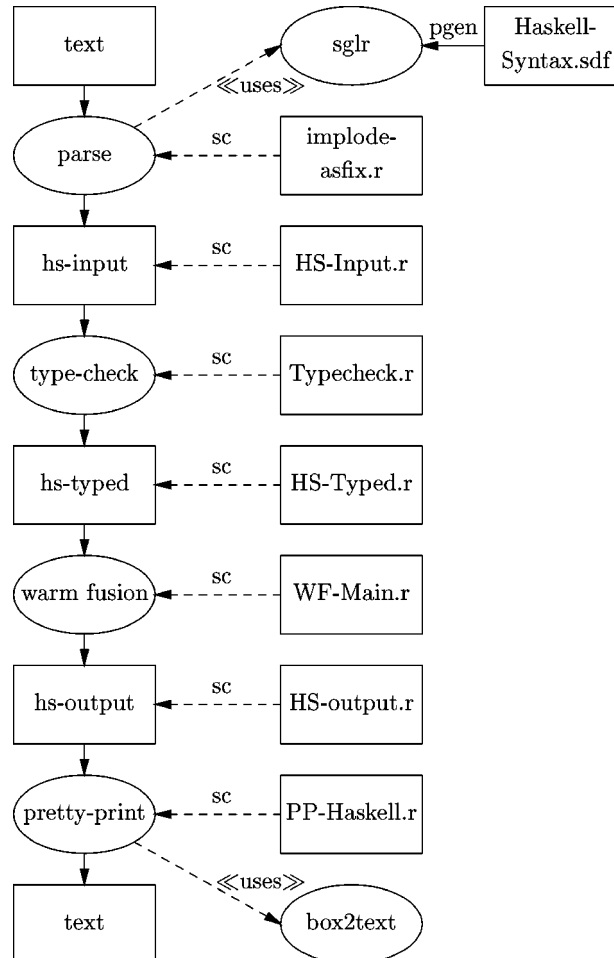


Figure 5. Architecture of the warm fusion transformation tool. Boxes represent data, ellipses represent components. Dashed arrows represent generation of components from specifications via the Stratego compiler (sc), the SDF2 parser generator (pgen) and a C compiler (gcc). The intermediate data-formats are also described in Stratego and format checkers are generated from their specifications.

a single component, but dividing it into separate components encourages separation of concerns during development and makes future application of the transformation tool in another setting – e.g., connection to a compiler front-end – easier.

The parser is generated from a specification of the full¹ Haskell98 syntax [21] in the syntax definition formalism SDF2 [30]. Although the parser supports the full syntax, currently only the kernel subset of Haskell is supported by the subsequent components. A Haskell desugaring component can be added in the future to extend the transformer to full Haskell.

¹ The syntax definition is complete up to layout.

Note that SDF2 based parsers are not required for Stratego. Parsing front-ends can also be written using YACC or any other parser generator, as long as the generated parsers output abstract syntax trees in the ATerm format. The SDF2 parser that we use actually outputs parse trees. These are transformed to abstract syntax trees by a generic – i.e., grammar independent – tool (*implode-asfix*) written in Stratego.

The current typechecker is basically a preprocessor that distributes type information from signature declarations to variable uses. This could be enhanced to a tool that does full type inference, but for the purposes of our case study this was not necessary; types of variables are declared explicitly in input programs. Note that this is not too much of a restriction. In Haskell it is customary to declare the types of functions anyway.

The intermediate data structures that are exchanged between components are represented in the generic ATerm format [28]. Furthermore, each component consumes and produces a different subset of the general abstract syntax of the language. These formats are also described in Stratego by means of strategies that check the structure of a term. These strategies can be used by components to verify their input.

The warm fusion transformer processes each of the function definitions in a program and tries to transform it into build-cata form. It also inlines previously transformed functions in the definitions it is processing to achieve deforestation by the short cut.

The pretty-printer is a formatter that translates abstract syntax to strings. A Stratego specification (PP-Haskell) defines the translation from abstract syntax to Box terms. These are translated to formatted text by a generic Box formatter [4,29].

In the next sections we will discuss various aspects of the specification of the warm fusion transformation system. In section 5 we discuss the specification of the abstract syntax, checking subsets of an abstract syntax, and the specification of bound variable renaming and substitution by instantiating generic language independent algorithms. In section 6 we present the overall structure of the transformer. In section 7 we discuss the details of some of the transformations.

5. Abstract syntax

The warm fusion transformation is performed on the abstract syntax of kernel Haskell, or `AHaskell`. The signature of the language is shown in figure 4. It is a standard functional language with abstraction, application, data type deconstruction by means of case expressions, and a recursive let binding. The language is explicitly typed, which entails that types of variables in bindings can be declared, and that atomic expressions (variables, constructors and constants) can be annotated with their types. Polymorphic expressions are constructed by means of type abstraction and instantiated by means of type application. A program consists of a list of type and function definitions.

5.1. Format checking

In the course of the transformation we encounter three intermediate formats that are subsets of `AHaskell` (figure 4). The input format `hs-input` allows atomic expressions without type annotations because requiring annotations would clutter the source code. It also allows infix operators as syntactic sugar for prefix application. In the intermediate format `hs-typed` all atomic expressions are annotated with their types and are type correct. In addition, all operators are in prefix form. Like `hs-typed`, the output format `hs-output` requires fully annotated atoms, but it also allows expressions constructed using the `Build` and `Cata` operators. The latter are not allowed in the input to the transformation.

These three expression formats could be described by introducing three separate signatures with different constructors. This would, however, require three sets of names for the same constructs and trivial translations from one set to the next. Instead, we use one signature and the recursive patterns of section 3 to characterize the three restrictions. These recursive patterns document the formats and can be used to check the inputs to the transformation components.

We now consider in turn the forms of expressions in each of the three subformats of `AHaskell`. Atomic expressions in the `hs-typed` format consist of a variable, constructor or literal and a type annotation as described by the patterns

```
AExp = Var(id) + Constr(id) + Lit(id)
atom(t) = Typed(AExp, t)
TypedVar = Typed(Var(id), Type)
TypedAtom = atom(Type)
```

where `Type` is a recursive pattern which describes the structure of `AHaskell`'s types. Type annotations are represented by means of the constructor `Typed`, which represents the `e :: t` notation in Haskell. Note that these patterns are parameterized with the format for types `t`. The basic shape of a `hs-typed` expression is described by the patterns

```
exp(e, t, pat, var) =
  Abs(list(var), option(t), e) +
  Case(e, list(alt(e, t, pat))) +
  Let(list(decl(e, t)), e) +
  App(e, list(e)) +
  TAbs(list(TVar(id)), e) +
  TInst(e, list(t))

alt(e, t, pat) =
  Alt(pat, option(t), e)

simple-pattern(var) =
  Constr(id) +
  App(Constr(id), list(var))
```

```
TypedPat =
  simple-pattern(TypedVar)
```

and a typed expression is characterized by the recursive pattern

```
TypedExp =
  rec e(TypedAtom + exp(e, Type, TypedPat, TypedVar))
```

In the `hs-input` format, atomic expressions (variables, constructors and literals) can be untyped. Furthermore, infix operator applications in addition to prefix application and binary in addition to n -ary application are allowed. This is described by

```
PreVar =
  Var(id) +
  Typed(Var(id), PreType)
```

```
PrePat =
  simple-pattern(PreVar)
  + rec x(AppBin(x, PreVar) + Constr(id))
```

```
pre-exp(e) =
  OpApp(e, id, e) +
  AppBin(e, e) +
  Negation(e) +
  If(e, e, e)
```

```
PreExp =
  rec e(AExp + atom(PreType) + pre-exp(e) +
        exp(e, PreType, PrePat, PreVar))
```

The typechecker normalizes infix and binary applications to n -ary applications and annotates all atomic expressions with their types.

Finally, the expressions in the output format `hs-output` are typed expressions extended with `Build` and `Cata` operators:

```
ext-exp(e, t) =
  Cata(t, t, list(e)) +
  Build(t, e)
```

```
ExtExp =
  rec e(TypedAtom + exp(e, Type, TypedPat, TypedVar) +
        ext-exp(e, Type))
```

5.2. Variable renaming and substitution

AHaskell has variable binding constructs. The Stratego library defines (using standard Stratego) the generic, language independent strategies `rename` for renaming bound variables, `substitute` for parallel substitution of expressions for variables,

and `free-vars` for the extraction of the free variables from an expression. These operations are instantiated by declaring the shape of variables, indicating the binding constructs, and identifying the binding positions. We illustrate their instantiation for `AHaskell`. The implementation of the generic algorithms is presented in [33].

The following rules are used to describe the shape of variables.

```

IsVar(s)   : Var(x) -> Var(<s> Var(x))
ExpVar     : Typed(Var(x),_) -> Var(x)
ExpVar     : Var(x) -> Var(x)
ExpVars    : Var(x) -> [Var(x)]

```

The binding constructs of expressions are lambda abstraction, case alternatives, and let binding. The rules `ExpBnd` define the projection from these constructs to the list of variables that they bind.

```

ExpBnd     : Abs(xs, _, _) -> <map(ExpVar)> xs
ExpBnd     : Alt(App(c, xs), t, e) -> <map(ExpVar)> xs
ExpBnd     : Let(decls, e) -> <filter(DeclVar)> decls
DeclVar    : Valdef(Var(x), e) -> Var(x)

```

Using the rules above the instantiations of `free-vars`, `substitute`, and `rename` for expressions are

```

expvars    = free-vars(ExpVars, ExpBnd)
exprename  = rename(IsVar, ExpBnd)
expsubst   = substitute(Typed(Var(id),id) + Var(id), etrename)

```

Proper substitution entails that bound type variables in expressions that are substituted for term variables are also renamed, and so an exercise similar to that above must be carried out for type variables. This gives rise to the corresponding operators `tpvars`, `tpsubst`, and `tprename` for types. The strategy `etrename` is the sequential composition of `exprename` and `tprename`.

6. Transformer: Big picture

In this section we discuss the specification of the top-level of the warm fusion transformer. The reader is directed to [13], from which the following code is excerpted, for a complete code listing.

6.1. Transforming a program

The main strategy takes a program, i.e., a list of type and function declarations, and transforms each in turn. This is achieved by a transition step for each declaration:

```

Main = etrename;
      where (collect-data-defs);
            InitWF;

```

```
repeat(TransformDecl <+ NormD);
ExitWF
```

Note that all bound variables in the entire program are first renamed to establish the unique variable invariant. Furthermore, the strategy `collect-data-defs` finds the data type definitions in the program and stores them in a symbol table for later reference. The initial configuration is created from a list of declarations and the final configuration derives a transformed list of declarations:

```
InitWF :
  ds -> ([], [], ds)
ExitWF :
  (ds1, ds2, []) -> <reverse> ds2
```

The first accumulator list stores the functions that have been transformed to build-cata form. These are used for inlining in other functions. The second accumulator list stores all functions, including the non-transformed ones.

A definition is transformed by first inlining functions that were transformed earlier (in the list `ds1`) and then applying the warm fusion transformation to it.

```
TransformDecl :
  (ds1, ds2, [d | ds3]) -> ([d' | ds1], [d' | ds2], ds3)
  where <ior(inline(!ds1), Transform)> d => d'
```

Inlining and transformation can fail. If at least one succeeds then the result is considered to be transformed and is added to both accumulator lists. (The rule `ior` computes the inclusive or of two strategies, i.e., `ior(s1, s2)` applies `s1`, `s2` or both.) If both fail then the function is added only to the list of non-transformed functions using the rule

```
NormD :
  (ds1, ds2, [d | ds3]) -> (ds1, [d | ds2], ds3)
```

Inlining is achieved by replacing calls to functions in a given list of declarations by (renamings of) their bodies and then simplifying the resulting expressions using the rules of section 7. Inlining replaces as many calls as possible, but at least one call must be replaced in order for it to succeed:

```
inline(mkenv) = manytd(Inline(mkenv)); simplify
```

The function to be inlined is looked up in the list of declarations passed to the rule `Inline`. The strategy `<not(in)>` checks for recursion in the definition of the function. Recursive functions are not inlined.

```
Inline(mkenv) :
  Typed(Var(x), t) -> <tpsubst; etrename> (sbs, e)
  where mkenv; fetch(?Valdef(Var(x), e)); <not(in)> (Var(x), e);
  <tpunify> [(<type> e, t)] => sbs
```

6.2. Transforming a definition

The basic algorithm for transforming a recursive definition to build-cata form – as defined in [14] and illustrated in section 2 – is the following:

```

Transform =
  IntroBuildCata;
  simplify;
  SplitBodyCP;
  Unfoldlin2;
  [id, simplify;
   MakeCataBody];
  Unfold2in1;
  simplify

```

This strategy introduces the build-cata identity, splits the body into a wrapper and a worker, unfolds the wrapper in the worker, transforms the worker into a catamorphism, and unfolds the worker back in the wrapper. In between it simplifies the definitions.

As we remarked in section 2 this procedure applies only to functions that both consume and produce data structures. To accommodate functions that either only consume or only produce data structures we refine the algorithm using the same building blocks to the following:

```

Transform =
  ((IntroBuildCata;
   simplify;
   (ConsumerProducer
    <+ Producer
    <+ NonRecursiveProducer))
   <+ Consumer);
  simplify

```

The strategies `ConsumerProducer`, `Consumer`, `Producer`, and `NonRecursiveProducer` represent the different possible ways of transforming a function. The strategy `Consumer` is applied when introduction of the outer `build` and `cata` fails. In this case the output type of the function is not a data type and so the function does not produce a data structure. It may, however, still be a consumer. If, on the other hand, the introduction of the outer `build` and `cata` succeeds, then `ConsumerProducer` splits the body of the function into a wrapper and a worker and tries to derive a catamorphism for the worker. If deriving a catamorphism from the worker fails, then the function is only a producer.

Although it is not apparent at this level of abstraction, the introduction of the outer `build` and `cata` is governed by the input and output types of the function being transformed. We consider the details of the above transformation in section 7.

The derivation of a catamorphism for the worker and unfolding it back in the wrapper is defined in the strategy `BodyToCata`:


```

BodyToCata =
  UnfoldLin2;
  [id, simplify;
   SplitBodyP;
   UnfoldLin2;
   [id, simplify;
    MakeCataBody];
   Unfold2in1];
  Unfold2in1

```

Unlike `Transform`, this strategy splits and unfolds the worker twice in order to recognize the abstracted constructors as static parameters.

7. Transformer: Details

In this section we go into the details of some of the transformations mentioned above.

7.1. Simplification

The simplifier consists of a number of standard simplification rules for functional programs such as beta reduction:

```

BetaOne :
  App(Abs([x|xs], t, e), [a|as]) ->
  App(Abs(xs, t, <expsubst> ([x], [a], e)), as)
  where <value> a + <linear> (x, e)

```

Here, `value` and `linear` are strategies that prevent duplication of work during reduction. An expression is a value if it represents either a function or a data object; a variable `v` appears linearly in the expression `b` if reduction of `b` can never cause duplication of any term substituted for `v`. Terms which do not encode computation are literally copied regardless of whether or not the variables they instantiate occur linearly in their host terms.

The beta reduction rule `BetaOne` reduces an application of a function to its first argument. The following rule reduces such an application as far as possible, either exhausting all formal or all actual parameters.

```

Beta :
  App(Abs(xs, t, e), as) ->
  App(Abs(ys, t, <expsubst> (sbs, e)), bs)
  where <rest-zip(id)> (xs, as) => (ys, bs, sbs);
        (<lzip((id,value) + (fst,id); linear)> (sbs, e))

```

Other simplification rules include elimination of dead `let` bindings, inlining of `let` bindings, case specialization, distribution of application over cases, uncurrying of

expression and type applications; see the definition of `basic-rules` below. A particularly important rule for the warm fusion transformation is, of course, `cata-build` fusion:

```
CataBuild :
  App(Cata(t1, t2, fs), [Build(t1, g)]) ->
  App(TInst(g, [t2]), fs)
```

Here `t1` is the input type for the catamorphism and `t2` is its return type. Similarly, `t1` is the type of `build`'s output.

These basic rules can be combined in various ways to build simplifiers, depending on the desired effect. We use the following configuration in the warm fusion transformer.

```
basic_rules =
  Beta + Eta + (Inl; Dead) + TEta + TBeta +
  CaseConstr + CaseDistL + CaseDistR + Uncurry

basic-cata = CataConstr + CataBuild + basic_rules

simplify = innermost(basic-cata)
```

The strategy `innermost` is defined by

```
innermost(s) = rec x(all(x); (s; x <+ id))
```

Although the definition of `simplify` here uses `innermost` reduction, Stratego's separation of logic from control make it particularly convenient to change the term reduction strategy used in the simplifier.

7.2. Build-cata introduction

The initial `build-cata` identity is introduced into the body of the function definition under its leading abstractions:

```
IntroBuildCata = Valdef(id, under-abs(MkBuildCata))
```

where the notion "under its leading abstractions" can be expressed by the recursive pattern

```
under-abs(s) = rec x((Abs(id, id, x) + TAbs(id, x)) <+ s)
```

In concrete syntax the `build-cata` identity has the form

```
build[t1](/\t2 -> \fs :: t2 -> (cata[t1][t2](fs) e))
```

where `t1` is the type of the expression `e`, `t2` is a new type variable and the `fs` are the abstract constructors corresponding to the constructors of the data type. Generation of this form is defined by the following rule:

```
MkBuildCata :
  e -> Build(t1, TAbs([t2], Abs(fs, Some(t2),
  App(Cata(t1, t2, fs), [e])))
```

```

where new-tvar => t2; <type> e => t1;
  <get-constructors> t1 => cdecls;
  <lzip(AbsConstr)> (cdecls, (t1, t2)) => fs

```

Type information plays a crucial role in `build-cata` introduction and subsequent processing. It is used not only to determine which instances of the `Cata` and `Build` functions to introduce, but also to generate arguments of the appropriate types for these instances. The strategy `type` derives the type from an expression. The strategy `get-constructors` obtains the constructor declarations corresponding to the type of `e`. For each constructor of the data type an abstract constructor (variable) with the appropriate type is constructed by rule `AbsConstr`:

```

AbsConstr :
  (ConstrDecl(_, _, c, ts), (t1, t2)) ->
  Typed(Var(f), TFun(ts', t2))
  where new => f; <map(try(?t1;!t2))> ts => ts'

```

The rule creates a variable expression with new variable `f` and its type. The function has the same number of arguments as the original constructor. The output of the function is of type `t2`. Where the constructor has a recursive argument, indicated by the recursion type `t1`, the output type `t2` is instantiated. The other arguments remain the same type.

7.3. Splitting function definitions

Splitting a function into a wrapper and a worker involves determining where in the body the split is performed, which variables the worker is abstracted over, creating the definition of the worker and replacing the expression in the wrapper body by a call to the worker. There are several ways to do this. We discuss one of them.

The strategy `SplitBodyP` first computes the non-static parameters `vs` of the function definition and then splits the body. This is achieved by instantiating `SplitBody` with a strategy for splitting expressions:

```

SplitBodyP =
  where(NonStaticParams => vs);
  SplitBody(SplitExpr(!vs))

```

`NonStaticParams` extracts the nonstatic parameters from a function definition; the function's case selector must be the head of the list of nonstatic parameters in order to satisfy the strictness requirement of the promotion theorem. Given any list `xs` of value and type variables, the rule `SplitExpr` creates a definition for a function with a new name `f` that has the expression as its body and abstracts over `xs`. It also creates a call to `f` with `xs` as arguments. The definition of `SplitExpr` assumes that the type parameters to a function are always static.

```

SplitExpr(mkxs) :
  e -> (App(Typed(Var(f), t), xs), Valdef(Var(f), body))
  where mkxs => xs; new => f;

```

```
<etrename> Abs(xs, Some(<type> e), e) => body;
<type> body => t
```

Given a strategy `split` for splitting an expression, rule `SplitBody` splits the body of a function definition by creeping under its leading abstractions and splitting the expression it encounters there.

```
SplitBody(split) :
  Valdef(Var(x), body) -> [Valdef(Var(x), body'), def]
  where <under-abs-build(split => (e, def); !e)> body => body'
```

The `split` results in an expression (the call) and a new definition. The expression `split => (e, def); !e` matches the result of splitting against the pattern `(e, def)` and then replaces it by just the expression. The binding to `def` is used in the right-hand side of the rule, where a list of two definitions is created.

Since we want to split off the worker under the build expression, if present, we use a variant of the `under-abs` pattern that we saw before.

```
under-abs-build(split) =
  rec x((Abs(id,id,x) + TAbs(id,x) + Build(id,split)) <+ split)
```

Similar patterns can be used to describe other contexts in which a transformation has to take place.

Parameterizing over `under-abs-build` as well as `split` would make `SplitBody` a completely generic splitting strategy. However, even as defined here, `SplitBody` is a general strategy for splitting under any type and term abstractions and any builds in a function definition. Our splitting mechanism therefore generalizes that from [22] upon which the wrapper-worker decomposition in [14] is based. The extra generality is useful: splitting a function definition into a wrapper and a worker sometimes requires splitting under a function's leading build, while at other times no builds are present. The strategy `under-abs-build` given here is general enough to accommodate both situations.

7.4. Unfolding

Unfolding is defined by the following contextual rules [35] that replace all occurrences of atoms with the name of the function being unfolded by its body.

```
Unfoldlin2 :
  [Valdef(Var(x),body1), Valdef(Var(y),body2[Typed(Var(x),_)])]
-> [Valdef(Var(x),body1), Valdef(Var(y),body2[body1'])(alltd)]
  where <exprename> body1 => body1'
```

```
Unfold2in1 :
  [Valdef(Var(x),body1[Typed(Var(y),_)]), Valdef(Var(y),body2)]
-> Valdef(Var(x),body1[body2'])(alltd)
  where <not(in)> (Var(y), body2); <exprename> body2 => body2'
```

7.5. Cata promotion

In section 2 we discussed how a catamorphism can be derived from a recursive definition using the promotion theorem. The core of the promotion is the creation of a function

$$h = \lambda z_1 \dots z_n \rightarrow e(c(y_1) \dots (y_n))$$

for each constructor c with n arguments. The function e is then unfolded exactly once, and the result is simplified using the standard rules, together with a dynamically generated set of rules that rewrite recursive applications involving the y_i s to the appropriate variables z_i . The abstract syntax of the initial form of the function h is

```
Abs(zs, App(e, [App(Typed(Constr(c), TFunc(ts, t)), ys)]))
```

The rule `DynRules` creates for a specific constructor, the lists of y and z variables and the corresponding dynamic rewrite rules. The strategy `dsimplify` extends the normal simplification with the application of these dynamic rules.

```
dsimplify(mkrls) = innermost(AppDynRule(mkrls) <+ basic_rules)
```

Putting this together the rule `MkH` creates the replacement function corresponding to a constructor of the original function's input data type.

```
MkH :
  (ConstrDecl(_, _, c, ts), (g, e, t)) -> h
  where
    <DynRules> (t, g, c) => (ys, zs, rls);
    !Typed(Constr(c), TFunc(ts, t)) => ct;
    <dsimplify(!rls)>
      Abs(zs, None, App(<etrename> e, [App(ct, ys)])) => h;
    <not(once t d({y : ?Var(y);
                  where(<fetch(Typed(Var(?y), id))> ys)}))> h
```

Note that the bound variables in expression e are renamed to maintain the unique variable invariant.

These replacement functions are then used by `MakeCataBody` to construct the catamorphic version of that function's worker. Unfolding the worker in the wrapper yields the build-cata form of the function definition being transformed.

```
MakeCataBody :
  Valdef(Var(g), e) -> Valdef(Var(g), Cata(t1, t2, hs))
  where <type> e => tg;
        <split(dom, range)> tg => (t1, t2);
        <get-constructors> t1 => cdecls;
        <lzip(MkH)> (cdecls, (Typed(Var(g), tg), e, t1)) => hs
```

This concludes our sample of the specification. The complete text of the specification can be found in [13].

8. Related work

The first ideas for rewriting strategy operators with general traversal operators are described in [15]. In [35] these ideas are formalized by means of an operational semantics and are extended to the full set of System S operators by splitting simple rewrite rules into match, build and scope. This allows easy expression of contextual rules. An application to the specification of optimizers is discussed. In [34] it is shown how System S can be used to describe various features and evaluation strategies of traditional conditional rewriting systems. In [31] three programming idioms for *strategic pattern matching* are studied: recursive patterns, contextual rules, and overlays. The implementation of generic algorithms such as used for variable renaming and substitution is discussed in [33]. For a discussion of related work on rewriting strategies see [34]. The relation to other systems for program transformation is discussed in [35].

Techniques for program fusion can be classified into two broad categories: *search-based* and *calculation-based*. The earliest techniques for program fusion [2,3,27,36] were search-based. These rely on analyses of the *fold-unfold* transformation process of Burstall and Darlington to fuse compositions of recursive functions. In search-based fusion it is necessary to keep track at each step of the transformation process of all function calls that have been made. New function definitions to be used in unfolding must then be introduced. Search-based fusion is systematic, but relies on clever control mechanisms to avoid the possibility of infinite sequences of transformations by repeated unfolding of function definitions. As a result, good implementations of search-based fusion techniques have been somewhat difficult to achieve.

The warm fusion method and the short cut to deforestation which it facilitates are in the more recent tradition of calculation-based fusion [8,10,14,24,25]. In calculation-based fusion the recursive structure of each component participating in the fusion is made explicit. This enables fusion by direct application of simple transformation laws like the cata-build rule and the acid rain theorem [25]. The theoretical basis for calculation-based fusion lies in the study of *constructive algorithmics* [5,17,18].

9. Future work

The implementation of program fusion algorithms offers many additional opportunities for investigation. Among the issues pertaining directly to the Stratego implementation and meriting attention are: experimenting with various orders and strategies for applying the simplification transformations; experimenting with more unfolding of function definitions when converting recursion to catamorphisms via fold promotion so that fusion is not unnecessarily blocked; making inlining more context sensitive, so that build-cata forms are inlined only when there is the possibility of fusion via the short cut; and extending the transformations with Gill's *augment*. Benchmarking to determine the sense(s) in which deforested programs are "better" than their monolithic counterparts is also appropriate for the current warm fusion implementation. So is comparison of the Stratego specification with other implementations of warm fusion.

Other lines of inquiry involve the integration of automatic fusion tools into existing systems. Candidate systems include the optimizer of the RML compiler discussed in [26, 35], as well as state-of-the-art functional language compilers. Nemeth [19] has recently implemented warm fusion in the Glasgow Haskell Compiler and reported benchmarks on programs from the nofib suite [20].

Finally, rather than using Stratego as a tool to help deepen our understanding of program fusion techniques, we can turn the relationship between strategy-based languages and program fusion on its head and ask about possible applications of fusion to strategy-based languages. Can we formalize our intuition that certain combinations of strategies should themselves be amenable to suitable forms of strategy fusion? Is it possible, for example, to make precise the observation that

$$!C(t_1, \dots, t_n); ?C(t_1', \dots, t_n') = !t_1; ?t_1'; \dots; !t_n; ?t_n'$$

assuming that the term that is built is not used again?

10. Conclusion

We have presented a case study of the application of Stratego to build a complete, non-trivial program transformation system. Table 1 shows the sizes of the main components of the transformation system in number of modules, lines of code (text including comments), number of rules and number of strategies. Note that these figures do not include the signature and the pretty-printing modules. Distributed over time, it took us about 30 days to develop the entire transformation tool from scratch including a syntax definition for full Haskell. The development time included finding out how to program in Stratego and developing programming idioms. That is, when undertak-

Table 1
Size metrics of main components of the specification. Measuring number of modules (mod), lines of code (LOC) including documentation, number of constructors (cons), rules and strategies (strat).

Language	Component	mod	LOC	cons	rules	strat
SDF	Haskell.sdf		650		300	–
Stratego	Warm fusion	11	739	0	60	31
Stratego	Format checking	1	202	1	1	32
Stratego	Haskell library	4	246	0	32	21
Stratego	Haskell normalize	1	75	0	17	3
Stratego	Haskell typecheck	1	120	1	13	6
Stratego	Subtotal specification	18	1382	2	123	93
Stratego	Signature	28	544	103	0	0
Stratego	Pretty-printer	28	671	0	90	7
Stratego	Total specification	74	2597	105	213	100
Stratego	Stratego library	48	3634	65	131	317

ing this case study, Stratego was a new language, even for its author, and discovering idioms of use beyond the basic paradigm takes time. The development was aided by the wealth of generic, language independent rules and strategies in the Stratego library [32].

This case study strengthens our view that rewriting strategies are a good paradigm for the implementation of program transformation systems. The specification is highly modular at all levels and can easily be modified or extended with new transformations. It will serve as the basic infrastructure for further experimentation with transformations on full Haskell. The specification also provides examples of several Stratego idioms that can be used in the implementation of transformation systems for other languages. In particular the specification shows the use of compound rules, recursive patterns, distributed patterns, exchange of information between transformation rules through parameterized strategies and the compact specification of variable renaming, substitution, and free variable projection.

Acknowledgements

The authors would like to thank various anonymous referees for their comments on earlier versions of this paper.

References

- [1] A.W. Appel, *Compiling with Continuations* (Cambridge University Press, 1992).
- [2] R.M. Burstall and J. Darlington, A transformational system for developing recursive programs, *J. of the ACM* 24(1) (1977) 44–67.
- [3] W.N. Chin, Safe fusion of functional expressions, *ACM Lisp Pointers* 5(1) (1992) 11–20, Proceedings of the 1992 ACM Conference on Lisp and Functional Programming.
- [4] M. de Jonge, A pretty-printer for every occasion, in: *Proceedings of the 2nd International Symposium on Constructing Software Engineering Tools (CoSET2000)*, eds. I. Ferguson, J. Gray and L. Scott, Limerick, Ireland (June 2000); also: Technical Report, University of Wollongong, Australia.
- [5] M. Fokkinga, Law and order in algorithmics, Ph.D. thesis, Twente University (1992).
- [6] P. Fradet and D. Le Métayer, Compilation of functional languages by program transformation, *ACM Trans. Programming Languages Syst.* 13(1) (1991) 21–51.
- [7] A.J. Gill, Cheap deforestation for non-strict functional languages, Ph.D. thesis, University of Glasgow (January 1996).
- [8] A. Gill, J. Launchbury and S.L. Peyton Jones, A short cut to deforestation, in: *Functional Programming Languages and Computer Architecture (FPCA'93)*, ed. Arvind (ACM Press, 1993) pp. 223–232.
- [9] T. Hagino, A categorical programming language, Ph.D. thesis, University of Edinburgh (1987).
- [10] Z. Hu, H. Iwasaki and M. Takeichi, Deriving structural hylomorphisms from recursive definitions, *ACM SIGPLAN Notices* 31(6) (1996) 73–82, Proceedings of the International Conference on Functional Programming (ICFP'96), Philadelphia.
- [11] P. Johann, An implementation of warm fusion, available at <ftp://ftp.cse.ogi.edu/pub/pacsoft/wf/> (1997).
- [12] P. Johann, Short cut fusion is correct, Technical report, Bates College, Lewiston, Maine, USA (2000).

- [13] P. Johann and E. Visser, Warm fusion in Stratego: A case study in the generation of program transformation systems, Technical report, Institute of Information and Computing Sciences, Universiteit Utrecht, Utrecht, The Netherlands (2000).
- [14] J. Launchbury and T. Sheard, Warm fusion: Deriving build-catas from recursive definitions, in: *Functional Programming Languages and Computer Architecture (FPCA'95)*, ed. S.L. Peyton Jones (ACM Press, 1995) pp. 314–323.
- [15] B. Luttik and E. Visser, Specification of rewriting strategies, in: *2nd International Workshop on the Theory and Practice of Algebraic Specifications (ASF+SDF'97)*, ed. M.P.A. Sellink, Electronic Workshops in Computing, Berlin, November 1997 (Springer, Berlin, 1997).
- [16] G. Malcolm, Homomorphisms and promotability, in: *Mathematics of Program Construction*, Lecture Notes in Computer Science, Vol. 375 (Springer, Berlin, 1989) pp. 335–347.
- [17] G.J. Malcolm, Data structures and program transformation, *Sci. Comput. Programming* 14 (1990) 255–279.
- [18] E. Meijer, M. Fokkinga and R. Paterson, Functional programming with bananas, lenses, envelopes and barbed wire, in: *Functional Programming and Computer Architecture (FPCA'91)*, ed. R.J.M. Hughes, Lecture Notes in Computer Science, Vol. 523 (Springer, Berlin, 1991) pp. 124–144.
- [19] L. Németh, Catamorphism based program transformation for non-strict functional languages, Draft, Ph.D. thesis, University of Glasgow (2000).
- [20] W. Partain, The `nofib` benchmark suite of Haskell programs, in: *Functional Programming*, eds. J. Launchbury and P.M. Sansom (Springer, Berlin, 1992) pp. 195–202.
- [21] S. Peyton Jones, J. Hughes et al., Report of the programming language Haskell98 a non-strict, purely functional language (February 1999).
- [22] S.L. Peyton Jones and J. Launchbury, Unboxed values as first class citizens in a non-strict functional language, in: *Functional Programming and Computer Architecture (FPCA'91)*, ed. R.J.M. Hughes, Lecture Notes in Computer Science, Vol. 523 (Springer, Berlin, 1991) pp. 636–666.
- [23] S.L. Peyton Jones and A.L.M. Santos, A transformation-based optimiser for Haskell, *Sci. Comput. Programming* 32(1–3) (1998) 3–47.
- [24] T. Sheard and L. Fegaras, A fold for all seasons, in: *Functional Programming and Computer Architecture (FPCA'93)*, ed. Arvind, Copenhagen, Denmark (ACM Press, 1993) pp. 233–242.
- [25] A. Takano and E. Meijer, Shortcut deforestation in calculational form, in: *Functional Programming and Computer Architecture (FPCA'95)*, ed. S.L. Peyton-Jones, San Diego, CA (June 1995).
- [26] A. Tolmach and D. Oliva, From ML to Ada: Strongly-typed language interoperability via source translation, *J. Funct. Programming* 8(4) (1998) 367–412.
- [27] V. Turchin, The concept of a supercompiler, *ACM Trans. Programming Languages Syst.* 8(3) (1986) 292–326.
- [28] M.G.J. van den Brand, H.A. de Jong, P. Klint and P.A. Olivier, Efficient annotated terms, *Software Practice Exper.* 30 (2000) 259–291.
- [29] M.G.J. van den Brand and E. Visser, Generation of formatters for context-free languages, *ACM Trans. Software Engrg. Methodol.* 5(1) (1996) 1–41.
- [30] E. Visser, Syntax definition for language prototyping, Ph.D. thesis, University of Amsterdam (September 1997).
- [31] E. Visser, Strategic pattern matching, in: *Rewriting Techniques and Applications (RTA'99)*, eds. P. Narendran and M. Rusinowitch, Lecture Notes in Computer Science, Vol. 1631, Trento, Italy, July 1999 (Springer, Berlin, 1999) pp. 30–44.
- [32] E. Visser, *The Stratego Library*, Institute of Information and Computing Sciences, Universiteit Utrecht, Utrecht, The Netherlands (1999).
- [33] E. Visser, Language independent traversals for program transformation, in: *Workshop on Generic Programming (WGP2000)*, ed. J. Jeuring, Ponte de Lima, Portugal, July 6 (2000); also: Technical Report UU-CS-2000-19, Universiteit Utrecht.
- [34] E. Visser and Zine-el-Abidine Benaissa, A core language for rewriting, in: *Proceedings of the Second International Workshop on Rewriting Logic and its Applications (WRLA'98)*, eds. C. Kirchner and H. Kirchner, Electronic Notes in Theoretical Computer Science, Vol. 15 (1998).

- [35] E. Visser, Zine-el-Abidine Benaissa and A. Tolmach, Building program optimizers with rewriting strategies, *ACM SIGPLAN Notices* 34(1) (1999) 13–26, Proceedings of the International Conference on Functional Programming (ICFP'98).
- [36] P. Wadler, Deforestation: Transforming programs to eliminate trees, *Theor. Comput. Sci.* 73 (1990) 231–248.
- [37] www.stratego-language.org.