


# Using Transformations in the Implementation of Higher-order Functions

Hanne Riis Nielson  
Flemming Nielson

DAIMI PB – 337  
October 1990

COMPUTER SCIENCE DEPARTMENT  
AARHUS UNIVERSITY  
Ny Munkegade, Building 540  
DK-8000 Aarhus C, Denmark



# Using Transformations in the Implementation of Higher-order Functions

Hanne Riis Nielson    Flemming Nielson

Computer Science Department,  
Aarhus University,  
Ny Munkegade 116, DK-8000 Aarhus C, Denmark.

## Abstract

Traditional functional languages do not have an explicit distinction between binding times. It arises implicitly, however, as typically one instantiates a higher-order function with the arguments that are known whereas the unknown arguments remain to be taken as parameters. The distinction between ‘*known*’ and ‘*unknown*’ is closely related to the distinction between *binding times* like ‘*compile-time*’ and ‘*run-time*’. This leads to using a combination of *polymorphic type inference* and *binding time analysis* for obtaining the required information about which arguments are known.

Following the current trend in the implementation of functional languages we then transform the run-time level of the program (*not* the compile-time level) into *categorical combinators*. At this stage we have a natural distinction between two kinds of program transformations: *partial evaluation* which involves the compile-time level of our notation and *algebraic transformations* (i.e. the application of algebraic laws) which involves the run-time level of our notation.

By reinterpreting the combinators in suitable ways we obtain specifications of *abstract interpretations* (or data flow analyses). In particular, the use of combinators makes it possible to use a general framework for specifying both *forward* and *backward analyses*. The results of these analyses will be used to enable program transformations that are not applicable under all circumstances.

Finally, the combinators can also be reinterpreted to specify *code generation* for various (abstract) machines. To improve the efficiency of the code generated, one may apply abstract interpretations to collect extra information for guiding the code generation. Peephole optimizations may be used to improve the code at the lowest level.

## 1 Introduction

The *functional programming style* is closely related to the use of higher-order functions. In particular, the functional programming style suggests that many function definitions are instances of the same general computational pattern and that this pattern is defined by a higher-order function. The various instances of the pattern are then obtained by supplying the higher-order function with some of its arguments.

One of the benefits of this programming style is the reuse of function definitions and, more importantly, the reuse of properties proved to hold for them: usually a property of a higher-order function carries over to an instance by verifying that the arguments satisfy some simple properties.

One of the disadvantages is that the efficiency in a compiler that reuses code generated for higher-order functions is often rather poor. The reason is that when generating code for the higher-order function it is impossible to make any assumptions about its arguments and to optimize the code accordingly. Also conventional machine architectures make it rather costly to use functions as data.

We shall therefore be interested in *transforming* instances of higher-order functions into functions that can be implemented more efficiently. The key observation in the approach to be presented here is that an instance of a higher-order function is a function where some of the arguments are known and others are not. To be able to exploit this we shall introduce an *explicit* distinction between *known* and *unknown* values or, using traditional compiler terminology, between *compile-time* entities and *run-time* entities. This leads to the following approach to implementing functional languages:

- annotate the programs with *type information* so that they can be uniquely typed in a monotypic type system (Section 3),
- annotate the programs with *binding time information* so that there is an explicit distinction between the computations that involve known (compile-time) data and those that involve unknown (run-time) data (Section 4),
- transform the programs into *combinator form* so that the computations involving unknown (run-time) data are expressed as combinators (Section 5),
- specify various abstract interpretations by *reinterpreting* the combinators and use the results to enable *program transformations* that are not applicable under all circumstances (Sections 6, 7 and 8),
- specify a simple-minded code generation scheme for an abstract machine by *reinterpreting* the combinators (Section 9), and
- specify an abstract interpretation by reinterpreting the combinators and use this to improve the code generation (Section 10).

We shall illustrate our approach on an example program in the enriched  $\lambda$ -calculus to be presented in Section 2. We shall assume that the semantics is *non-strict*, i.e. lazy, although most of the analyses and transformations equally well apply to a language with a strict semantics. Finally, Section 11 contains a discussion of our approach including efficiency, correctness and automation.

## 2 Operations made explicit

In MIRANDA<sup>1</sup> [47] the functions `reduce` and `sum` may be written as

---

<sup>1</sup>MIRANDA is a trademark of Research Software Limited.

```

reduce f u = g
      where g [] = u
            g (x:xs) = f x (g xs)
sum = reduce (+) 0

```

The left hand side of an equation specifies the name of the function and a list of patterns for its parameters and the right hand side is the body of the function. If more than one equation is given for the same function (as is the case for `g`) there is an *implicit* conditional testing the form of the patterns in the parameter list. Recursion is left *implicit* because a function name on the right hand side of an equation that defines that function name indicates recursion (as is the case for `g`). Finally, function application is left *implicit* as the function is just juxtaposed with its argument(s).

A similar equational definition of functions is allowed in STANDARD ML [24]. But here one has the possibility of making some of the implicit operations or concepts more explicit as is illustrated in

```

val reduce = fn f => fn u => let val rec g = fn xs =>
                        if xs = [] then u
                        else f (hd xs) (g (tl xs))
                    in g end;
val sum = reduce (fn x => fn y => x+y) 0;

```

Function abstraction is now expressed *explicitly* by the construct `fn ... => ...` and the recursive structure of `g` is expressed by the *explicit* occurrence of `rec`. Also the test on the form of the list argument is expressed *explicitly*.

Our approach requires *all* the operations to be expressed explicitly. We shall therefore define a small language, an *enriched  $\lambda$ -calculus*, that captures a few of the more important constructs present in modern functional languages like MIRANDA and STANDARD ML but in a *explicit* way.

The general form of a program in the enriched  $\lambda$ -calculus is a sequence of definitions (prefixed by DEF) followed by an expression (prefixed by VAL) and its overall type (prefixed by HAS). In the enriched  $\lambda$ -calculus we shall write the program `sum` as

```

DEF reduce =  $\lambda f. \lambda u. \text{fix } (\lambda g. \lambda xs. \text{if isnil } xs \text{ then } u$ 
                         $\text{else } f \text{ (hd } xs) \text{ (g (tl } xs))})$ 
VAL reduce ( $\lambda x. \lambda y. +(\langle x, y \rangle)$ ) (0)
HAS Int list  $\rightarrow$  Int

```

Here we use the special parentheses ( and ) for the explicit function application. Furthermore, we use `fix` to make recursive definitions and the angle brackets `<` and `>` to construct pairs. The constructs `isnil e`, `hd e`, `tl e`, `fix e` and `if  $e_1$  then  $e_2$  else  $e_3$`  are built-in constructs of the language and thus we omit the explicit parantheses. Additionally we have used the (ordinary) parentheses ( and ) to indicate the parsing of the expression.

### 3 Types made explicit

Both MIRANDA and STANDARD ML enjoy the property that a programmer need not specify the types of the entities defined in the program. The implementations of the languages are able to infer those types if the program can be consistently typed at all.

This is important for the functional programming style because then the higher-order functions can be instantiated much more freely.

As an example, implementations of MIRANDA and STANDARD ML will infer that the type of **reduce** is

$$(\alpha \rightarrow \beta \rightarrow \beta) \rightarrow \beta \rightarrow \alpha \text{ list} \rightarrow \beta$$

where  $\alpha$  and  $\beta$  are type variables. The occurrence of **reduce** in the definition of **sum** has the type

$$(\text{Int} \rightarrow \text{Int} \rightarrow \text{Int}) \rightarrow \text{Int} \rightarrow \text{Int list} \rightarrow \text{Int}$$

because it is applied to arguments of type  $\text{Int} \rightarrow \text{Int} \rightarrow \text{Int}$  and  $\text{Int}$ .

The enriched  $\lambda$ -calculus is equipped with a type inference system closely related to that found in MIRANDA and STANDARD ML [25,11] but somewhat more monomorphic. As an example, the function application  $e(e')$  is only well-formed if the type of  $e$  has the form  $t \rightarrow t'$  and if the type of  $e'$  is  $t$  and then the type of the application is  $t'$ . Another example is the construct **isnil**  $e$  which is well-formed of type **Bool** provided that  $e$  is well-formed and has a type of the form  $t \text{ list}$ . Similar rules exist for the other composite constructs of the language. For the constants we have axioms stating e.g. that  $+$  has type  $\text{Int} \times \text{Int} \rightarrow \text{Int}$  and that  $0$  has type  $\text{Int}$ . Based upon such axioms and rules we can infer that the program **sum** is well-formed and we can determine the types of the various subexpressions. For some programs these types may contain type variables and these type variables will then be replaced by a trivial type called **Void**. The purpose of the type following the keyword **HAS** is to avoid all type variables to be replaced by **Void**: those type variables that must match a subtype of the type following **HAS** will be instantiated accordingly.

The next step is to *annotate* the program with the inferred type information: we shall add the actual types to the constants and the bound variables of  $\lambda$ -abstractions. This means that the type analysis will transform the untyped program for **sum** (of Section 2) into the following typed program (to be called **sum<sub>t</sub>**)

```
DEF reduce =  $\lambda f[\text{Int} \rightarrow \text{Int} \rightarrow \text{Int}]. \lambda u[\text{Int}].$ 
              fix ( $\lambda g[\text{Int list} \rightarrow \text{Int}]. \lambda xs[\text{Int list}].$ 
                  if isnil xs then u
                  else f(hd xs)(g(tl xs)))
VAL reduce ( $\lambda x[\text{Int}]. \lambda y[\text{Int}]. +[\text{Int} \times \text{Int} \rightarrow \text{Int}](\langle x, y \rangle)$ ) (0[Int])
HAS Int list  $\rightarrow$  Int
```

Note that the type of **reduce** is fixed. This is possible because there is only *one* application of **reduce** in the program. Otherwise we may have to duplicate the definition as the development to be performed does not (yet) allow the use of type variables. (This is what was meant above by the type system being more monomorphic than those for MIRANDA and STANDARD ML.)

## 4 Binding time made explicit

Neither MIRANDA, STANDARD ML nor the enriched  $\lambda$ -calculus has an explicit distinction between binding times. However, for higher-order functions we can distinguish between the parameters that are known and those that are not. The idea is now to capture this implicit distinction between binding times and then annotate the operations of the enriched  $\lambda$ -calculus accordingly.

## 4.1 2-level syntax

$\text{Reduce}_1 = (\text{Int} \rightarrow \text{Int} \rightarrow \text{Int}) \rightarrow \text{Int} \rightarrow \text{Int list} \rightarrow \text{Int}$
$\text{Reduce}_2 = (\text{Int} \rightarrow \underline{\text{Int}} \rightarrow \text{Int}) \rightarrow \text{Int} \rightarrow \text{Int list} \rightarrow \text{Int}$
$\text{Reduce}_3 = (\text{Int} \rightarrow \text{Int} \rightarrow \text{Int}) \rightarrow \text{Int} \rightarrow \underline{\text{Int list}} \rightarrow \underline{\text{Int}}$
$\text{Reduce}_4 = (\underline{\text{Int}} \rightarrow \text{Int} \rightarrow \text{Int}) \rightarrow \text{Int} \rightarrow \text{Int list} \rightarrow \text{Int}$
$\text{Reduce}_5 = (\text{Int} \rightarrow \underline{\text{Int}} \rightarrow \text{Int}) \rightarrow \text{Int} \rightarrow \underline{\text{Int list}} \rightarrow \underline{\text{Int}}$
$\text{Reduce}_6 = (\text{Int} \rightarrow \text{Int} \rightarrow \text{Int}) \rightarrow \underline{\text{Int}} \rightarrow \underline{\text{Int list}} \rightarrow \underline{\text{Int}}$
$\text{Reduce}_7 = (\underline{\text{Int}} \rightarrow \underline{\text{Int}} \rightarrow \underline{\text{Int}}) \rightarrow \text{Int} \rightarrow \underline{\text{Int list}} \rightarrow \underline{\text{Int}}$
$\text{Reduce}_8 = (\text{Int} \rightarrow \underline{\text{Int}} \rightarrow \text{Int}) \rightarrow \underline{\text{Int}} \rightarrow \underline{\text{Int list}} \rightarrow \underline{\text{Int}}$
$\text{Reduce}_9 = (\underline{\text{Int}} \rightarrow \text{Int} \rightarrow \text{Int}) \rightarrow \underline{\text{Int}} \rightarrow \underline{\text{Int list}} \rightarrow \underline{\text{Int}}$
$\text{Reduce}_{10} = (\underline{\text{Int}} \rightarrow \underline{\text{Int}} \rightarrow \underline{\text{Int}}) \rightarrow \underline{\text{Int}} \rightarrow \underline{\text{Int list}} \rightarrow \underline{\text{Int}}$

Table 1: Well-formed 2-level types for **reduce**

We shall use the types of functions to record when their parameters will be available and their results produced. For **sum** it is clear that the list parameter is not available at compile-time and we shall record this by underlining the corresponding component of the type

$$\text{Sum}_t = \underline{\text{Int list}} \rightarrow \text{Int}$$

The fact that the list argument is not available at compile-time will have consequences for when the parameters are available for **reduce**. Again we shall record this by underlining parts of the type

$$\text{Reduce}_t = (\text{Int} \rightarrow \text{Int} \rightarrow \text{Int}) \rightarrow \text{Int} \rightarrow \underline{\text{Int list}} \rightarrow \text{Int}$$

$\text{Sum}_t$  and  $\text{Reduce}_t$  are examples of *2-level types*.

We shall interpret an underlined type as meaning that values of that type cannot be expected to be available until run-time. On the other hand, a type that is not completely underlined will denote values that definitely will be available at compile-time. With this intuition in mind we claim that  $\text{Sum}_t$  and  $\text{Reduce}_t$  are unacceptable annotations because they denote functions that operate on actual run-time data but at compile-time. Our point of view shall therefore be that  $\text{Sum}_t$  and  $\text{Reduce}_t$  are *incomplete* annotations and that we need to make them well-formed.

To motivate our definition of a *well-formed* 2-level type we need to take a closer look at the interplay between the compile-time level and the run-time level. Thinking in terms of a compiler it is quite clear that at compile-time we can manipulate pieces of code (to be executed at run-time) but we cannot manipulate entities computed at run-time. Hence at compile-time we cannot directly manipulate objects of type  $\underline{\text{Int list}}$  whereas we can manipulate objects of type  $\underline{\text{Int list}} \rightarrow \underline{\text{Int}}$  because the latter type may be regarded as the type of code for functions (to be executed at run-time). So  $\text{Int} \rightarrow \underline{\text{Int list}} \rightarrow \underline{\text{Int}}$  will be the type of a function that given an integer at compile-time will produce a piece of code that has to be executed at run-time. Similarly  $\underline{\text{Int}} \rightarrow \underline{\text{Int list}} \rightarrow \underline{\text{Int}}$  will be the type of a function to be executed at run-time whereas  $\text{Int} \rightarrow \text{Int list} \rightarrow \text{Int}$  will be the type of a function to be executed at compile-time.

To be a bit more precise we shall say that an ‘all underlined’ function type is well-formed and we shall call it a *run-time type*. A type that does not contain any underlinings




Figure 1: Compatible 2-level types for `reduce`

will also be well-formed since it will be the type of values fully computed at compile-time. Well-formed types can then be combined using the type constructors  $\rightarrow$ ,  $\times$  and `list`.

Corresponding to the type of `sum` we have two well-formed 2-level types

$\text{Sum}_1 = \text{Int list} \rightarrow \text{Int}$

$\text{Sum}_2 = \underline{\text{Int}} \underline{\text{list}} \rightarrow \underline{\text{Int}}$

For `reduce` we have the ten well-formed 2-level types shown in Table 1.

We shall say that two 2-level types have the same *underlying type* if they are equal except for the underlinings. One 2-level type is then *compatible* with another if they have the same underlying type and if underlined occurrences in the latter also are underlined in the former. Thus the intention is that compatible types may be obtained by ‘moving’ data, computations or results from compile-time to run-time but not vice versa. As an example `Reduce3` is compatible with `Reducet` but `Reduce2` is not. This is expressed by the Hasse diagram of Figure 1.

We shall say that the *best completion* of `Reducet` is `Reduce3` because all other well-formed 2-level types that are compatible with `Reducet` also will be compatible with `Reduce3`. (Shortly we shall see that the actual definition of `reduce` will impose further restrictions upon the type.) Similarly, the best completion of `Sumt` is `Sum2`. A formal account may be found in [37].

## 4.2 Binding time analysis

We shall now annotate complete programs so that it becomes explicit which computations should be performed at compile-time (namely those that are not underlined) and which should be postponed until run-time (namely those that are underlined). The resulting program is called a *2-level program* and similarly, an annotated expression is called a *2-level expression*.

The above discussion suggests that the annotation of `reduce` should reflect the binding time information given by `Reduce3`. However, it is not possible to annotate the definition of `reduce` so that it is both well-formed and will have type `Reduce3`. The reason is that the binding times do *not* match: according to `Reduce3` the function argument `f` has type `Int → Int → Int` and the list argument `xs` has type `Int list` so `f(hd xs)` will supply `f` with a run-time argument even though it expects a compile-time argument. This will be formalized by defining a *well-formedness condition* on 2-level expressions [38]<sup>2</sup>. There are two ingredients in this. One is that the underlying expression (obtained by removing the underlinings) must type-check properly. The other ingredient is that the binding times must agree. Thus a function with a run-time type must be given a run-time argument and we express this by underlining the  $\lambda$ -abstraction of the function definition, as in  $\underline{\lambda}xs[ ] \dots$ , and by underlining the parentheses in the function application, as in  $f(\underline{\dots})$ . Similarly, a function with a compile-time type must be supplied with a compile-time argument and in that case we neither underline the  $\lambda$ -abstraction nor the parentheses of the function application.

Given an incomplete annotation `sumt` of `sum`

```
DEF reduce =  $\lambda f[Int \rightarrow Int \rightarrow Int]. \lambda u[Int].$ 
              fix ( $\lambda g[Int\ list \rightarrow Int]. \lambda xs[Int\ list].$ 
                  if isnil xs then u
                  else f(hd xs) (g(tl xs)))
VAL reduce ( $\lambda x[Int]. \lambda y[Int]. +[Int \times Int \rightarrow Int] (\langle x, y \rangle)$ ) (0[Int])
HAS Int list  $\rightarrow$  Int
```

the *binding time analysis* [36,37,45] will complete the annotation so that the resulting 2-level program is

- well-formed (as explained above),
- postpones as few computations as possible to run-time, and
- is compatible with the incomplete annotation (where compability for expressions is very much as for types).

The program obtained<sup>3</sup> from `sumt` is `sum9`

```
DEF reduce9 =  $\lambda f[\underline{Int \rightarrow Int \rightarrow Int}]. \lambda u[\underline{Int}].$ 
               fix ( $\lambda g[\underline{Int\ list \rightarrow Int}]. \lambda xs[\underline{Int\ list}].$ 
                   if isnil xs then u
                   else f(hd xs) (g(tl xs)))
VAL reduce9 ( $\lambda x[\underline{Int}]. \lambda y[\underline{Int}]. +[\underline{Int \times Int \rightarrow Int}] (\underline{\langle x, y \rangle})$ ) (0[Int])
HAS Int list  $\rightarrow$  Int
```

where `reduce9` has type `Reduce9`. Note that we cannot use  $\lambda u[\underline{Int}]. \dots$  instead of  $\underline{\lambda}u[\underline{Int}]. \dots$  as the well-formedness conditions [38] on types only allow us to manipulate run-time *function* types at compile-time and not run-time data types like `Int`.

The remaining well-formed 2-level programs are less interesting. Corresponding to the 2-level type `Reduce1` we have a program with no underlinings at all and corresponding

<sup>2</sup>Alternative well-formedness conditions are studied in [36,37,41,45].

<sup>3</sup>Actually we apply an adapted version of the binding time analysis algorithm of [37] in order to comply with the well-formedness predicate of [38].



to  $\text{Reduce}_{10}$  we have a program where all operations are underlined. These are the only well-formed 2-level programs with  $\text{sum}$  as the underlying program<sup>4</sup>. Note that in each of these programs the binding times for  $\text{reduce}$  is fixed and this is possible because there is only *one* application of  $\text{reduce}$  in the program. Otherwise we may have to duplicate the definition of the function.

### 4.3 Improving the binding time analysis

The annotation obtained from the binding time analysis is *optimal* in the sense that as few computations as possible are postponed until run-time. However, it is often the case that a slight rewriting of the program will produce an even better distinction between the binding times. As an example, the order of the parameters of a function may be changed or the representation of data types may be modified. To illustrate this consider the function lookup of type

$$(\text{Name} \times \underline{\text{Int}}) \text{ list} \rightarrow \text{Name} \rightarrow \text{Int}$$

where we assume that the second components of all pairs in the first argument are unknown at compile-time. We then have a situation where known and unknown data are mixed and the binding time analysis of [37] will return a function with the annotated type

$$(\underline{\text{Name}} \times \underline{\text{Int}}) \underline{\text{list}} \rightarrow \underline{\text{Name}} \rightarrow \underline{\text{Int}}$$

so that all computations will be postponed until run-time. However, we may split the list of pairs into two lists and rearrange the order of the parameters so that the type becomes

$$\text{Name list} \rightarrow \text{Name} \rightarrow \underline{\text{Int list}} \rightarrow \text{Int}$$

We then get a much better distinction between the binding times because now only the elements of the *second* list will be unknown at compile-time and the binding time analysis will return a function with the annotated type

$$\text{Name list} \rightarrow \text{Name} \rightarrow \underline{\text{Int list}} \rightarrow \underline{\text{Int}}$$

Hence some of the computations can be performed at compile-time and this idea is further explored in [42].

The above example is rather involved in that it changes the overall type of the function. However, it may also be useful to perform some more local transformations. Consider the program  $\text{sum}_9$  of the previous subsection. Here the fixed point operation of  $\text{reduce}_9$  is underlined and intuitively this means that we cannot use the recursive structure of its body at compile-time, e.g. during abstract interpretation and when generating code. We may therefore want to replace the run-time fixed point by a compile-time fixed point. Since  $\text{sum}_9$  is the best completion of  $\text{sum}_t$  we cannot obtain this effect by simply changing the annotation. The idea is therefore first

- to *transform* the underlying program, and then
- to *repeat* the binding time analysis.

In our case we shall apply the transformation

---

<sup>4</sup>This statement holds for the well-formedness predicate of [38]. If the predicate of [37,36] is used then there will be two more well-formed programs.

$$\lambda x[t_1].\text{fix}(\lambda f[t_2 \rightarrow t_3].e) \Rightarrow \text{fix}(\lambda g[t_1 \rightarrow t_2 \rightarrow t_3].\lambda x[t_1].e[g(x)/f])$$

that replaces the first pattern with the second. Here  $e[g(x)/f]$  is  $e$  with all occurrences of  $f$  replaced by  $g(x)$  and  $g$  is assumed to be a fresh identifier. We then repeat the binding time analysis with  $\text{Sum}_t$  as the ‘goal’ type and get

```
DEF reduce9a = λf[Int → Int → Int].
  fix (λg'[Int → Int list → Int]. λu[Int]. λxs[Int list].
    if isnil xs then u
    else f(hd xs)(g'(u)(tl xs)))
VAL reduce9a (λx[Int]. λy[Int]. +[Int × Int → Int](⟨x,y⟩)) (0[Int])
HAS Int list → Int
```

where the fixed point now is computed at compile-time and the function  $g'$  is bound at compile-time.

As a side effect the functionality of the fixed point has been changed and, in particular,  $g'$  has become a higher-order function. For some stack-based implementations it is expensive to handle higher-order functions and we may therefore want to transform the program further to change the functionality of  $g'$ . To do that we first apply the following transformation to the underlying program of  $\text{sum}_{9a}$

$$\begin{aligned} \text{fix}(\lambda f[t_1 \rightarrow t_2 \rightarrow t_3].\lambda x[t_1].\lambda y[t_2].e) \Rightarrow \\ \lambda x[t_1].\lambda y[t_2].(\text{fix}(\lambda g[t_1 \times t_2 \rightarrow t_3].\lambda z[t_1 \times t_2]. \\ e[\lambda x[t_1].\lambda y[t_2].g(\langle x,y \rangle)/f] [\text{fst } z/x][\text{snd } z/y])) \\ (\langle x,y \rangle)) \end{aligned}$$

Here  $g$  and  $z$  are assumed to be fresh identifiers. Next we apply the  $\beta$ -transformation

$$(\lambda x[t].e)(e') \Rightarrow e[e'/x]$$

(recalling that our semantics is non-strict). The binding time analysis is then applied to the resulting program with  $\text{Sum}_t$  as the overall annotated type and we get the program  $\text{sum}_{9b}$

```
DEF reduce9b = λf[Int → Int → Int]. λu[Int]. λxs[Int list].
  (fix (λg[Int × Int list → Int]. λz[Int × Int list].
    if isnil (snd z) then fst z
    else f(hd (snd z))(g(⟨fst z, tl (snd z)⟩))))
  (⟨u,xs⟩)
VAL reduce9b (λx[Int]. λy[Int]. +[Int × Int → Int](⟨x,y⟩)) (0[Int])
HAS Int list → Int
```

Note, however, that the motivation for this transformation is not that fewer computations are postponed until run-time. Rather the transformation will allow us to specify a backward abstract interpretation in Section 8. Also note that the two transformations preserve total correctness.

## 4.4 Remarks

We have experimented with different ways in which the two binding levels may interact [36,37,38] but the restrictions have always been rather conservative. As an example,

we have not allowed values computed at compile-time to be used directly at run-time because of the computational consequences of allowing this for function types. Also our binding time analysis does not change the underlying program. Flexibility is introduced by allowing arbitrary (semantics preserving) program transformations, e.g. expressed as simple pattern matching templates or as in [42] within the unfold/fold framework of [7].

There are alternative approaches to binding time analysis that allow greater interaction between the levels and which incorporates certain program transformations (see e.g. [21, 22, 27]). The price to be paid is that it becomes harder to reason about the correctness properties. Our approach gives a better separation of concerns: It is fairly straightforward to prove the correctness and optimality of our binding time analysis algorithm [37] and at the same time we are not restricted in the kind of program transformations that can be made.

Our approach to binding time analysis is purely *syntactic* whereas e.g. Mogensen [26] takes a more semantic approach using projections. In fact our approach is strongly motivated by the (syntactic) approach to polymorphic type inference of [25]. The structure of our binding time analysis algorithm is very close to the algorithm  $\mathcal{W}$  of [25], the main difference being that we have no counterpart of substitutions and therefore have to perform some extra recursive calls to ‘unify’ the binding time annotations.

Finally, we should point out that we only allow explicit monotypes in Section 3 because it is unclear how to do binding time annotations for type variables. In particular the identity function  $\lambda x.x$  might have type  $\alpha \multimap \alpha$  and  $\alpha \rightarrow \alpha$  depending upon the kind of 2-level types that  $\alpha$  may range over.

## 5 Combinators made explicit

The binding time information of a 2-level program clearly indicates which computations should be carried out at compile-time and which should be carried out at run-time. The compile-time computations should be executed by a compiler and it is well-known how to do this. The run-time computations should give rise to code instead. We may also want to perform some data flow analyses either in order to validate some program transformations or to improve the efficiency of the code generated. It is important to observe that it is the run-time computations, not the compile-time computations, that should be analysed just as it is the run-time computations, not the compile-time computations, that should give rise to code. This then calls for the ability to *interpret the run-time constructs* in different ways depending upon the task at hand.

It is not straightforward to do so when the run-time computations are expressed in the form of  $\lambda$ -expressions. As an example, the usual meaning of

$$\lambda x[\text{Int} \times \text{Int}]. f (g (x))$$

is  $\lambda v.f(g v)$ . However, we may be interested in an analysis determining whether both components of  $x$  are needed in order to compute the result. This is an example of a backward data flow analysis and, as we shall see later, the natural interpretation of the expression will then be  $\lambda v.g(f v)$ . It is not straightforward to interpret function abstraction and function application so as to be able to obtain both meanings. The idea is therefore to focus on functions and functionals (expressed as combinators as in [2, 10]) rather than values and functions. We then write

$$f \sqcap g$$

for the expression above and the effect of both  $\lambda v.f(g\ v)$  and  $\lambda v.g(f\ v)$  can be obtained by suitably reinterpreting the functional  $\square$ .

This observation calls for transforming the run-time computations into combinator form. Similar considerations motivate the use of combinators in the implementation of functional languages [46,9,17] and the use of categorical combinators when interpreting the typed  $\lambda$ -calculus in an arbitrary cartesian closed category [23]. However, it is important to stress that we shall leave the compile-time computations in the form of  $\lambda$ -expressions and only transform the run-time computations into combinator form.

## 5.1 Combinator introduction

A 2-level program (or 2-level expression) is in *combinator form* whenever all the run-time computations are expressed as *categorical combinators*. For the program  $\text{sum}_{9b}$  the corresponding program  $\text{sum}_{9B}$  in combinator form will be

```
DEF reduce9B =  $\lambda f[ ]$ .Curry (fix( $\lambda g[ ]$ .
    Cond(Isnil[ ]  $\square$  Snd[ ], Fst[ ],
        Apply[ ]  $\square$  Tuple(f  $\square$  Hd[ ]  $\square$  Snd[ ],
            g  $\square$  Tuple(Fst[ ], Tl[ ]  $\square$  Snd[ ]))))))
VAL Apply[ ]  $\square$  Tuple((reduce9B (Curry +[ ]))  $\square$  (Const[ ] 0[ ]), Id[ ])
HAS Int list  $\rightarrow$  Int
```

as we will explain below. For the sake of readability we have omitted the type information (in square brackets).

The conditional of  $\text{reduce}_{9B}$  has been replaced by the combinator  $\text{Cond}$  where the intention is that

$$\text{Cond}(f,g,h) = \lambda x. \text{if } f(x) \text{ then } g(x) \text{ else } h(x)$$

Within the conditional the test isnil (snd  $z$ ) is replaced by  $\text{Isnil}[ ] \square \text{Snd}[ ]$  because the run-time parameter will be implicit. Similarly, the true-branch fst  $z$  is replaced by  $\text{Fst}[ ]$ . Here  $\text{Isnil}$ ,  $\text{Fst}$  and  $\text{Snd}$  are combinators corresponding to isnil  $e$ , fst  $e$  and snd  $e$ , respectively. To understand the transformation of the false-branch note that the intention with the combinators  $\text{Apply}$  and  $\text{Tuple}$  is given by

$$\text{Apply} = \lambda \langle f, x \rangle. f\ x$$

$$\text{Tuple}(f, g) = \lambda x. \langle f(x), g(x) \rangle$$

and that  $\text{Hd}$  and  $\text{Tl}$  correspond to hd  $e$  and tl  $e$ , respectively.

The function  $\text{reduce}$  takes its arguments one at a time and this is achieved using the combinator  $\text{Curry}$

$$\text{Curry } f = \lambda x. \lambda y. f\ (\langle x, y \rangle)$$

In the expression part of the program we use the additional combinators  $\text{Id}$  and  $\text{Const}$  and the intention is that  $\text{Id}$  is the identity function and  $\text{Const } f$  equals  $\lambda x.f$ .

To transform well-formed 2-level programs into combinator form we may use the algorithm of [38]<sup>5</sup>. It is an extension of the usual algorithm for translating the typed  $\lambda$ -calculus into categorical combinators [10]. Thus for each of the programs  $\text{sum}_1$ ,  $\text{sum}_9$ ,  $\text{sum}_{9a}$ ,  $\text{sum}_{9b}$  and  $\text{sum}_{10}$  we will have equivalent programs in combinator form. In the remainder of this paper we shall restrict our attention to  $\text{sum}_{9B}$ .

<sup>5</sup>The algorithm in [41] may also be used if the more general well-formedness predicate of [36,37] is used (in Section 4) but unfortunately it uses a rather unnatural combinator in certain cases.

## 5.2 Partial evaluation

The program `sum9B` can be further simplified because we can supply the occurrence of `reduce` with its first parameter which is known at compile-time. For this we need the two transformations

$$\begin{aligned} \text{DEF } f = e \text{ VAL } e' \text{ HAS } t &\Rightarrow \text{VAL } e'[e/f] \text{ HAS } t \\ (\lambda x. e')(e) &\Rightarrow e'[e/x] \end{aligned}$$

which preserve total correctness (due to the use of a non-strict semantics). This sort of program transformation is often called *partial evaluation* [14,21,36] and in our case it gives rise to the program

```
VAL Apply[ ] □ Tuple((Curry (fix(λg[ ].
    Cond(Isnil[ ] □ Snd[ ], Fst[ ],
        Apply[ ] □ Tuple((Curry +[ ]) □ Hd[ ] □ Snd[ ],
            g □ Tuple(Fst[ ], Tl[ ] □ Snd[ ]))))))
    □ (Const[ ] 0[ ], Id[ ]))
HAS Int list → Int
```

This transformation could equally well have been performed before the combinators were introduced.

## 5.3 Algebraic transformation

The run-time counterpart of partial evaluation is called *algebraic transformation* [2]. An example is the transformation

$$\text{Apply}[ ] \square \text{Tuple}((\text{Curry } e) \square e', e'') \Rightarrow e \square \text{Tuple}(e', e'')$$

If we apply this transformation twice to the above program then we get the program `sum'9B`

```
VAL fix(λg[ ].Cond(Isnil[ ] □ Snd[ ], Fst[ ],
    +[ ] □ Tuple(Hd[ ] □ Snd[ ],
        g □ Tuple(Fst[ ], Tl[ ] □ Snd[ ]))))
    □ Tuple(Const[ ] 0[ ], Id[ ])
HAS Int list → Int
```

Note that by now all higher-order run-time functions have disappeared. Also note that the transformation preserves total correctness.

The algebraic transformations play an important role in simplifying programs. The algorithm for combinator introduction proceeds by structural induction and the algebraic transformations can be used to reduce certain unnecessarily complicated expressions that arise in this process. Also the application of partial evaluation and algebraic transformations can be intertwined. In Section 8 we shall see how the algebraic transformations can simplify results obtained from further transformations.

## 5.4 Remarks

We have used the categorical combinators for mainly two reasons. One is that they are *functionally complete*, meaning that their expressive power is equivalent to that of the typed  $\lambda$ -calculus [10]. As a consequence of this it is possible to translate expressions in the typed  $\lambda$ -calculus into expressions of the categorical combinatory logic. The other reason is that we get a *fixed set of combinators* and this is important in the next section where we introduce the concept of a parameterized semantics.

We believe that the development of this section can be performed for any fixed set of combinators that is sufficiently expressive. The traditional algorithms for combinator introduction must then be extended so that the combinators are introduced only for the run-time constructs of the notation. The concept of partial evaluation is unchanged as it does not depend on the run-time level at all whereas the algebraic transformations must be modified so as to use algebraic laws that hold for the actual set of combinators.

Some of the most popular approaches to implementing functional languages use supercombinators [17,44]. The actual set of combinators will then depend upon the program at hand so that our approach does not directly carry over. It should be fairly straightforward to extend the algorithm for introducing supercombinators [17] so that supercombinators are generated only for the run-time constructs. The concept of algebraic transformation will have to be modified since we do not have algebraic laws in the same sense as before. For this one may explore some variant of partial evaluation that only considers how to rewrite the supercombinators.

## 6 Parameterized semantics

We want to interpret the run-time constructs in different ways depending on the task at hand and at the same time we want the meaning of the compile-time constructs to be fixed. To make this possible we shall *parameterize* the semantics on an *interpretation* [32,33,35,39] specifying the meaning of the run-time level. The interpretation will define

- the meaning of the run-time function types, and
- the meaning of the combinators (and the compile-time fixed point).

Relative to an interpretation one can then define the semantics of all well-formed 2-level types, of all well-formed 2-level programs and of all well-formed 2-level expressions in combinator form. In the next sections we shall give some example interpretations that specify forward and backward data flow analyses and code generation. The correctness of these interpretations will be expressed relative to a *non-strict* (i.e. lazy) standard interpretation to be given below.

### 6.1 The type part of an interpretation

The *type part* of an interpretation  $\mathcal{I}$  will specify a set  $\mathcal{I}_{t \multimap t'}$  of values for each well-formed type  $t \multimap t'$ . (Technically, this set is extended with a partial ordering expressing when one value is less defined than another but we shall dispense with this in the present paper.) For the standard interpretation  $\mathcal{S}$  we have

$$\mathcal{S}_{t \multimap t'} = \llbracket t \rrbracket(\mathcal{S}) \rightarrow \llbracket t' \rrbracket(\mathcal{S})$$

where  $\llbracket t \rrbracket(\mathcal{S})$  and  $\llbracket t' \rrbracket(\mathcal{S})$  are the sets of values of type  $t$  and  $t'$ , respectively, and  $\rightarrow$  constructs the appropriate function space. We can then define

$$\begin{aligned}\llbracket \text{Int} \rrbracket(\mathcal{S}) &= \mathbf{Z} && \text{(the set of integers)} \\ \llbracket \text{Bool} \rrbracket(\mathcal{S}) &= \mathbf{T} && \text{(the set of truth values)}\end{aligned}$$

and more interestingly

$$\begin{aligned}\llbracket t \rightarrow t' \rrbracket(\mathcal{S}) &= \llbracket t \rrbracket(\mathcal{S}) \rightarrow \llbracket t' \rrbracket(\mathcal{S}) && \text{(functions)} \\ \llbracket t \times t' \rrbracket(\mathcal{S}) &= \llbracket t \rrbracket(\mathcal{S}) \times \llbracket t' \rrbracket(\mathcal{S}) && \text{(pairs of values)} \\ \llbracket t \text{ list} \rrbracket(\mathcal{S}) &= \llbracket t \rrbracket(\mathcal{S})^* && \text{(sequences of values)}\end{aligned}$$

so that  $\llbracket t \rrbracket(\mathcal{S})$  is defined for all run-time types.

Given the meaning of the run-time function types it is straightforward to extend it to all well-formed compile-time types. In general, the set  $\llbracket t \rrbracket(\mathcal{I})$  of values associated with the type  $t$  will be defined by structural induction and we have

$$\begin{aligned}\llbracket \text{Int} \rrbracket(\mathcal{I}) &= \mathbf{Z} \\ \llbracket \text{Bool} \rrbracket(\mathcal{I}) &= \mathbf{T} \\ \llbracket t \rightarrow t' \rrbracket(\mathcal{I}) &= \llbracket t \rrbracket(\mathcal{I}) \rightarrow \llbracket t' \rrbracket(\mathcal{I}) \\ \llbracket t \times t' \rrbracket(\mathcal{I}) &= \llbracket t \rrbracket(\mathcal{I}) \times \llbracket t' \rrbracket(\mathcal{I}) \\ \llbracket t \text{ list} \rrbracket(\mathcal{I}) &= \llbracket t \rrbracket(\mathcal{I})^* \\ \llbracket t \rightarrow t' \rrbracket(\mathcal{I}) &= \mathcal{I}_{t \rightarrow t'}\end{aligned}$$

As an example,  $\llbracket \text{Sum}_1 \rrbracket(\mathcal{S}) = \llbracket \text{Sum}_2 \rrbracket(\mathcal{S}) = \mathbf{Z}^* \rightarrow \mathbf{Z}$  so that the distinction between the compile-time function type and the run-time function type disappears (as must be expected in a standard semantics).

## 6.2 The expression part of an interpretation

The *expression part* of an interpretation specifies a value or a function for each of the combinators  $\square$ , **Tuple**, **Fst**, etc. In Table 2 we give parts of the specification of  $\mathcal{S}$  using an appropriate mathematical notation (shown in boldface). The meaning of the compile-time constructs is predetermined except for that of **fix**  $e$ . Here the interpretation  $\mathcal{I}$  specifies a function  $\mathcal{I}_{\text{fix}}^t$  for each type  $t$  of **fix**  $e$ . In Table 2 we write **FIX** for the (least) fixed point operator.

Given the meanings of the combinators and of **fix** it is now straightforward to extend it to expressions and programs. For expressions we define a value  $\llbracket e \rrbracket(\mathcal{I})$  in  $\llbracket t \rrbracket(\mathcal{I})$  for every well-formed expression  $e$  of type  $t$ . To cater for the free variables we need an environment  $env$  that associates the variables with their values. We shall not give all the details here but some of the more interesting clauses are

$$\begin{aligned}\llbracket \lambda x[t].e \rrbracket(\mathcal{I}) \text{ env } \mathbf{v} &= \llbracket e \rrbracket(\mathcal{I}) \text{ env } [\mathbf{v}/x] \\ \llbracket e(e') \rrbracket(\mathcal{I}) \text{ env } &= (\llbracket e \rrbracket(\mathcal{I}) \text{ env}) (\llbracket e' \rrbracket(\mathcal{I}) \text{ env}) \\ \llbracket x \rrbracket(\mathcal{I}) \text{ env } &= \text{env}(x)\end{aligned}$$

where  $\text{env}[\mathbf{v}/x]$  is as  $env$  except that  $x$  has the value  $\mathbf{v}$ . The next example clauses illustrate the use of the interpretation  $\mathcal{I}$

$\mathcal{S}_\square (f,g) v = f(g v)$
$\mathcal{S}_{\text{Tuple}} (f,g) v = \langle f v, g v \rangle$
$\mathcal{S}_{\text{Fst}} \langle u, v \rangle = u$
$\mathcal{S}_{\text{Snd}} \langle u, v \rangle = v$
$\mathcal{S}_{\text{Hd}} (v:l) = v$
$\mathcal{S}_{\text{Tl}} (v:l) = l$
$\mathcal{S}_{\text{Isnil}} l = (l=[])\rightarrow\text{true} \mid (l=v:l')\rightarrow\text{false}$
$\mathcal{S}_{\text{Cond}} (f,g,h) v = (f v=\text{true})\rightarrow g v \mid (f v=\text{false})\rightarrow h v$
$\mathcal{S}_{\text{Const}} f v = f$
$\mathcal{S}_{\text{Id}} v = v$
$\mathcal{S}_+ \langle u, v \rangle = u+v$
$\mathcal{S}_{\text{fix}}^t f = \text{FIX } f \quad \text{for all } t$

Table 2: Fragments of the expression part of  $\mathcal{S}$

$$\begin{aligned}
\llbracket e \square e' \rrbracket(\mathcal{I}) \text{ env} &= \mathcal{I}_\square (\llbracket e \rrbracket(\mathcal{I}) \text{ env}, \llbracket e' \rrbracket(\mathcal{I}) \text{ env}) \\
\llbracket \text{Tuple}(e, e') \rrbracket(\mathcal{I}) \text{ env} &= \mathcal{I}_{\text{Tuple}} (\llbracket e \rrbracket(\mathcal{I}) \text{ env}, \llbracket e' \rrbracket(\mathcal{I}) \text{ env}) \\
\llbracket \text{Fst} [ ] \rrbracket(\mathcal{I}) \text{ env} &= \mathcal{I}_{\text{Fst}} \\
\llbracket \text{fix } e \rrbracket(\mathcal{I}) \text{ env} &= \mathcal{I}_{\text{fix}}^t (\llbracket e \rrbracket(\mathcal{I}) \text{ env}) \quad \text{where } \text{fix } e \text{ has type } t
\end{aligned}$$

To summarize,  $\llbracket e \rrbracket(\mathcal{I}) \text{ env}$  is defined by structural induction on  $e$  and in order to compose the results obtained for the subexpressions we will either consult the interpretation  $\mathcal{I}$  (if the constructor is a combinator or **fix**) or we will use a standard approach.

Finally, the semantics of a program is defined to be that of the expression after ‘**VAL**’. The specification of  $\mathcal{S}$  given in Table 2 is sufficient to determine that

$$\llbracket \text{sum}'_{9B} \rrbracket(\mathcal{S}) l = \text{FIX}(\lambda g. \lambda \langle v, l \rangle. (l=[])\rightarrow v \mid (l=a:l')\rightarrow a+g\langle v, l' \rangle) \langle 0, l \rangle$$

### 6.3 Remarks

In the previous sections we have presented type inference, binding time analysis and combinator introduction as purely syntactic manipulations of programs. We have only referred to (a vague and unspecified notion of) semantics when talking about program transformations and the extent to which they are correctness preserving. Having introduced the parameterized semantics and the standard interpretation  $\mathcal{S}$  we can now take at least two approaches.

One approach is to view the parameterized semantics as the ultimate semantics for the original enriched  $\lambda$ -calculus of Section 2. Taking **sum** as an example we thus have to

- transform it into an explicitly typed program (i.e.  $\text{sum}_t$ ),
- then into an explicitly annotated program without doing any program transformations (i.e.  $\text{sum}_9$  rather than  $\text{sum}_{9b}$ ) and
- finally into combinator form without doing any partial evaluation or algebraic transformations (much as  $\text{sum}_{9B}$  was obtained from  $\text{sum}_{9b}$ ).

The semantics of **sum** then amounts to the parameterized semantics of the resulting program, say  $\text{sum}_C$ . We can then express the correctness of the program transformations,



partial evaluations and algebraic transformations as the condition that they do not change the final semantics regardless of the interpretation considered.

Another approach is to give a direct definition of the semantics of the enriched  $\lambda$ -calculus. Then one must ensure that the semantics of an untyped program (e.g. `sum`) corresponds to the parameterized semantics with respect to  $\mathcal{S}$  of the combinator version of the program (e.g. `sumC`). We shall not go deeply into this issue here but we favour the first approach above as our ultimate interest is in non-standard interpretations, e.g. interpretations for code generation.

The combination of using a denotational as well as a parameterized semantics facilitates the development of various *meta-theories*, i.e. frameworks for correctness proofs of abstract interpretations [33,39] and code generation [35]. Also note that the notion of a parameterized semantics would have been a very challenging task to define if we had not restricted our attention to a fixed set of combinators.

## 7 Forward abstract interpretation

In the previous sections we have seen some very simple program transformations that are *universally* applicable. This means that whenever the left hand side pattern matches an expression one may replace the expression by the right hand side pattern. However, in other cases it is necessary to analyse the expression to ensure that certain conditions are fulfilled before the transformation can be applied. This can be illustrated by `sum'9B` where the first component of `g`'s parameter will have the value 0 always and therefore we may want to replace the true-branch of the conditional by `Const[ ] 0[ ]`. In the program so obtained it will then be possible to remove the first component of the parameter since it will never be used. So for the `sum'9B` example we shall proceed in two stages

- apply a *forward* analysis called *constant propagation* to verify that the true-branch always will return 0 — this will enable a program transformation that replaces it by `Const[ ] 0[ ]`, and then
- apply a *backward* analysis called *liveness analysis* to verify that the fixed point expression of the program only needs the second component of its parameter in order to compute the result — this will enable a program transformation that removes the first component of the parameter.

### 7.1 Constant propagation: the type part of $\mathcal{P}$

The purpose of *constant propagation* [1] is to determine whether an expression will evaluate to a constant always and to determine that constant. The analysis will be specified as an interpretation  $\mathcal{P}$  following the pattern described in the previous section. So we shall define

- the meaning  $\mathcal{P}_{t \rightarrow t'}$  of the run-time types  $t \rightarrow t'$ , and
- the meanings  $\mathcal{P}_{\square}$ ,  $\mathcal{P}_{\text{Tuple}}$ , etc of the combinators and the fixed point.

The details of  $\mathcal{P}$  will be very much as for  $\mathcal{S}$  except that  $\mathcal{P}$  will operate on *properties of values* rather than the ‘real’ values computed by  $\mathcal{S}$ . An expression of type `Int` will evaluate to an integer in  $\mathcal{S}$  but in  $\mathcal{P}$  it will evaluate to one of the properties




Figure 2: Hasse diagram for  $\mathbf{Z}^\top$

- $\mathbf{n}$ , an integer, meaning that the ‘real’ value always will be equal to  $\mathbf{n}$  (unless it is undefined),
- $\top$ , meaning that we cannot determine a constant that the ‘real’ value always will be equal to, or
- $\perp$ , meaning that the ‘real’ value always will be undefined.

We shall write  $\mathbf{Z}^\top$  for this selection of properties and a pictorial representation is given in Figure 2. We next define the *type part* of  $\mathcal{P}$  by

$$\mathcal{P}_{t \mapsto t'} = \llbracket t \rrbracket(\mathcal{P}) \rightarrow \llbracket t' \rrbracket(\mathcal{P})$$

so that the meaning of a run-time function will map properties of the argument type  $\llbracket t \rrbracket(\mathcal{P})$  to properties of the result type  $\llbracket t' \rrbracket(\mathcal{P})$ . The properties of values of type  $t$  are defined according to the structure of  $t$  and we have

$$\begin{aligned} \llbracket \mathbf{Int} \rrbracket(\mathcal{P}) &= \mathbf{Z}^\top \\ \llbracket \mathbf{Bool} \rrbracket(\mathcal{P}) &= \mathbf{T}^\top \\ \llbracket t \times t' \rrbracket(\mathcal{P}) &= \llbracket t \rrbracket(\mathcal{P}) \times \llbracket t' \rrbracket(\mathcal{P}) \\ \llbracket t \text{ list} \rrbracket(\mathcal{P}) &= (\llbracket t \rrbracket(\mathcal{P})^*)^\top \end{aligned}$$

Thus the properties of pairs of values are pairs of properties. The properties of lists of values will be lists of properties but may also be the special value  $\top$  that will be used as a property of lists with different lengths. With this definition it turns out that for all run-time types  $t$  the set  $\llbracket t \rrbracket(\mathcal{P})$  will contain an element (ambiguously denoted  $\top$ ) representing that nothing can be determined about the ‘real’ value. It will also contain an element (ambiguously denoted  $\perp$ ) representing that the ‘real’ value definitely will be undefined. In fact each  $\llbracket t \rrbracket(\mathcal{P})$  will be a complete lattice whenever  $t$  is a run-time type. As in the standard semantics we have  $\llbracket \mathbf{Sum}_1 \rrbracket(\mathcal{P}) = \mathbf{Z}^* \rightarrow \mathbf{Z}$  whereas  $\llbracket \mathbf{Sum}_2 \rrbracket(\mathcal{P}) = ((\mathbf{Z}^\top)^*)^\top \rightarrow \mathbf{Z}^\top$  reflecting the difference between the annotations of  $\mathbf{Sum}_1$  and  $\mathbf{Sum}_2$ .

## 7.2 Constant propagation: the expression part of $\mathcal{P}$

The *expression part* of  $\mathcal{P}$  will specify how to operate on these properties. In Table 3 we give a few illustrative clauses. The clauses for  $\square$ , **Tuple**, **Fst**, **Snd**, **Const** and **Id** are as in  $\mathcal{S}$ . The clauses for **Hd**, **Tl**, **Isnll** and **+** are as in  $\mathcal{S}$  but extended to cope with the special properties  $\top$  and  $\perp$ . In the clause for the conditional we distinguish between whether the test evaluates to **true**, **false**,  $\top$  or  $\perp$ . In the first two cases we choose the property of the appropriate branch. If the test evaluates to  $\top$  then the ‘real’ value may be **true**

$\mathcal{P}_{\square} (f,g) p = f(g p)$
$\mathcal{P}_{\text{Tuple}} (f,g) p = \langle f p, g p \rangle$
$\mathcal{P}_{\text{Fst}} \langle p, q \rangle = p$
$\mathcal{P}_{\text{Snd}} \langle p, q \rangle = q$
$\mathcal{P}_{\text{Hd}} p = (p = [] \rightarrow \perp \mid (p = q : p') \rightarrow q \mid (p = \top) \rightarrow \top \mid \perp$
$\mathcal{P}_{\text{Tl}} p = (p = [] \rightarrow \perp \mid (p = q : p') \rightarrow p' \mid (p = \top) \rightarrow \top \mid \perp$
$\mathcal{P}_{\text{Isnll}} p = (p = [] \rightarrow \text{true} \mid (p = q : p') \rightarrow \text{false} \mid (p = \top) \rightarrow \top \mid \perp$
$\mathcal{P}_{\text{Cond}} (f,g,h) p = (f p = \text{true}) \rightarrow g p \mid (f p = \text{false}) \rightarrow h p \mid$ $(f p = \top) \rightarrow (g p) \sqcup (h p) \mid \perp$
$\mathcal{P}_{\text{Const}} f p = f$
$\mathcal{P}_{\text{Id}} p = p$
$\mathcal{P}_{+} \langle p, q \rangle = (p = \perp \vee q = \perp) \rightarrow \perp \mid (p = \top \vee q = \top) \rightarrow \top \mid p + q$
$\mathcal{P}_{\text{fix}}^t f = f(\top) \quad \text{for } t = t' \rightarrow t''$

Table 3: Fragments of the expression part of  $\mathcal{P}$

or it may be false and we shall therefore combine the properties obtained from the two branches using the least upper bound operation  $\sqcup$ . For  $Z^\top$ ,  $p \sqcup q$  is defined to be  $\top$  if  $p$  and  $q$  are distinct integers whereas it is  $p$  if they are equal. Furthermore  $p \sqcup \top = \top \sqcup p = \top$  and  $p \sqcup \perp = \perp \sqcup p = p$  for all  $p$ . Finally, the specification of  $\mathcal{P}_{\text{fix}}^t$  will determine how the fixed point is approximated. We shall take a rather crude approach and assume that all recursive calls in the fixed point return non-constant values (i.e.  $\top$ ).

If we apply  $\mathcal{P}$  to  $\text{sum}'_{9B}$  we then get

$$\llbracket \text{sum}'_{9B} \rrbracket(\mathcal{P}) p = (p = [] \rightarrow 0 \mid (p = q : p') \rightarrow \top \mid (p = \top) \rightarrow \top \mid \perp$$

### 7.3 The enabled program transformation

We cannot deduce very much from  $\llbracket \text{sum}'_{9B} \rrbracket(\mathcal{P})$  about the values arising internally in  $\text{sum}'_{9B}$ . We shall therefore need a *sticky* variant [31,33,20] of the analysis that will record the environments and the arguments supplied to  $\llbracket e \rrbracket(\mathcal{P})$  for the various subexpressions  $e$  of  $\text{sum}'_{9B}$ . For the true-branch  $\text{Fst}[]$  of the conditional we then get that  $\llbracket \text{Fst}[] \rrbracket(\mathcal{P})$  is called with one environment and two different parameters, namely  $\langle 0, [] \rangle$  (if the test evaluates to **true**) and  $\langle 0, \top \rangle$  (if the test evaluates to  $\top$ ). In both cases  $\llbracket \text{Fst}[] \rrbracket(\mathcal{P})$  will evaluate to  $0$  so we see that the analysis gives the expected result.

The constant propagation analysis enables a program transformation called *constant folding* [1]:

*Assume that  $e$  has a subexpression  $e'$  such that during the computation of  $\llbracket e \rrbracket(\mathcal{P}) \text{env}$  all calls of  $\llbracket e' \rrbracket(\mathcal{P})$  with the possible environments return the constant  $v$ .*

*Then  $e$  can safely be changed to the expression  $e''$  obtained by replacing the subexpression  $e'$  by  $\text{Const}[] v[]$ .*

Applying this transformation to  $\text{sum}'_{9B}$  we get the program  $\text{sum}_{9B}^P$ .

```

VAL fix( $\lambda g[] . \text{Cond}(\text{Isnll}[] \square \text{Snd}[], \text{Const}[] 0[],$ 
 $\quad \quad \quad +[] \square \text{Tuple}(\text{Hd}[] \square \text{Snd}[], g \square \text{Tuple}(\text{Fst}[], \text{Tl}[] \square \text{Snd}[])))$ 
 $\quad \square \text{Tuple}(\text{Const}[] 0[], \text{Id}[]))$ 
HAS Int list  $\rightarrow$  Int

```

In a similar way the remaining occurrence of `Fst[]` could be replaced by `Const[] 0[]` but this is not necessary for the next transformation to take place. This transformation only preserves *partial correctness* in the sense that the original program may loop in situations where the transformed program will not loop.

## 7.4 Remarks

The interpretation  $\mathcal{P}$  has been presented as a *non-standard* interpretation. However, in order to be sure that the transformation enabled by  $\mathcal{P}$  is correct it is crucial to make sure that  $\mathcal{P}$  only collects *safe* (i.e. correct) information. Because of decidability issues  $\mathcal{P}$  will in general not collect precise information and the remedy then is to *err on the safe side*.

Safety can be formulated by specifying  $\mathcal{P}$  as an *abstract interpretation*. There are several approaches to how to achieve this. One possibility is to define a concretization function mapping abstract values (i.e. those used by  $\mathcal{P}$ ) to sets of concrete values (i.e. those used by  $\mathcal{S}$ ). Another possibility is to use a representation function mapping concrete values to abstract values. The main use of these functions is to express safety which then must be proved to hold. In our case the proof will amount to verifying safety for all expressions in our notation (using structural induction) and thereby the result will follow for all programs.

We have developed this approach into a *meta-theory* that will ease the specification of the analyses and their safety proofs [33,39]. The idea is that *representation functions* are specified for the run-time base types together with ways to pass between the run-time type constructors. This is sufficient to specify the type part of the interpretation and the representation functions for the various domains. The representation functions may be used to express safety but they also provide means for constructing ‘*best induced analyses*’, i.e. safe analyses that are as precise as possible among all safe interpretations with the same type part. As these analyses are not decidable in general the *meta-theory* also studies certain ‘*expected forms*’ that may be used instead of the ones generated automatically and without compromising safety. It still remains to incorporate sticky analyses into this meta-theory.

To overcome the absence of total correctness in the program transformation enabled by constant propagation one may investigate formulations of constant propagation based upon the approach in [30]. This would allow to let a property like `n` stand for the integer `n` *without* the need to allow undefinedness. The theory, however, becomes much more complex.

## 8 Backward abstract interpretation

In *liveness analysis* [1] we want to know whether or not values are *live*, i.e. may be needed in future computations, or *dead*, i.e. definitely will not be needed in future computations. This analysis differs from the previous one in two important aspects. One is that we are not going to talk about properties of values but rather *properties of the future use of values*. Another difference is that the information about liveness propagates through the program in the *opposite direction* of the flow of control. Therefore liveness analysis is often called a backward analysis whereas constant propagation is called a forward analysis [1].

## 8.1 Liveness analysis: type part of $\mathcal{L}$

The liveness analysis will be specified by an interpretation  $\mathcal{L}$ . A property of a value of type Int (or Bool) will either be

- **dead** if the value is definitely not needed in future computations, or
- **live** if the value may be needed later.

The *type part* of  $\mathcal{L}$  will e.g. have

$$\begin{aligned}\llbracket \text{Int} \rrbracket(\mathcal{L}) &= \{\text{dead}, \text{live}\} \\ \llbracket \text{Bool} \rrbracket(\mathcal{L}) &= \{\text{dead}, \text{live}\} \\ \llbracket t \times t' \rrbracket(\mathcal{L}) &= \llbracket t \rrbracket(\mathcal{L}) \times \llbracket t' \rrbracket(\mathcal{L}) \\ \llbracket t \text{ list} \rrbracket(\mathcal{L}) &= \{\text{dead}, \text{live}\}\end{aligned}$$

So a property of the future use of a pair of values will be a pair of properties, one for each component. For the sake of simplicity, the properties of the future use of lists are defined to be  $\{\text{dead}, \text{live}\}$  so that it will not be possible to see if only parts of the list will be needed in future computations. If a more refined analysis is wanted we may e.g. replace the definition above with  $\llbracket t \text{ list} \rrbracket(\mathcal{L}) = \{L \mid L \subseteq \llbracket t \rrbracket(\mathcal{L})^*\}$  (see e.g. [39]). With the definition above each set  $\llbracket t \rrbracket(\mathcal{L})$  will contain an element  $\perp_t$  (ambiguously denoted **dead**) representing that the ‘real’ value will definitely not be needed in the future computations and it will contain an element  $\top_t$  (ambiguously denoted **live**) representing that the ‘real’ value may be needed in the future computations. Because the analysis is a backward analysis a property of functions will be a function that maps properties in the opposite direction:

$$\mathcal{L}_{t \rightarrow t'} = \llbracket t' \rrbracket(\mathcal{L}) \rightarrow \llbracket t \rrbracket(\mathcal{L})$$

Then  $\llbracket \text{Sum}_1 \rrbracket(\mathcal{L}) = \mathbf{Z}^* \rightarrow \mathbf{Z}$  and  $\llbracket \text{Sum}_2 \rrbracket(\mathcal{L}) = \{\text{dead}, \text{live}\} \rightarrow \{\text{dead}, \text{live}\}$ . The analysis is only specified for the subset of our notation where only first-order run-time functions are allowed. In fact, this was one of the motivations for the second transformation in Subsection 4.3.

## 8.2 Liveness analysis: the expression part of $\mathcal{L}$

Turning to the *expression part* of  $\mathcal{L}$  we shall specify how to operate on the properties above. In Table 4 we give a few illustrative cases. The clause for  $\square$  reflects that functions have to be composed in the opposite order of the standard one. In the clause for **Tuple** we are given a pair  $\langle p, q \rangle$  of properties for the future uses of the result of **Tuple** $(\dots, \dots)$  and we use the least upper bound operation  $\sqcup$  to join the results. On the set  $\{\text{dead}, \text{live}\}$  the operation is defined by  $\text{dead} \sqcup \text{live} = \text{live} \sqcup \text{dead} = \text{live} \sqcup \text{live} = \text{live}$  and  $\text{dead} \sqcup \text{dead} = \text{dead}$ . (This corresponds to the least upper bound in a partially ordered set where  $\text{dead} \sqsubseteq \text{live}$ .) The clauses for **Fst**, **Snd**, **Hd**, **Tl**, **Isnil**, **Const**, **Id** and **+** should be straightforward. For the conditional we simply combine the properties obtained from the test and the two branches. Finally, the meaning of the fixed point is defined to be the least fixed point as in the standard interpretation  $\mathcal{S}$ .

We can now apply  $\mathcal{L}$  to the program  $\text{sum}_{\text{9B}}^{\text{P}}$  and get

$$\llbracket \text{sum}_{\text{9B}}^{\text{P}} \rrbracket(\mathcal{L}) \text{ p} = \text{p}$$

which simply states that if we need the sum of the elements of the list then we also need the list.

$\mathcal{L}_\square (f,g) p = g(f p)$
$\mathcal{L}_{\text{Tuple}} (f,g) \langle p,q \rangle = (f p) \sqcup (g q)$
$\mathcal{L}_{\text{Fst}} p = \langle p, \text{dead} \rangle$
$\mathcal{L}_{\text{Snd}} p = \langle \text{dead}, p \rangle$
$\mathcal{L}_{\text{Hd}} p = p$
$\mathcal{L}_{\text{Tl}} p = p$
$\mathcal{L}_{\text{Isnil}} p = p$
$\mathcal{L}_{\text{Cond}} (f,g,h) p = (f p) \sqcup (g p) \sqcup (h p)$
$\mathcal{L}_{\text{Const}} f p = \text{dead}$
$\mathcal{L}_{\text{Id}} p = p$
$\mathcal{L}_+ p = \langle p, p \rangle$
$\mathcal{L}_{\text{fix}}^t f = \text{FIX} (f) \quad \text{for all } t$

Table 4: Fragments of the expression part of  $\mathcal{L}$

### 8.3 The enabled program transformation

As in the case of constant propagation, we shall need a sticky variant of this analysis. In the invocation of  $\llbracket \text{sum}_{\text{9B}}^P \rrbracket(\mathcal{L}) p$  we then get that

$$\llbracket \text{fix} (\lambda g [ \ ] \dots) \rrbracket(\mathcal{L}) \text{env}$$

only is applied to the property  $p$  and always evaluates to  $\langle \text{dead}, p \rangle$ . This shows that the first component of the argument never is used.

Live variable analysis enables the following transformation on fixed points

Assume that  $e$  has a subexpression  $\text{fix}(\lambda g[t_1 \times t_2 \rightarrow t_3].e')$  such that during the computation of  $\llbracket e \rrbracket(\mathcal{L})\text{env}$  each call of  $\llbracket \text{fix}(\lambda g[t_1 \times t_2 \rightarrow t_3].e') \rrbracket(\mathcal{L})$  with the possible environments return the property  $\langle \text{dead}, p \rangle$  for some  $p$ .

Then  $e$  can safely be changed to  $e''$  obtained by replacing the subexpression

$$\text{fix}(\lambda g[t_1 \times t_2 \rightarrow t_3].e')$$

by

$$\begin{aligned} & \text{fix}(\lambda g'[t_2 \rightarrow t_3].e'[g' \sqcup \text{Snd}[t_1 \times t_2]/g] \sqcup \\ & \quad \text{Tuple}(\text{Fix}[ \ ] \sqcup (\text{Curry Snd}[t_2 \times t_1]), \text{Id}[t_2])) \sqcup \text{Snd}[t_1 \times t_2] \end{aligned}$$

where  $g'$  is a fresh identifier.

Here  $\text{Fix}[ \ ] \sqcup (\text{Curry Snd}[t_2 \times t_1])$  is a function that transforms the argument of type  $t_2$  to an ('unneeded') value ( $\perp$ ) of type  $t_1$ . Applying this transformation to  $\text{sum}_{\text{9B}}^P$  we get the program

```
VAL fix(λg'[ ].Cond(Isnil[ ] ⊔ Snd[ ], Const[ ] 0[ ],
  +[ ] ⊔ Tuple(Hd[ ] ⊔ Snd[ ],
    g' ⊔ Snd[ ] ⊔ Tuple(Fst[ ], Tl[ ] ⊔ Snd[ ]))) ⊔
  Tuple(Fix[ ] ⊔ (Curry Snd[ ]), Id[ ]))
  ⊔ Snd[ ] ⊔ Tuple(Const[ ] 0[ ], Id[ ])
HAS Int list → Int
```

At this stage it is worthwhile to apply a few *algebraic transformations* (Section 5)

$$\begin{aligned}
& \text{Snd}[ ] \sqsubseteq \text{Tuple}(e, e') \Rightarrow e' \\
& e \sqsubseteq \text{Id}[ ] \Rightarrow e \\
& \text{Cond}(e_1, e_2, e_3) \sqsubseteq e \Rightarrow \text{Cond}(e_1 \sqsubseteq e, e_2 \sqsubseteq e, e_3 \sqsubseteq e) \\
& (\text{Const}[ ] \ 0[ ]) \sqsubseteq e \Rightarrow \text{Const}[ ] \ 0[ ] \\
& \text{Tuple}(e_1, e_2) \sqsubseteq e \Rightarrow \text{Tuple}(e_1 \sqsubseteq e, e_2 \sqsubseteq e)
\end{aligned}$$

and the resulting program is  $\text{sum}_{\text{SB}}^L$

```

VAL fix( $\lambda g'[ ]$ .Cond(Isnil[ ], Const[ ] 0[ ], +[ ]  $\sqsubseteq$  Tuple(Hd[ ], g'  $\sqsubseteq$  Tl[ ])))
HAS Int list  $\rightarrow$  Int

```

These transformations preserve the total correctness of the program.

## 8.4 Remarks

As in the previous section the analysis  $\mathcal{L}$  is only specified as a non-standard interpretation so nothing has been said about its safety properties. Unfortunately only *some* of the techniques mentioned in Section 7 can be applied to  $\mathcal{L}$ . The main reason is that  $\mathcal{L}$  is not concerned with properties of values but, as we mentioned earlier, properties of their future use. Thus it does not make sense to define representation functions as before. It is worth noticing that it is not the fact that  $\mathcal{L}$  is a backward analysis that limits the applicability of the previous approach – indeed there are forward analyses (e.g. available expressions [1]) that cannot be proved safe with these techniques neither. We refer to [39] for further details about the correctness of  $\mathcal{L}$ .

It is also worth pointing out the relationship to the ‘strictness continuations’ used in [43] and to the ‘evaluators’ of [6]. There also appears to be a close relationship to the ‘demand analysis’ of [4].

## 9 Code generation

One of the motivations behind the 2-level notation is that we should only generate code for run-time computations, not compile-time computations. We shall illustrate this by showing how to reinterpret the combinators so as to specify code generation for a simple abstract machine.

### 9.1 An abstract machine

The abstract machine we shall consider is closely related to the categorical abstract machine [9] and the abstract machine of [32]. The *configurations* of the machine has two components

- a *control stack* CS containing the code to be executed, and
- an *evaluation stack* ST containing the intermediate results.

In addition there is an *environment* ENV containing named pieces of code. During the execution of a program the configurations (CS,ST) will be modified whereas the environment will remain unchanged. We shall write

(ENTER:CS,v:ST) → <sub>ENV</sub> (CS,v:v:ST)
(SWITCH:CS,v:u:ST) → <sub>ENV</sub> (CS,u:v:ST)
(TUPLE:CS,v:u:ST) → <sub>ENV</sub> (CS,⟨v,u⟩:ST)
(FST:CS,⟨v,u⟩:ST) → <sub>ENV</sub> (CS,v:ST)
(SND:CS,⟨v,u⟩:ST) → <sub>ENV</sub> (CS,u:ST)
(HD:CS,[v:u]:ST) → <sub>ENV</sub> (CS,v:ST)
(TL:CS,[v:u]:ST) → <sub>ENV</sub> (CS,u:ST)
(ISNIL:CS,[ ]:ST) → <sub>ENV</sub> (CS,true:ST)
(ISNIL:CS,[v:u]:ST) → <sub>ENV</sub> (CS,false:ST)
(BRANCH(C,C'):CS,true:ST) → <sub>ENV</sub> (C:CS,ST)
(BRANCH(C,C'):CS,false:ST) → <sub>ENV</sub> (C':CS,ST)
(CONST(v):CS,u:ST) → <sub>ENV</sub> (CS,v:ST)
(SKIP:CS,ST) → <sub>ENV</sub> (CS,ST)
(ADD:CS,⟨v,u⟩:ST) → <sub>ENV</sub> (CS,(v+u):ST)
(CALL(l):CS,ST) → <sub>ENV</sub> (ENV(l):CS,ST)
(DELAY(C):CS,v:ST) → <sub>ENV</sub> (CS,{C,v}:ST)
(RESUME:CS,{C,v}:ST) → <sub>ENV</sub> (C:CS,v:ST)
(RESUME:CS,v:ST) → <sub>ENV</sub> (CS,v:ST)    otherwise

Table 5: Abstract machine instructions

$$(CS,ST) \rightarrow_{ENV} (CS',ST')$$

if one execution step of the machine will change the configuration from  $(CS,ST)$  to  $(CS',ST')$ .

In Table 5 we list some of the *instructions* of the machine together with their operational meaning. The control stack is a list of instructions and we shall (ambiguously) write ‘:’ for appending two lists and for prepending an element to a list. The evaluation stack is a list of values and we shall distinguish between three different kinds of values

- primitive values such as the truth values and the integers,
- composite values such as pairs (denoted by  $\langle \dots, \dots \rangle$ ) and lists (denoted by  $[\dots]$ ), and
- delay closures (denoted by  $\{ \dots, \dots \}$ ) containing a list of instructions and a value.

A *delay closure*  $\{C,v\}$  is used to delay the execution of the code  $C$  on the stack with  $v$  on its top. The instruction  $DELAY(\dots)$  constructs a delay closure whereas  $RESUME$  initiates the execution of the code of a delay closure on the given argument.

The overall operation of the machine upon input  $v$  and with program  $C$  is to start in the configuration  $(C,v:nil)$ . It will then repeatedly execute the instructions of  $C$  and update the configuration accordingly. If at some state there is no next configuration then the machine stops. If one of the stacks of the configuration becomes empty or it calls a name not defined in the environment then the machine also stops. It is of course also possible that the machine will continue executing instructions forever. However, if it stops and the configuration has the form  $(nil,u:nil)$  then we say that the machine has terminated with output  $u$ . In all other cases where the machine stops we shall say that it has terminated with failure.

As an example consider the program  $ENTER:TUPLE:ADD$ . Assuming that the input is 7, the initial configuration will be



(ENTER:TUPLE:ADD:nil,7:nil)

Execution of the program with the empty environment  $\emptyset$  will modify the configuration as follows

(ENTER:TUPLE:ADD:nil,7:nil)  
 $\rightarrow_{\emptyset}$  (TUPLE:ADD:nil,7:7:nil)  
 $\rightarrow_{\emptyset}$  (ADD:nil,⟨7,7⟩:nil)  
 $\rightarrow_{\emptyset}$  (nil,14:nil)

The machine has now terminated with output 14. The same effect is obtained by executing the program DELAY(ENTER:TUPLE:ADD):RESUME

(DELAY(ENTER:TUPLE:ADD):RESUME:nil,7:nil)  
 $\rightarrow_{\emptyset}$  (RESUME:nil,{ENTER:TUPLE:ADD:nil,7}:nil)  
 $\rightarrow_{\emptyset}$  (ENTER:TUPLE:ADD:nil,7:nil)  
 $\rightarrow_{\emptyset}$  (TUPLE:ADD:nil,7:7:nil)  
 $\rightarrow_{\emptyset}$  (ADD:nil,⟨7,7⟩:nil)  
 $\rightarrow_{\emptyset}$  (nil,14:nil)

## 9.2 Coding interpretation: the type part of $\mathcal{K}$

Recall that to specify the semantics it is sufficient to specify the meaning of the run-time function types, the combinators and the compile-time fixed point. Let **Code** be the set of instruction sequences for the stack machine. The *type part* of the interpretation  $\mathcal{K}$  will then have

$$\mathcal{K}_{t \rightarrow t'} = \mathbf{Code}$$

for all types  $t \rightarrow t'$  since we want to generate a sequence of instructions for any run-time function. Note that we do not use structural induction to define the meaning of the run-time types as in the previous sections. As an example, we have  $\llbracket \text{Sum}_1 \rrbracket(\mathcal{K}) = \mathbf{Z}^* \rightarrow \mathbf{Z}$  and  $\llbracket \text{Sum}_2 \rrbracket(\mathcal{K}) = \mathbf{Code}$ . Perhaps more interestingly we have  $\llbracket \text{Reduce}_9 \rrbracket(\mathcal{K}) = \mathbf{Code} \rightarrow \mathbf{Code}$  reflecting that given the code for the function argument of **reduce** we obtain the code for the specialized function. Intuitively,  $\mathbf{Code} \rightarrow \mathbf{Code}$  is the type of code fragments with holes for where the argument code should be inserted.

## 9.3 Coding interpretation: the expression part of $\mathcal{K}$

The *expression part* of  $\mathcal{K}$  will specify how to generate code for the combinators and the fixed point. So e.g.  $\mathcal{K}_{\text{Fst}}$  will specify the code to be generated for **Fst** [ ] and  $\mathcal{K}_{\square}$  will specify how to compose the two pieces of code for the arguments of  $\square$ .

The clauses of Table 6 reflect that our language has a non-strict semantics in that the execution of the various operations are postponed (using the DELAY instruction) until they are really needed. We use a simple-minded approach where everything is enclosed in a delay closure whenever the result might possibly not be needed. As an example, the code for  $e_1 \square e_2$  will contain the code for  $e_1$  but encapsulated in a DELAY instruction because it is not known whether it needs to be executed. On the other hand, the code for  $e_2$  needs to be executed in order to get a result and is therefore not encapsulated. Similarly, we do not know whether the subexpressions  $e_1$  and  $e_2$  of **Tuple**( $e_1, e_2$ ) need

$\mathcal{K}_{\square}(\mathbf{f}, \mathbf{g}) = \text{DELAY}(\mathbf{g}) : \mathbf{f}$
$\mathcal{K}_{\text{Tuple}}(\mathbf{f}, \mathbf{g}) = \text{ENTER} : \text{DELAY}(\mathbf{g}) : \text{SWITCH} : \text{DELAY}(\mathbf{f}) : \text{TUPLE}$
$\mathcal{K}_{\text{Fst}} = \text{RESUME} : \text{FST} : \text{RESUME}$
$\mathcal{K}_{\text{Snd}} = \text{RESUME} : \text{SND} : \text{RESUME}$
$\mathcal{K}_{\text{Hd}} = \text{RESUME} : \text{HD} : \text{RESUME}$
$\mathcal{K}_{\text{Tl}} = \text{RESUME} : \text{TL} : \text{RESUME}$
$\mathcal{K}_{\text{Isnil}} = \text{RESUME} : \text{ISNIL}$
$\mathcal{K}_{\text{Cond}}(\mathbf{f}, \mathbf{g}, \mathbf{h}) = \text{ENTER} : \mathbf{f} : \text{BRANCH}(\mathbf{g}, \mathbf{h})$
$\mathcal{K}_{\text{Const}} \mathbf{f} = \text{CONST}(\mathbf{f})$
$\mathcal{K}_{\text{Id}} = \text{RESUME}$
$\mathcal{K}_{+} = \text{RESUME} : \text{ENTER} : \text{SND} : \text{RESUME} :$ $\text{SWITCH} : \text{FST} : \text{RESUME} : \text{TUPLE} : \text{ADD}$
$\mathcal{K}_{\text{fix}}^t \mathbf{f} = \text{CALL}(f)$ where $t$ is $t' \rightarrow t''$ and ENV is extended with $\text{ENV}(f) = \mathbf{f}(\text{CALL}(f))$ and where $f$ is a fresh identifier

Table 6: Fragments of the expression part of  $\mathcal{K}$

to be evaluated and therefore the corresponding fragments of code are enclosed in DELAY instructions.

Whenever we want to recover a value that may be a delay closure we use the instruction RESUME. As an example, the code for  $\mathcal{K}_{\text{Fst}}$  reflects that we must first evaluate the argument to a pair and then we select the first component. The actual pair could very well be a pair of delay closures and in that case the result will be further evaluated. The code for  $\mathcal{K}_{+}$  shows how the RESUME instruction is used to force the evaluation of the argument to a pair and then to force the evaluation of each of the components to an integer such that the addition instruction can be applied.

Finally, the code generated for the fixed point is simply  $\text{CALL}(f)$  where  $f$  is a labeled piece of code specified in the environment. The argument  $\mathbf{f}$  of  $\mathcal{K}_{\text{fix}}^t$  is intuitively a code fragment with *holes* in it for where the recursive calls have to be inserted, and in  $\mathbf{f}(\text{CALL}(f))$  the holes have been filled with  $\text{CALL}(f)$ .

We can now apply  $\mathcal{K}$  to the program  $\text{sum}_{\text{9B}}^L$  and we obtain the code

$\text{CALL}(\text{sum})$

The environment ENV will associate the label *sum* with the code

```

ENTER:RESUME:ISNIL:
BRANCH(CONST(0),
  DELAY(ENTER:
    DELAY(DELAY(RESUME:TL:RESUME):CALL(sum)):
    SWITCH:
    DELAY(RESUME:HD:RESUME):
    TUPLE):
  RESUME:ENTER:SND:RESUME:SWITCH:FST:RESUME:TUPLE:ADD)

```

## 9.4 Peephole optimization

The code generated by the coding interpretation  $\mathcal{K}$  is rather inefficient and we shall therefore show that a few optimization rules can be used to improve the quality of the code.

We shall use the following transformation rules, often called peep-hole optimizations, in order to reduce the number of delay closures:

$$\text{DELAY}(C):\text{RESUME} \Rightarrow C$$

$$\begin{aligned} &\text{ENTER}:C_2:\text{SWITCH}:C_1:\text{TUPLE}: \text{ENTER}:\text{SND}:C'_2:\text{SWITCH}:\text{FST}:C'_1:\text{TUPLE} \\ &\Rightarrow \text{ENTER}:C_2:C'_2:\text{SWITCH}:C_1:C'_1:\text{TUPLE} \end{aligned}$$

Applying these transformations to the code of  $\text{ENV}(\text{sum})$  we get

```
ENTER:RESUME:ISNIL:
BRANCH(CONST(0),
      ENTER:DELAY(RESUME:TL:RESUME):CALL(sum)
      SWITCH:RESUME:HD:RESUME:TUPLE:ADD)
```

Before each recursive call of the function the argument has been encapsulated in a delay closure. This is because the code generation reflects the non-strict semantics and in the next section we shall see how a strictness analysis can be used to avoid generating these delay closures.

## 9.5 Remarks

The abstract machine has been chosen so as to make code generation straightforward. In our previous work [32,35] we have used a more complicated machine where the code is linearized and the flow of control is modified by jumping between (uniquely generated) labels. However, the non-standard interpretation can also be used to generate code for full-fledged abstract machines like FAM [8] and the G-machine [44].

The semantics of the abstract machine is specified operationally and when proving the correctness of the code generation we therefore have to relate a denotational semantics (specified by  $\mathcal{S}$ ) to this operational semantics. Such a correctness proof is given in [35] for the slightly more complicated machine mentioned above and with a strict semantics (and thereby a code generation avoiding the generation of delay closures).

When the actions of the abstract machine are specified operationally it may be more convenient to prove the correctness of the code generation with respect to an *operational* semantics. This has been done in [9,13] for the categorical abstract machine and it is interesting to note that the overall structure of that proof and the one in [35] have much in common.

## 10 Optimization of the code generation

A closer inspection of the program `sum` shows that it is *strict* in its parameter. This is not reflected in the code generated in the previous section so we shall want to improve the code generation. The idea is now

- to *specify* a strictness analysis, and

- to *modify* the coding interpretation to generate special code for strict functions.

A strictness analysis is just another *forward abstract interpretation*. Thus we shall proceed very much as in Section 7 when we studied the constant propagation analysis.

## 10.1 Strictness analysis

The purpose of strictness analysis [28,29,5,18,34,48] is to determine whether an expression always will need its argument, i.e. whether it needs to be evaluated. The technical details of how  $\mathcal{T}$  is specified are very much as those of  $\mathcal{P}$  (and  $\mathcal{S}$ ) except that the domains of properties are even simpler. An expression of type Int will now have one of two properties

- **1** meaning that the value may be defined or undefined (much as  $\top$  in  $\llbracket \text{Int} \rrbracket(\mathcal{P})$ ), or
- **0** meaning that the value is undefined (much as  $\perp$  in  $\llbracket \text{Int} \rrbracket(\mathcal{P})$ ).

We shall write  $\Delta$  for this set of properties. The *type part* of  $\mathcal{T}$  will now define

$$\mathcal{T}_{t \mapsto t'} = \llbracket t \rrbracket(\mathcal{T}) \rightarrow \llbracket t' \rrbracket(\mathcal{T})$$

so that the meaning of a run-time function will map properties of the argument type  $\llbracket t \rrbracket(\mathcal{T})$  to properties of the result type  $\llbracket t' \rrbracket(\mathcal{T})$ . The properties of values of type  $t$  are defined according to the structure of  $t$

$$\llbracket \text{Int} \rrbracket(\mathcal{T}) = \Delta$$

$$\llbracket \text{Bool} \rrbracket(\mathcal{T}) = \Delta$$

$$\llbracket t \times t' \rrbracket(\mathcal{T}) = \llbracket t \rrbracket(\mathcal{T}) \times \llbracket t' \rrbracket(\mathcal{T})$$

$$\llbracket t \text{ list} \rrbracket(\mathcal{T}) = \Delta$$

with the exception that the list types are interpreted as the simple domain  $\Delta$ . This means that the analysis will not be sufficiently precise to determine whether some function is strict in, say, the head of the list or the spine of the list. This restriction is merely for the sake of simplicity.

The *expression part* of  $\mathcal{T}$  will specify how to compute with these properties. In Table 7 we give some of the more interesting cases. The operation  $\mathbf{p} \vee \mathbf{q}$  is defined on  $\Delta$  to be **1** if either  $\mathbf{p}$  or  $\mathbf{q}$  is **1** and otherwise **0**. In general it is the least upper bound operation on the complete lattice  $\llbracket t \rrbracket(\mathcal{T})$  for  $t$  a run-time type. Similarly,  $\mathbf{p} \wedge \mathbf{q}$  is defined on  $\Delta$  to be **0** if either  $\mathbf{p}$  or  $\mathbf{q}$  is **0** and otherwise **1**. In general it is the greatest lower bound operation. In the clause for  $\mathcal{T}_{\text{fix}}^t$  we exploit that the complete lattice has a greatest element  $\top$  (ambiguously denoted **1**).

If we apply  $\mathcal{T}$  to  $\text{sum}_{\text{9B}}^L$  we then get the expected

$$\llbracket \text{sum}_{\text{9B}}^L \rrbracket(\mathcal{T}) \mathbf{p} = \mathbf{p}$$

and taking  $\mathbf{p}$  to be **0** this shows that  $\text{sum}_{\text{9B}}^L$  is strict in its argument.

$\mathcal{T}_{\square} (f,g) p = f(g p)$
$\mathcal{T}_{\text{Tuple}} (f,g) p = \langle f p, g p \rangle$
$\mathcal{T}_{\text{Fst}} \langle p, q \rangle = p$
$\mathcal{T}_{\text{Snd}} \langle p, q \rangle = q$
$\mathcal{T}_{\text{Hd}} p = p$
$\mathcal{T}_{\text{Tl}} p = p$
$\mathcal{T}_{\text{Isnil}} p = p$
$\mathcal{T}_{\text{Cond}} (f,g,h) p = (f p) \wedge ((g p) \vee (h p))$
$\mathcal{T}_{\text{Const}} f p = 1$
$\mathcal{T}_{\text{Id}} p = p$
$\mathcal{T}_{+} \langle p, q \rangle = p \wedge q$
$\mathcal{T}_{\text{fix}}^t f = f(1) \quad \text{for } t=t' \rightarrow t''$

Table 7: Fragments of the expression part of  $\mathcal{T}$

## 10.2 Modifying the coding interpretation

The idea is now to *combine* the strictness information and the code generation. We shall do that by defining an interpretation  $\mathcal{C}$  that will compute the strictness information at the same time as it computes the code to be generated. The *type part* of  $\mathcal{C}$  will have

$$\mathcal{C}_{t \rightarrow t'} = (\llbracket t \rrbracket(\mathcal{T}) \rightarrow \llbracket t' \rrbracket(\mathcal{T})) \times \text{Code}$$

where **Code** is as in the previous section. Thus the meaning of a run-time function will be a pair where the first component contains the strictness information and the second component is the code for the function.

The *expression part* of the interpretation will combine the clauses for the strictness analysis with those for the coding interpretation. In the clause specifying the code to be generated for  $e_1 \square e_2$  we use that if  $e_1$  is strict then  $e_2$  *has to* be computed and so we can dispense with the DELAY instruction. The clauses are shown in Figure 8.

We can now apply the coding interpretation to the program  $\text{sum}_{\text{B}}^{\text{T}}$  and we get the code

CALL(*sum*)

The environment ENV will now associate the label *sum* with the code

```

ENTER:RESUME:ISNIL:
BRANCH(CONST(0),
    ENTER:DELAY(RESUME:TL:RESUME:CALL(sum)):
    SWITCH:DELAY(RESUME:HD:RESUME)
    TUPLE:
    RESUME:ENTER:SND:RESUME:SWITCH:FST:RESUME:TUPLE:ADD)

```

We can apply the peephole optimizations of the previous section together with the transformations

```

TUPLE:RESUME  $\Rightarrow$  TUPLE
ENTER:C:C1:SWITCH:C:C2  $\Rightarrow$  C:ENTER:C1:SWITCH:C2

```

and the result is

$\mathcal{C}_{\square} ((f',f),(g',g)) = (\lambda p.f'(g' p),$ $(f' 0 = 0) \rightarrow g:f \mid \text{DELAY}(g):f)$
$\mathcal{C}_{\text{Tuple}} ((f',f),(g',g)) = (\lambda p.\langle f' p, g' p \rangle,$ $\text{ENTER:DELAY}(g):\text{SWITCH:DELAY}(f):\text{TUPLE})$
$\mathcal{C}_{\text{Fst}} = (\lambda \langle p,q \rangle.p, \text{RESUME:FST:RESUME})$
$\mathcal{C}_{\text{Snd}} = (\lambda \langle p,q \rangle.q, \text{RESUME:SND:RESUME})$
$\mathcal{C}_{\text{Hd}} = (\lambda p.p, \text{RESUME:HD:RESUME})$
$\mathcal{C}_{\text{Tl}} = (\lambda p.p, \text{RESUME:TL:RESUME})$
$\mathcal{C}_{\text{Isnil}} = (\lambda p.p, \text{RESUME:ISNIL})$
$\mathcal{C}_{\text{Cond}} ((f',f),(g',g),(h',h)) = (\lambda p.(f' p) \wedge ((g' p) \vee (h' p)),$ $\text{ENTER:f:BRANCH}(g,h))$
$\mathcal{C}_{\text{Const}} f = (\lambda p.1, \text{CONST } f)$
$\mathcal{C}_{\text{Id}} = (\lambda p.p, \text{RESUME})$
$\mathcal{C}_{+} = (\lambda \langle p,q \rangle.p \wedge q,$ $\text{RESUME:ENTER:SND:RESUME:SWITCH:FST:RESUME:TUPLE:ADD})$
$\mathcal{C}_{\text{fix}}^t F = \text{let } (f',f) = F(1, \text{CALL}(f))$ $\text{in } (f', \text{CALL}(f) \text{ where } \text{ENV}(f) = f)$ $\text{for } t=t' \rightarrow t'' \text{ and where } f \text{ is a fresh identifier}$

Table 8: Fragments of the expression part of  $\mathcal{C}$

```

ENTER:RESUME:ISNIL:
BRANCH(CONST(0),
  RESUME:ENTER:TL:RESUME:CALL(sum):
  SWITCH:HD:RESUME:TUPLE:ADD)

```

### 10.3 Remarks

The code generation above can be further improved by introducing a test for *totality* [28]. In the case of  $e_1 \square e_2$  we can then dispense with the `DELAY` instruction both in the case where  $e_2$  is total and in the case where  $e_1$  is strict. Furthermore, the `DELAY` instructions encapsulating the code generated for  $e_i$  for `Tuple`( $e_1, e_2$ ) can be avoided if  $e_i$  is total. However, such an improvement will not affect the code generated for our example program because the relevant functions will not be total. The formulation of the totality analysis is along the lines of the strictness analysis.

In [43] it is shown how even better code can be generated. One scheme amounts to using ‘strictness continuations’ and another scheme amounts to using ‘evaluation degrees’.

We have taken the approach of *combining* two interpretations into one and then use results obtained by one part when defining the other part. An alternative approach would be to *annotate* the programs with strictness information and then extend the coding interpretation of Section 9 to cope with the annotated constructs.

## 11 Conclusion

To assess the applicability of our approach to the implementation of functional languages one will need to investigate the following aspects:

- *efficiency*: the extent to which the code is comparable in efficiency (e.g. time or space) to that of code produced by other means,
- *correctness*: that the generated code behaves as expected, i.e. as the program being compiled, and
- *automation*: the extent to which the analyses and transformations can be performed without user interaction.

## 11.1 Efficiency

To obtain efficient code we have borrowed techniques for program analysis (or data flow analysis) and program transformation (or code optimization) from traditional compiler technology [1]. The functional programming style introduces an *implicit* distinction between binding times as higher-order functions typically are supplied with some but not all of their arguments. This kind of distinction is often explicit in imperative languages and the efficient implementations of these languages rely on this distinction. However, using a *binding time analysis* we can make this distinction *explicit* in a functional language as well and then use it in the analysis and transformation of the programs. Following our approach the transformations at the source level may be grouped in four categories:

- *Transformations aiding the binding time analysis*: The purpose is to get an appropriate distinction between binding times. A similar problem has been discussed in the context of semantic specifications in [35,22].
- *Transformations at the compile-time level*: The purpose is to carry out some of the compile-time computations once and for all. These transformations are closely related to the partial evaluation described in [14] and a restricted form of the transformations described in the unfold/fold framework of [7].
- *Transformations at the run-time level that are universally applicable*: The run-time level of our notation is converted into combinator notation and we apply transformations similar to those developed for FP [2,3,16].
- *Transformations at the run-time level that are enabled by certain program analyses*: The enabling analyses are expressed as abstract interpretations [33,39] and they may express properties obtained by forward or backward analyses [1]. The use of transformations enabled by program analyses is discussed in [31,1].

The translation of the source program into target code is specified by a rather simple scheme and to improve the efficiency of the code generated we need to analyse and transform the program:

- *Analyse the program*: The purpose is here to collect information that can be used to select more appropriate code. The analyses may be specified as abstract interpretations (as in Section 10).
- *Transform the generated code*: We may use *peephole optimizations* to remove some of the superfluous instructions generated.

The example developed throughout this paper demonstrates that substantial improvements in the overall performance can be obtained using these techniques.

## 11.2 Correctness

Correctness involves

- validating that the (simple-minded) *code generation* scheme is correct,
- validating that the results of the *program analyses* are correct, and
- validating that the *program transformations* preserve the semantics of programs.

As we have seen the analyses and transformations can be applied both at the source level and at the target level, i.e. one can perform analyses and transformations on the original functional program and on the code generated for it. As in traditional compiler technology this means that the actual code generation is often kept rather simple-minded because the program may have been transformed so as to facilitate code generation and any shortcomings in the generated code may be alleviated by subsequent transformations. This is important when one is arguing about the correctness issues since it reduces the number of special cases to consider.

From the point of view of compiler construction one is particularly interested in program transformations that are not semantics-preserving in general but only in certain contexts. Information about the ‘contexts’ is provided by program analyses and this motivates a close relationship between the program transformations and their enabling program analyses. An underlying principle of abstract interpretation is to approximate perfect information by ‘erring on the safe side’. A number of abstract interpretations have been developed for functional languages [5,18,19,34] and a framework for developing abstract interpretations is described in [33,35,39]. The theoretical aspects of using information obtained from abstract interpretation to enable program transformations and to improve the code generation is discussed in [31,20].

## 11.3 Automation

In an implementation of a programming language it is important that the user does not have to be concerned with the choice of analyses and transformations to be performed and the order in which to perform them. At the current stage our approach is certainly not amenable to automation. We have experimented with binding time analyses, combinator introduction, algebraic transformations and the concept of parameterized semantics in our *PSI-system* and we expect that the system can be extended to include *specific* analyses and transformations. However, the range of analyses and transformations that one would want to perform is so broad that much more research is needed to determine which analyses and transformations should be built into an intelligent compiler for a functional language and in which order they should be performed. We hope to look further into this using more general but semi-automatic systems as e.g. [15,12].

## Acknowledgement

This work has been supported by the Danish Technical Research Council and the Danish Natural Science Research Council. A preliminary version of the earlier sections of this paper appeared as [40].



# References

- [1] A.V.Aho, R.Sethi, J.D.Ullman: *Compilers – Principles, Techniques and Tools*, Addison-Wesley, 1986.
- [2] J.Backus: Can Programming be Liberated from the von Neumann Style? A Functional Style and its Algebra of Programs, *Communications of the ACM* **21**, 1978.
- [3] F.Bellegarde: Rewriting Systems on FP Expressions to Reduce the Number of Sequences Yielded, *Science of Computer Programming* **6**, 1986.
- [4] B.Bjerner, S.Holmström: A compositional approach to time analysis of first order lazy functional programs, *Proceedings Functional Programming Languages and Computer Architectures*, ACM Press, 1989.
- [5] G.L.Burn, C.Hankin, S.Abramsky: Strictness analysis for higher-order functions, *Science of Computer Programming* **7**, 1986.
- [6] G.L.Burn: Evaluation Transformers – a model for the parallel evaluation of functional languages, *Proceedings of the Functional Programming Languages and Computer Architectures Conference*, Springer Lecture Notes in Computer Science **274**, 1987.
- [7] R.M.Burstall, J.Darlington: A Transformation System for Developing Recursive Programs, *Journal of the ACM* **24**, 1977.
- [8] L.Cardelli: The functional abstract machine, Bell Labs. Technical Report TR-107, 1983.
- [9] G.Cousineau, P.-L. Curien, M.Mauny: The Categorical Abstract Machine, *Science of Computer Programming* **8**, 1987.
- [10] P.-L. Curien: *Categorical Combinators, Sequential Algorithms and Functional Programming*, Pitman, London, 1986.
- [11] L.Damas: Type Assignment in Programming Languages, Ph.D.-thesis CST-33-85, University of Edinburgh, Scotland, 1985.
- [12] J.Darlington, H.Pull: A program development methodology based on a unified approach to execution and transformation, *Partial Evaluation and Mixed Computation*, D.Bjørner, A.P.Ershov and N.D.Jones (eds.), North-Holland, 1988.
- [13] J.Despeyroux: Proof of translation in natural semantics, *Symposium on Logic in Computer Science*, 1986.
- [14] A.P.Ershov: Mixed Computation: Potential Applications and Problems for Study, *Theoretical Computer Science* **18**, 1982.
- [15] M.S.Feather: A System for Assisting Program Transformation, *ACM Transactions on Programming Languages and Systems* **4**, 1982.
- [16] P.G.Harrison: Linearisation: An Optimisation for Nonlinear Functional Programs, *Science of Computer Programming* **10**, 1988.
- [17] J.Hughes: Supercombinators: a new Implementation Method for Applicative Languages, *Proc. of 1982 ACM Conf. on LISP and Functional Programming*, 1982.

- [18] J.Hughes: Strictness detection in non-flat domains, *Proc. Programs as Data Objects*, Springer Lecture Notes in Computer Science **217**, 1986.
- [19] J.Hughes: Backwards Analysis of Functional Programs, *Partial Evaluation and Mixed Computation*, D.Bjørner, A.P.Ershov and N.D.Jones (eds.), North-Holland, 1988.
- [20] P.Hudak, J.Young: A Collecting Interpretation of Expressions (without Powerdomains), *Proceedings of the 15th ACM Symposium on Principles of Programming Languages*, 1988.
- [21] N.D.Jones & P.Sestoft & H.Søndergaard: An experiment in Partial Evaluation: The Generation of a Compiler Generator, *Proceedings of Rewriting Techniques and Applications*, Springer Lecture Notes in Computer Science **202**, 1985.
- [22] U.Jørring, W.L.Scherlis: Compilers and Staging Transformations, *Proc. 13th ACM Symposium on Principles of Programming Languages*, 1986.
- [23] J.Lambek, P.J.Scott: *Introduction to Higher Order Categorical Logic*, Cambridge Studies in Advanced Mathematics **7**, 1986.
- [24] R.Milner: The Standard ML core language, *Proc. 1984 ACM Conference on LISP and Functional Programming*, 1984.
- [25] R.Milner: A Theory of Type Polymorphism in Programming, *Journal of Computer Systems* **17**, 1978.
- [26] T.Mogensen: Binding Time Analysis for Higher Order Polymorphically Typed Languages, *TAPSOFT 1989*, Springer Lecture Notes in Computer Science, 1989.
- [27] M.Montenyohl, M.Wand: Correct flow analysis in continuation semantics, *Proc. 15th ACM Symposium on Principles of Programming Languages*, 1988.
- [28] A.Mycroft: Abstract Interpretation and Optimizing Transformations for Applicative Programs, Ph.D.-thesis CTS-15-81, University of Edinburgh, Scotland, 1981.
- [29] A.Mycroft: The theory and practice of transforming call-by-need into call-by-value, *Proc. 4th International Symposium on Programming*, Springer Lecture Notes in Computer Science **83**, 1980.
- [30] A.Mycroft, F.Nielson: Strong abstract interpretation using powerdomains, *Proceedings ICALP 1983*, Lecture Notes in Computer Science **154**, 1983.
- [31] F.Nielson: Program Transformations in a Denotational Setting, *ACM Transactions on Programming Languages and Systems* **7**, 1985.
- [32] H.R.Nielson, F.Nielson: Semantics Directed Compiling for Functional Languages, *Proc. 1986 ACM Conference on LISP and Functional Programming*, 1986.
- [33] F.Nielson: Towards a Denotational Theory of Abstract Interpretation, *Abstract Interpretation of Declarative Languages*, S.Abramsky and C.Hankin (eds.), Ellis Horwood, 1987.
- [34] F.Nielson: Strictness Analysis and Denotational Abstract Interpretation (Extended Abstract), *ACM Conference on Principles of Programming Languages*, 1987.
- [35] F.Nielson, H.R.Nielson: Two-Level Semantics and Code Generation, *Theoretical Computer Science* **56**, 1988.

- [36] F.Nielson: A Formal Type System for Comparing Partial Evaluators, *Partial Evaluation and Mixed Computation*, North-Holland, 1988.
- [37] H.R.Nielson, F.Nielson: Automatic Binding Time Analysis for a Typed  $\lambda$ -calculus, *Science of Computer Programming* **10**, 1988. Also see *Proc. 15th ACM Symposium on Principles of Programming Languages*, 1988.
- [38] F.Nielson, H.R.Nielson: 2-level  $\lambda$ -lifting, *Proc. ESOP 1988*, Springer Lecture Notes in Computer Science **300**, 1988.
- [39] F.Nielson: Two-Level Semantics and Abstract Interpretation, *Theoretical Computer Science — Fundamental Studies* **69**, 1989.
- [40] H.R.Nielson, F.Nielson: Transformations on higher-order functions, *Proceedings of the Functional Programming and Computer Architecture Conference '89*, ACM Press, 1989.
- [41] H.R.Nielson, F.Nielson: Functional Completeness of the Mixed  $\lambda$ -Calculus and Combinatory Logic, *Theoretical Computer Science* **70**, 1990.
- [42] H.R.Nielson, F.Nielson: Eureka definitions for free! or Disagreement points for fold/unfold transformations, *Proceedings ESOP '90*, Springer Lecture Notes in Computer Science **432**, 1990.
- [43] H.R.Nielson, F.Nielson: Context information for lazy code generation, *Proceedings of the LISP and Functional Programming Conference Conference '90*, ACM Press, 1990.
- [44] S.Peyton Jones: *The Implementation of Functional Programming Languages*, Prentice Hall, 1987.
- [45] D.A.Schmidt: Static Properties of Partial Evaluation, *Partial Evaluation and Mixed Computation*, North-Holland, 1988.
- [46] D.Turner: A New Implementation Technique for Applicative Languages, *Software, Practice and Experience* **9**, 1979.
- [47] D.A.Turner: Miranda: A Non-strict Functional Language with Polymorphic Types, *Proc. Functional Programming Languages and Computer Architectures*, Springer Lecture Notes in Computer Science **201**, 1985.
- [48] P.Wadler: Strictness analysis on non-flat domains (by abstract interpretation over finite domains), *Abstract Interpretation of Declarative Languages*, S.Abramsky and C.Hankin (eds.), Ellis Horwood, 1987.