

# Accurate Step Counting

Catherine Hope and Graham Hutton

School of Computer Science and IT  
University of Nottingham, UK  
{cvh,gmh}@cs.nott.ac.uk

**Abstract** Starting with an evaluator for a language, an abstract machine for the same language can be mechanically derived using successive program transformations. This has relevance to studying both the time and space properties of programs because these can be estimated by counting transitions of the abstract machine and measuring the size of the additional data structures needed, such as environments and stacks. In this paper we will use this process to derive a function that accurately counts the number of steps required to evaluate expressions in a simple language, and illustrate this function with a range of examples.

## 1 Introduction

The problem of reasoning about intensional properties of functional programs, such as the time requirements, is a long-running one. It is complicated by different evaluation strategies and sharing of expressions, meaning that some parts of a program may not be run or only partially so. One of the issues involved in reasoning about the amount of time a program will take to complete is what to count as an atomic unit in evaluation, or an evaluation step.

An evaluator is usually an implementation of the denotational semantics [13] of a language — it evaluates an expression based on the meaning of its sub-expressions. This level of understanding helps us reason about extensional properties of the language, but it doesn't say anything about the underlying way that the evaluation is taking place. By contrast, an operational semantics [10] shows us the method that is being used to evaluate an expression, and the conventional approach is to use this to measure the number of steps that are required. This, however, may not be very accurate because what is usually being measured is  $\beta$ -reductions, each of which may take arbitrarily long. An example of this approach is using the tick monad [15], which counts  $\beta$ -reductions.

It is proposed that a more realistic measure would be to count transitions in an actual machine. The idea is to get more detailed information than by just counting  $\beta$ -reductions, but still have a principled way of obtaining this information. This would not be as accurate as actual time measurements, but would provide a useful half-way house between these two approaches. What is needed is a justification that the machine correctly implements the evaluation semantics and a way to reflect the number of steps required by the machine back to the semantic level.

Recently Danvy *et al* have explored the basis of abstract machines, and the process of deriving them from evaluators [1, 2, 3, 5]. This process uses, in particular, two program transformation techniques: transformation to continuation passing style, to make order of evaluation explicit; and defunctionalization, to eliminate the consequent use of higher-order functions.

In this paper, we will apply this process to the problem of accurately counting evaluation steps. In brief, this consists of introducing a simple language in which to write expressions, an evaluator for this language, and then deriving a corresponding abstract machine using successive program transformations. Simple step counting is then added to the machine, by threading and incrementing a counter that measures the number of transitions.

The next stage is to apply the same process but in the reverse order, resulting in an evaluator that additionally counts the number of steps, directly corresponding to the number of transitions of the underlying abstract machine. Finally, a direct step counting function will be calculated from this evaluator and will be used to reason about evaluation of some example computations expressed in the language.

A particular aspect of our derivation process is that the program at each stage is *calculated* directly from a specification of its behaviour [8]. All the programs are written in Haskell [9], but no specific features of this language are used, so they may be easily adapted to other functional languages.

## 2 A Simple Language

To start with we will consider a simple language with expressions consisting of integers and addition, and with integers as values:

```
data Expr = Add Expr Expr | Val Int
type Value = Int
```

Although this language is not powerful enough to be used to analyse the time requirements of any interesting computations, it will be sufficient to show the derivation process without over-complication. An extended language will be presented later to look at some example functions.

### 2.1 Evaluator

The initial evaluator takes an expression and evaluates it to a value:

```
eval          :: Expr -> Value
eval (Val v)  = v
eval (Add x y) = eval x + eval y
```

That is, evaluating an integer value returns that integer, and evaluating an addition evaluates both sides of the addition to an integer and then adds them together. The order of evaluation is not specified at this level but will be determined by the semantics of the underlying language; in particular, when the expression is an addition, *Add x y*, the evaluation of the *x* or *y* may occur first.

## 2.2 Tail-recursive Evaluator

Our aim is to turn the evaluator into an *abstract machine*, a term-rewriting system that makes explicit the step-by-step process by which evaluation can be performed. More precisely, we seek to construct an abstract machine implemented in Haskell as a first-order, tail-recursive function.

The evaluator is already first order, but it is not tail-recursive. It can be made so by transforming it to continuation passing style (CPS) [11]. A continuation is a function that represents the rest of a computation; this makes the evaluation order of the arguments explicit, so intermediate results need to be ordered using the continuation. A program can be transformed in to CPS by redefining it to take an extra argument, a function which is applied to the result of the original one. In our case, the continuation function will take an argument of type *Value* and its result is a *Value*:

**type** *Con* = *Value* → *Value*

The new tail-recursive evaluator can be calculated from the old one by using the specification:

*evalTail* :: *Expr* → *Con* → *Value*  
*evalTail* *e* *c* = *c* (*eval* *e*)

That is, the new evaluator has the same behaviour as simply applying the continuation to the result of the original evaluator. The definition of this function can be calculated by performing induction on the structure of the expression, *e*.

Case : *e* = *Val* *v*

$$\begin{aligned} & \textit{evalTail} (\textit{Val} \ v) \ c \\ = & \quad \{ \textit{specification} \} \\ & \ c \ (\textit{eval} \ v) \\ = & \quad \{ \textit{definition of eval} \} \\ & \ c \ v \end{aligned}$$

Case : *e* = *Add* *x* *y*

$$\begin{aligned} & \textit{evalTail} (\textit{Add} \ x \ y) \ c \\ = & \quad \{ \textit{specification} \} \\ & \ c \ (\textit{eval} \ (\textit{Add} \ x \ y)) \\ = & \quad \{ \textit{definition of eval} \} \\ & \ c \ (\textit{eval} \ x \ + \ \textit{eval} \ y) \\ = & \quad \{ \textit{reverse } \beta\text{-reduction, abstract over eval } x \} \\ & \ (\lambda m \rightarrow c \ (m \ + \ \textit{eval} \ y)) \ (\textit{eval} \ x) \\ = & \quad \{ \textit{inductive assumption for } x \} \\ & \ \textit{evalTail} \ x \ (\lambda m \rightarrow c \ (m \ + \ \textit{eval} \ y)) \\ = & \quad \{ \textit{reverse } \beta\text{-reduction, abstract over eval } y \} \\ & \ \textit{evalTail} \ x \ (\lambda m \rightarrow (\lambda n \rightarrow c \ (m \ + \ n)) \ (\textit{eval} \ y)) \\ = & \quad \{ \textit{inductive assumption for } y \} \\ & \ \textit{evalTail} \ x \ (\lambda m \rightarrow \textit{evalTail} \ y \ (\lambda n \rightarrow c \ (m \ + \ n))) \end{aligned}$$





The original semantics can be recovered by passing in the equivalent of the initial continuation, the *Top* constructor:

$$eval\ e = evalMach\ e\ Top$$

## 2.4 Step Counting Machine

The number of time steps required to evaluate an expression is to be measured by counting the number of transitions of the abstract machine. The abstract machine derived can be simply modified by adding a step count that is incremented each time a transition, a function call to *evalMach* or *apply*, is made. The step count is added as an accumulator, rather than just incrementing the count that the recursive call returns, so that it is still an abstract machine.

It is relevant here to note that this is just one possible derivation to produce an abstract machine. Different abstract machines may be generated by applying different program transformations, as demonstrated in [5].

$$\begin{aligned} \text{type } Step &= Int \\ stepMach &:: (Expr, Step) \rightarrow Cont \rightarrow (Value, Step) \\ stepMach\ (Val\ v,\ s)\ c &= apply'\ c\ (v,\ s + 1) \\ stepMach\ (Add\ x\ y,\ s)\ c &= stepMach\ (x,\ s + 1)\ (AddL\ c\ y) \\ apply' &:: Cont \rightarrow (Value, Step) \rightarrow (Value, Step) \\ apply'\ Top\ (v,\ s) &= (v,\ s + 1) \\ apply'\ (AddL\ c\ y)\ (m,\ s) &= stepMach\ (y,\ s + 1)\ (AddR\ m\ c) \\ apply'\ (AddR\ m\ c)\ (n,\ s) &= apply'\ c\ (m + n,\ s + 1) \end{aligned}$$

In this case we are only counting the machine transitions and the actual addition of the integers is defined to happen instantly, though this could be extended by introducing an additional factor that represents the number of steps to perform an addition.

The evaluation function now returns a pair, where the first part is the evaluated value, and the second is the number of steps taken, which is initialised to zero. Therefore, the semantics of the original evaluator can be recovered by taking the first part of the pair returned:

$$eval\ e = fst\ (stepMach\ e\ 0\ Top)$$

## 2.5 Step Counting Tail-recursive Evaluator

The aim now is to derive a function that counts the number of steps required to evaluate an expression. The specification for this is the second part of the pair returned by the abstract machine:

$$steps\ e = snd\ (stepMach\ (e,\ 0)\ Top)$$

The first stage in the reverse process is to refunctionalize the representation of the continuation. The original continuation was a function of type  $Value \rightarrow Value$ , so the type of the new one will be  $(Value, Step) \rightarrow (Value, Step)$ .





The language is implemented as the following Haskell data type:

```
data Expr = Var String | Abs String Expr | App Expr Expr
         | Add Expr Expr | Val Value
         | Cons Expr Expr | Foldr Expr Expr Expr
data Value = Const Int | ConsV Value Value | Nil | Clo String Expr Env
```

That is, an expression is either a variable, abstraction, application of two expressions, addition of two expressions, fold-right over an expression (where the function and empty list-case arguments are both expressions), a list containing further expressions or a value. In turn, a value is either an integer, list containing further values or a closure (an abstraction paired with an environment containing bindings for all free variables in the abstraction).

The data type differentiates between expressions and values, in particular lists that contain unevaluated and evaluated expressions, so that there is no need to iterate repeatedly over the list to check if each element is fully evaluated, which would introduce an artificial evaluation overhead.

The primitive functions, such as *Add* and *Foldr*, are implemented as fully saturated, in that they take their arguments all at once. The main reason for this is to make it easier to define what a value is: if they were introduced as constants then, for example, *App Add 1* would be a value, since it cannot be further evaluated. This doesn't affect what can be expressed in the language; partial application can be expressed by using abstractions, so an equivalent expression would be *Abs "x" (Add 1 (Var "x"))*, which would be evaluated to a closure.

Note that, for simplicity, the language has not been provided with a type system. Rather, in this article we assume that only well-formed expressions are considered.

## 4 Evaluator

For simplicity, we will consider evaluation using the call-by-value strategy, where arguments to functions are evaluated before the function is applied. Evaluation is performed using an environment, that is used to look up what a variable is bound to. This avoids having to substitute in expressions for variables, which is complicated by the need to deal with avoiding name-capture. Under call-by-value evaluation arguments are evaluated before function application, so variables will be bound to a value. The environment is represented as a list of pairs:

```
type Env = [(String, Value)]
```

The initial evaluator is given below:

```

eval                               :: Expr → Env → Value
eval (Val v)                       env = v
eval (Var x)                       env = fromJust (lookup x env)
eval (Abs x e)                     env = Clo x e env
eval (App f e)                     env = let Clo x e' env' = eval f env
                                       v = eval e env
                                       in eval e' ((x, v) : env')
eval (Add x y)                     env = let Const m = eval x env
                                       Const n = eval y env
                                       in Const (m + n)
eval (Cons x xs)                   env = ConsV (eval x env) (eval xs env)
eval (Foldr f v xs)                env = case eval xs env of
  Nil → eval v env
  ConsV z zs → let f' = eval f env
                 x = eval (Foldr (Val f') v (Val zs)) env
                 in eval (App (App (Val f') (Val z)) (Val x)) []

```

That is, values are already evaluated, so they are simply returned. Variables are evaluated by returning the value the variable is bound to in the environment. Under the call-by-value strategy, evaluation is not performed under  $\lambda$ -terms, so abstractions are turned in to values by pairing them with the current environment to make a closure. An application,  $App\ f\ e$ , is evaluated by first evaluating  $f$  to an abstraction, then evaluating the body of the abstraction, with the environment extended with the variable bound to the value that  $e$  evaluates to.

Addition is performed by first evaluating both sides to an integer and then adding them together. This will give another integer result and so does not require further evaluation. Evaluating a  $Cons$  consists of evaluating the first and second arguments (the head and the tail of the list) and then re-assembling them using the  $ConsV$  constructor to make an evaluated list.

Evaluation of the  $Foldr$  case proceeds by first evaluating the list argument and doing case analysis. If the list is  $Nil$  then the result is the evaluation of the second argument,  $v$ . In the non-empty case, first the function argument is evaluated, then the fold-right applied to the tail of the list, and finally the function is applied to the head of the list and the result of folding the tail of the list. This evaluation is performed with an empty environment, since all the expressions are values at that point.

The evaluation of the fold-right could have been specified in different ways. The completely call-by-value way would be to evaluate the arguments in left to right order, so that the first two arguments are evaluated before the list argument. However, for the  $Nil$  list argument case, the function argument to  $Foldr$  is evaluated even though it is not required. The approach in the evaluator is to evaluate the list argument first to allow pattern matching and then evaluate the other arguments depending on what the list evaluated to. So when the list evaluates to  $Nil$  only the second argument to  $Foldr$  is evaluated. The justification not to use the purely call-by-value way is that it would introduce

some artificial behaviour of the *Foldr* function. When the  $\lambda$ -calculus is extended with a conditional function, for example, it is not implemented to expand both branches under call-by-value evaluation, but more efficiently by evaluating the condition first and then one branch depending on the value of the condition. In practice, this has little effect because the function supplied to the fold is often an abstraction and therefore is already evaluated.

## 5 Complete Function

Performing the derivation process for this extended language (which, as previously, proceeds by calculation) yields the following *steps* function:

$$\begin{aligned}
\text{steps } e &= \text{steps}' e [] + 1 \\
\text{steps}' (\text{Val } v) & \text{env} = 1 \\
\text{steps}' (\text{Var } x) & \text{env} = 1 \\
\text{steps}' (\text{Abs } x e) & \text{env} = 1 \\
\text{steps}' (\text{App } f e) & \text{env} = \mathbf{let} (\text{Clo } x e' \text{env}') = \text{eval } f \text{env} \\
& \quad v = \text{eval } e \text{env} \\
& \quad \mathbf{in} \text{steps}' f \text{env} + \text{steps}' e \text{env} + \\
& \quad \text{steps}' e' ((x, v) : \text{env}') + 3 \\
\text{steps}' (\text{Add } x y) & \text{env} = \text{steps}' x \text{env} + \text{steps}' y \text{env} + 3 \\
\text{steps}' (\text{Cons } x xs) & \text{env} = \text{steps}' x \text{env} + \text{steps}' xs \text{env} + 3 \\
\text{steps}' (\text{Foldr } f v xs) & \text{env} = \text{steps}' xs \text{env} + \mathbf{case} \text{eval } xs \text{env} \mathbf{of} \\
\text{Nil} & \rightarrow \text{steps}' v \text{env} + 2 \\
\text{ConsV } y ys & \rightarrow \mathbf{let} f' = \text{Val} (\text{eval } f \text{env}) \\
& \quad x = \text{Val} (\text{eval} (\text{Foldr} (\text{Val } f') v (\text{Val } ys)) \text{env}) \\
& \quad \mathbf{in} \text{steps}' f \text{env} + \text{steps}' (\text{Foldr} (\text{Val } f') v (\text{Val } ys)) \text{env} + \\
& \quad \text{steps}' (\text{App} (\text{App } f' (\text{Val } y)) x) [] + 4
\end{aligned}$$

As mentioned earlier, the derived function has calls to the original evaluator, where the result of evaluation is required. For example, in the *Foldr* case, a case analysis has to be performed on the evaluated third argument, to know if it is empty or not, and so whether to supply the *v* argument, or to keep folding.

We want to be able to reason about how the time requirements of some example functions depends on the size of the arguments to the function. In the case of functions defined using fold-right, it would be easier to reason about the time usage if it was expressed as a function over the size of the list, rather than as a recursive function — the *steps* function above the *Foldr* case makes a recursive call to fold the tail of the list. This can naturally be expressed as a fold-right over the value list data structure, defined as:

$$\begin{aligned}
\text{foldrVal} & :: (\text{Value} \rightarrow b \rightarrow b) \rightarrow b \rightarrow \text{Value} \rightarrow b \\
\text{foldrVal } f v \text{Nil} & = v \\
\text{foldrVal } f v (\text{ConsV } x xs) & = f x (\text{foldrVal } f v xs)
\end{aligned}$$

Also, if the number of steps to apply the function  $f$  to two arguments does not depend on the value of the arguments, such as adding two expressions, then a useful further simplification is to express this as a function over the length of the list argument supplied, defined as:

$$\mathit{lengthVal} = \mathit{foldrVal} (\lambda\_ n \rightarrow n + 1) 0$$

## 6 Example Functions

We can now use the derived *steps* function to look at some examples. Each step counting function produced was simplified by hand and then QuickCheck [4] was used to verify that the result produced was equal to the original function, to check that no errors had been introduced during simplification.

### 6.1 Summing a List

Summing a list of integers can be expressed using fold-right:

$$\begin{aligned} \mathit{sum} [] &= 0 \\ \mathit{sum} (x : xs) &= x + \mathit{sum} xs \end{aligned} \quad \Leftrightarrow \quad \mathit{sum} xs = \mathit{foldr} (+) 0 xs$$

The fold-right definition replaces each list constructor ( $:$ ) with  $+$ , and the  $[]$  at the end of the list with  $0$ , the unit of addition. First we translate the definition of *sum* into the language syntax and then call the *steps* function on the application of *sum* to an argument. The number of steps required to evaluate applying the *sum* function to a list of integers,  $xs$ , is given below:

$$\mathit{steps} (\mathit{App} \mathit{sum} (\mathit{Val} xs)) = 21 * (\mathit{lengthVal} xs) + 10$$

The function can be expressed as a length because the steps required to evaluate the addition of two values is a constant. The step count is a constant multiplied by the length of the list argument plus a constant amount; it is directly proportional to the length of the list argument. Though, of more interest is that we can see the constant factors involved in the evaluation.

### 6.2 Sum with an Accumulator

An alternative definition of *sum* is to use an accumulator; the fold-right is used to generate a function which is applied to the identity function in the empty list case, and in the non-empty case adds the head of the list to the accumulator.

$$\begin{aligned} \mathit{sumAcc} [] a &= a \\ \mathit{sumAcc} (x : xs) a &= \mathit{sumAcc} xs (a + x) \end{aligned} \quad \Leftrightarrow \quad \begin{aligned} \mathit{sumAcc} xs &= \mathit{foldr} f \mathit{id} xs 0 \\ &\mathbf{where} f x g a = g (a + x) \end{aligned}$$

Using an accumulator could potentially save on space, because additions could be performed without having to expand the whole list first. It would be useful to know what effect an accumulator has on the number of steps taken.

Translating the accumulator version and applying *steps* gives a result of the same form, linear on the length of the list, but the constant values are larger, because there is an additional overhead in evaluating the extra abstractions:

$$\text{steps } (\text{App } \text{sumAcc } (\text{Val } xs)) = 26 * (\text{lengthVal } xs) + 15$$

### 6.3 Concatenation

Concatenating a list of lists can be defined by folding the *append* function over the list, and *append* can also be expressed as a fold-right:

$$\begin{aligned} \text{concat } xs &= \text{foldr } \text{append } [] xs \\ \text{append } xs \ ys &= \text{foldr } (:) \ ys \ xs \end{aligned}$$

First we need to analyse the *append* function. The number of steps to evaluate *append* applied to two list arguments is as follows:

$$\text{steps } (\text{App } (\text{App } \text{append } (\text{Val } xs)) (\text{Val } ys)) = 21 * (\text{lengthVal } xs) + 15$$

So the number of steps to evaluate an *append* is proportional to the length of the first list argument. The step count of the *concat* function can now be calculated using this function. With the step count from the *append* function inlined, the resulting *steps* function is:

$$\begin{aligned} \text{steps } (\text{App } \text{concat } (\text{Val } xss)) &= \text{foldrVal } f \ 10 \ xss \\ &\mathbf{where} \ f \ ys \ s = 21 * (\text{lengthVal } ys) + 20 + s \end{aligned}$$

The number of steps required in evaluation is the sum of the steps taken to apply the *append* function to each element in the list. If the argument to *concat* is a list where all the list elements are of the same length (so the number of steps taken in applying the *append* function will always be constant), then this can be simplified to:

$$\begin{aligned} \text{steps } (\text{App } \text{concat } (\text{Val } xss)) &= 20 + \mathbf{case} \ xss \ \mathbf{of} \\ &\quad \text{Nil} \rightarrow 0 \\ &\quad \text{ConsV } ys \ yss \rightarrow (\text{lengthVal } xss) * (21 * (\text{lengthVal } ys)) \end{aligned}$$

The number of steps is now proportional to the length of the input list multiplied by the number of steps to evaluate appending an element of the list, which is proportional to the length of that element.

## 6.4 Reversing a List

Reversing a list can be expressed directly as a fold-right by appending the reversed tail of the list to the head element made in to a singleton list:

$$\begin{array}{l} \text{reverse []} = [] \\ \text{reverse (x : xs)} = \text{reverse xs} \# [x] \end{array} \Leftrightarrow \begin{array}{l} \text{reverse xs} = \text{foldr f [] xs} \\ \mathbf{where} \text{ f x xs} = \text{xs} \# [x] \end{array}$$

Translating the definition of *reverse* in to the language syntax and applying the *steps* function gives the step count function:

$$\begin{array}{l} \text{steps (App reverse (Val xs))} = \text{fst (foldrVal g (10, Nil) xs)} \\ \mathbf{where} \text{ g z (s, zs)} = (s + 21 * (\text{lengthVal zs}) + 34, \\ \text{eval (App (App append (Val zs)) (Val (ConsV z Nil))) []}) \end{array}$$

The steps function for *reverse* is dependent on the steps required to perform the *append* function for each element, which is proportional to length of the first argument to *append*. The size of this argument is increased by one each time, so the function is a sum up to the length of the list:

$$10 + \sum_{x=0}^{\text{length xs}-1} 21x + 34$$

Expanding this sum gives the following expression:

$$10 + 34 * (\text{length xs}) + \frac{21 * (\text{length xs} - 1) * (\text{length xs})}{2} \leq c (\text{length xs})^2$$

This is less than a constant multiplied by the square of the length of the list, for example when  $c = 11$  for all lists of length greater than 47. This means that the time requirements are quadratic on the length of the list [14].

## 6.5 Fast Reverse

The reverse function can also be expressed using an accumulator:

$$\begin{array}{l} \text{fastrev [] a} = a \\ \text{fastrev (x : xs) a} = \text{fastrev xs (x : a)} \end{array} \Leftrightarrow \begin{array}{l} \text{fastrev xs} = \text{foldr f id xs 0} \\ \mathbf{where} \text{ f x g a} = \text{g (x : a)} \end{array}$$

This definition should have better time properties because, as shown above, the steps required in evaluating the *append* function is proportional to the length of the first argument, so appending the tail of the list would be inefficient. The *steps* function produced for the accumulator version is directly proportional to the length of the list:

$$\text{steps (App fastrev (Val xs))} = 26 * (\text{lengthVal xs}) + 15$$

## 7 Conclusion and Further Work

We have presented a process that takes an evaluator for a language, and derives a function that gives an accurate count of the number of steps required to evaluate expressions using an abstract machine for the language. Moreover, all the steps in the derivation process are purely calculational, in that the required function at each stage is calculated directly from a specification of its desired behaviour.

Using an extended  $\lambda$ -calculus under call-by-value evaluation, the examples in the previous section give the expected complexity results, but also show the constants involved. This is useful to know because of the additional overheads that functions of the same complexity may have, for example in summing a list with and without an accumulator. They also show the boundaries at which one function with a lower growth rate but larger constants becomes quicker than another, for example in the two different definitions of the reverse function.

Ultimately, the most accurate information is to do real timing of programs. The approach taken in the profiler [12] that is distributed with the Glasgow Haskell Compiler, is similar to that here, in that abstract machines are used to bridge the gap between a big-step cost semantics and a small-step implementation. A regular clock interrupt is used to collect time information and assign to the costs of functions. However, this is a profiler, and we want to be able to express costs as a function over the arguments supplied, and not just in terms of hard time measurements. The thesis of this paper is that we can obtain this useful information relatively easily.

This work could be extended by considering more complicated evaluation strategies, such as call-by-name or lazy evaluation. It would also be interesting to apply the same technique to look at the space requirements for functions, by measuring the size of the additional data structures produced at the abstract machine level. Finally, a useful addition to this work would be to develop a calculus to automate deriving the step functions.

### Acknowledgement

Finally, the authors would like to thank the anonymous referees for their comments and suggestions, which have considerably improved the paper.

### References

- [1] Mads Sig Ager. From natural semantics to abstract machines. In *Logic Based Program Synthesis and Transformation, 14th International Symposium (LOPSTR 2004), Revised Selected Papers*.
- [2] Mads Sig Ager, Dariusz Biernacki, Olivier Danvy, and Jan Midtgaard. A functional correspondence between evaluators and abstract machines. Technical Report RS-03-13, March 2003. 28 pp. Appears in , pages 8–19.
- [3] Mads Sig Ager, Olivier Danvy, and Jan Midtgaard. A functional correspondence between call-by-need evaluators and lazy abstract machines. *Information Processing Letters*, 90(5):223–232, 2004. Extended version available as the technical report BRICS-RS-04-3.

- [4] Koen Claessen and John Hughes. Quickcheck: a lightweight tool for random testing of haskell programs. In *ICFP*, pages 268–279, 2000.
- [5] Olivier Danvy. A rational deconstruction of Landin’s SECD machine. Technical Report RS-03-33, October 2003.
- [6] Olivier Danvy. On evaluation contexts, continuations, and the rest of the computation. Number CSR-04-1, pages 13–23, Birmingham B15 2TT, United Kingdom, 2004. Invited talk.
- [7] Graham Hutton. A Tutorial on the Universality and Expressiveness of Fold. *Journal of Functional Programming*, 9(4):355–372, July 1999.
- [8] Graham Hutton and Joel Wright. Calculating an Exceptional Machine. To appear in the Proceedings of the Fifth Symposium on Trends in Functional Programming, 2005.
- [9] S. Peyton Jones. Haskell 98 language and libraries: The revised report. Technical report.
- [10] G. D. Plotkin. A Structural Approach to Operational Semantics. Technical Report DAIMI FN-19, University of Aarhus, 1981.
- [11] John C. Reynolds. Definitional interpreters for higher-order programming languages. *Higher Order Symbol. Comput.*, 11(4):363–397, 1998.
- [12] Patrick M. Sansom and Simon L. Peyton Jones. Formally based profiling for higher-order functional languages. *ACM Trans. Program. Lang. Syst.*, 19(2):334–385, 1997.
- [13] David A. Schmidt. *Denotational semantics: a methodology for language development*. William C. Brown Publishers, Dubuque, IA, USA, 1986.
- [14] Clifford A. Shaffer. *A Practical Introduction to Data Structures and Algorithm Analysis*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 2000.
- [15] Philip Wadler. Monads for functional programming. In M. Broy, editor, *Program Design Calculi: Proceedings of the 1992 Marktoberdorf International Summer School*. Springer-Verlag, 1993.