



Basic Research in Computer Science

On the Dynamic Extent of Delimited Continuations

Dariusz Biernacki
Olivier Danvy

**Copyright © 2005, Dariusz Biernacki & Olivier Danvy.
BRICS, Department of Computer Science
University of Aarhus. All rights reserved.**

**Reproduction of all or part of this work
is permitted for educational or research use
on condition that this copyright notice is
included in any copy.**

**See back inner page for a list of recent BRICS Report Series publications.
Copies may be obtained by contacting:**

**BRICS
Department of Computer Science
University of Aarhus
Ny Munkegade, building 540
DK-8000 Aarhus C
Denmark
Telephone: +45 8942 3360
Telefax: +45 8942 3255
Internet: BRICS@brics.dk**

**BRICS publications are in general accessible through the World Wide
Web and anonymous FTP through these URLs:**

`http://www.brics.dk`
`ftp://ftp.brics.dk`
This document in subdirectory RS/05/2/

On the Dynamic Extent of Delimited Continuations

Dariusz Biernacki and Olivier Danvy

BRICS*

Department of Computer Science

University of Aarhus[†]

January 2005

Abstract

We show that breadth-first traversal exploits the difference between the static delimited-control operator `shift` (alias \mathcal{S}) and the dynamic delimited-control operator `control` (alias \mathcal{F}). For the last 15 years, this difference has been repeatedly mentioned in the literature but it has only been illustrated with one-line toy examples. Breadth-first traversal fills this vacuum.

We also point out where static delimited continuations naturally give rise to the notion of control stack whereas dynamic delimited continuations can be made to account for a notion of ‘control queue.’

Keywords

Delimited continuations, direct style, continuation-passing style (CPS), CPS transformation, defunctionalization, control operators, shift and reset, control and prompt.

*Basic Research in Computer Science (www.brics.dk),
funded by the Danish National Research Foundation.

[†]IT-parken, Aabogade 34, DK-8200 Aarhus N, Denmark.
Email: {dabi,danvy}@brics.dk

Contents

1	Introduction	1
1.1	Background	1
1.2	Overview	1
2	An operational characterization	2
2.1	An abstract machine for <code>shift</code> and <code>reset</code>	2
2.2	An abstract machine for <code>control</code> and <code>prompt</code>	4
2.3	Simulating <code>shift</code> in terms of <code>control</code> and <code>prompt</code>	5
2.4	Simulating <code>control</code> in terms of <code>shift</code> and <code>reset</code>	5
2.5	Three examples in ML	6
3	Programming with delimited continuations	7
4	The samefringe problem	8
4.1	Depth first	9
4.1.1	An eager traversal	9
4.1.2	A lazy traversal	10
4.1.3	A continuation-based traversal	10
4.1.4	A direct-style traversal with <code>shift</code> and <code>reset</code>	10
4.1.5	A stack-based traversal	11
4.2	Breadth first	13
4.2.1	A queue-based traversal	13
4.2.2	A direct-style traversal with <code>control</code> and <code>prompt</code>	14
4.3	Summary and conclusion	15
5	Labeling a tree	15
5.1	Breadth-first numbering	15
5.1.1	A queue-based traversal	16
5.1.2	A direct-style traversal with <code>control</code> and <code>prompt</code>	16
5.2	Depth-first numbering	18
5.2.1	A stack-based traversal	18
5.2.2	A continuation-based traversal	19
5.2.3	A direct-style traversal with <code>shift</code> and <code>reset</code>	20
5.3	Summary and conclusion	21
6	Conclusion and issues	22
A	An implementation of <code>shift</code> and <code>reset</code>	23
B	An implementation of <code>control</code> and <code>prompt</code>	24

1 Introduction

To distinguish between the static extent and the dynamic extent of delimited continuations, let us first review the notions of continuation and of delimited continuation.

1.1 Background

Continuation-passing style (CPS) is a time-honored and logic-based format for functional programs where all intermediate results are named, all calls are tail calls, and programs are evaluation-order independent [29, 42, 49, 53, 59]. While this format has been an active topic of study [5, 19, 26, 28, 31, 37, 39, 45, 47, 50, 54, 57, 62], it also has been felt as a straightjacket both from a semantics point of view [19, 20, 22, 23, 34, 35, 58] and from a programming point of view [11, 12, 14, 15], where one would like to relax the tail-call constraint and compose continuations.

In direct style, continuations are accessed with a variety of control operators such as Landin’s `J` [38], Reynolds’s `escape` [53], Scheme’s `call/cc` [10, 36], and Standard ML of New Jersey’s `callcc` and `throw` [18]. These control operators give access to the current continuation as a first-class value. Activating such a first-class continuation has the effect of resuming the computation at the point where this continuation was captured; the then-current continuation is *abandoned*. Such first-class continuations *do not return to the point of their activation*—they model jumps, i.e., tail calls [59, 60].

In direct style, composable continuations are also accessed with control operators such as Felleisen et al.’s `control` (alias \mathcal{F}) [22, 23, 58] and Danvy and Filinski’s `shift` (alias \mathcal{S}) [15, 16]. These control operators also give access to the current continuation as a first-class value; activating such a first-class continuation also has the effect of resuming the computation at the point where this continuation was captured; the then-current continuation, however, *is then resumed*. Such first-class continuations *return to the point of their activation*—they model non-tail calls.

For a first-class continuation to return to the point of its activation, one must declare its point of completion, since this point is no longer at the very end of the overall computation, as with traditional, undelimited first-class continuations. In direct style, this declaration is achieved with a new kind of operator, due to Felleisen [19, 20]: a control delimiter. The control delimiter corresponding to `control` is called `prompt` (alias `#`). The control delimiter corresponding to `shift` is called `reset` (alias `<·>`) and its continuation-passing counterpart is a classical backtracking idiom in functional programming [1, 7, 9, 41, 55, 61]. Other, more advanced, delimited-control operators exist [30, 33, 44, 46, 52]; we return to them in the conclusion.

In the present work, we focus on `shift` and `control`.

1.2 Overview

In Section 2, we present an environment-based abstract machine that specifies the behaviors of `shift` and `control`, and we show how the extent of a `shift`-abstracted delimited continuation is static whereas that of a `control`-abstracted delimited continuation is dynamic. We show how `shift` can be trivially simulated in terms of

`control` and `prompt`, which is a well-known result [15, Section 5], and we review Shan’s simulation of `control` in terms of `shift` and `reset`, which is a recent result [56]. In Section 3, we present a roadmap of Sections 4 and 5, where we show how the static extent of a delimited continuation is compatible with a control stack and depth-first traversal, and how the dynamic extent of a delimited continuation can be made to account for a ‘control queue’ and breadth-first traversal.

Prerequisites and preliminaries: Besides some awareness of CPS and the CPS transformation [16, 49, 59], we assume a passing familiarity with defunctionalization [17, 53].

Our programming language of discourse is Standard ML [43]. In the following sections, we will make use of the notational equivalence of expressions such as

```
x1 :: x2 :: xs
(x1 :: x2 :: nil) @ xs
[x1, x2] @ xs
```

where `::` denotes infix list construction and `@` denotes infix list concatenation. In an environment where `x1` denotes 1, `x2` denotes 2, and `xs` denotes [3, 4, 5], each of the three expressions above evaluates to [1, 2, 3, 4, 5].

2 An operational characterization

In our previous work [6], we derived an environment-based abstract machine for the λ -calculus with `shift` and `reset` by defunctionalizing the corresponding definitional interpreter [15]. We use this abstract machine to explain the static extent of the delimited continuations abstracted by `shift` and the dynamic extent of the delimited continuations abstracted by `control`.

2.1 An abstract machine for `shift` and `reset`

The abstract machine is displayed in Figure 1; `reset` is noted $\langle \cdot \rangle$ and `shift` is noted \mathcal{S} . The set of possible values consists of closures and captured contexts. The machine extends Felleisen et al.’s CEK machine [21] with a meta-context C_2 , the two transitions for $\langle \cdot \rangle$ and \mathcal{S} , and the transition for applying a captured context to a value in an evaluation context and a meta-context. Intuitively, an evaluation context represents the rest of the computation up to the nearest enclosing delimiter, and a meta-context represents all of the remaining computation [13].

Given a term t , the machine is initialized in an *eval*-state with an empty environment e_{empty} , an empty context end , and an empty meta-context \bullet . The transitions out of an *eval*-state are defined by cases on its first component:

- a variable x is looked up in the current environment and the machine switches to a $cont_1$ -state;
- an abstraction $\lambda x.t$ is evaluated into a closure $[x, t, e]$ and the machine switches to a $cont_1$ -state;

- Terms: $t ::= x \mid \lambda x.t \mid t_0 t_1 \mid \langle t \rangle \mid \mathcal{S}k.t$
- Values (closures and captured continuations): $v ::= [x, t, e] \mid C_1$
- Environments: $e ::= e_{empty} \mid e[x \mapsto v]$
- Evaluation contexts: $C_1 ::= \text{end} \mid \text{arg}((t, e), C_1) \mid \text{fun}(v, C_1)$
- Meta-contexts: $C_2 ::= \bullet \mid C_1 \cdot C_2$
- Initial transition, transition rules, and final transition:

t	\Rightarrow	$\langle t, e_{empty}, \text{end}, \bullet \rangle_{eval}$
$\langle x, e, C_1, C_2 \rangle_{eval}$	\Rightarrow	$\langle C_1, e(x), C_2 \rangle_{cont_1}$
$\langle \lambda x.t, e, C_1, C_2 \rangle_{eval}$	\Rightarrow	$\langle C_1, [x, t, e], C_2 \rangle_{cont_1}$
$\langle t_0 t_1, e, C_1, C_2 \rangle_{eval}$	\Rightarrow	$\langle t_0, e, \text{arg}((t_1, e), C_1), C_2 \rangle_{eval}$
$\langle \langle t \rangle, e, C_1, C_2 \rangle_{eval}$	\Rightarrow	$\langle t, e, \text{end}, C_1 \cdot C_2 \rangle_{eval}$
$\langle \mathcal{S}k.t, e, C_1, C_2 \rangle_{eval}$	\Rightarrow	$\langle t, e[k \mapsto C_1], \text{end}, C_2 \rangle_{eval}$
$\langle \text{end}, v, C_2 \rangle_{cont_1}$	\Rightarrow	$\langle C_2, v \rangle_{cont_2}$
$\langle \text{arg}((t, e), C_1), v, C_2 \rangle_{cont_1}$	\Rightarrow	$\langle t, e, \text{fun}(v, C_1), C_2 \rangle_{eval}$
$\langle \text{fun}([x, t, e], C_1), v, C_2 \rangle_{cont_1}$	\Rightarrow	$\langle t, e[x \mapsto v], C_1, C_2 \rangle_{eval}$
$\langle \text{fun}(C'_1, C_1), v, C_2 \rangle_{cont_1}$	\Rightarrow	$\langle C'_1, v, C_1 \cdot C_2 \rangle_{cont_1}$
$\langle C_1 \cdot C_2, v \rangle_{cont_2}$	\Rightarrow	$\langle C_1, v, C_2 \rangle_{cont_1}$
$\langle \bullet, v \rangle_{cont_2}$	\Rightarrow	v

Figure 1: A call-by-value environment-based abstract machine for the λ -calculus extended with **shift** (\mathcal{S}) and **reset** ($\langle \cdot \rangle$)

- an application $t_0 t_1$ is processed by pushing t_1 and the environment onto the context and switching to an *eval*-state to process t_0 ;
- a reset-expression $\langle t \rangle$ is processed by pushing the current context on the current meta-context and switching to an *eval*-state to process t in an empty context, as an intermediate computation;
- a shift-expression $\mathcal{S}k.t$ is processed by capturing the context C_1 and binding it to k , and switching to an *eval*-state to process t in an empty context.

The transitions of a $cont_1$ -state are defined by cases on its first component:

- an empty context `end` specifies that an intermediate computation is completed; it is processed by switching to a $cont_2$ -state;
- a context `arg` $((t, e), C_1)$ specifies is processed by switching to an $eval$ -state to process t in a new context;
- a context `fun` $([x, t, e], C_1)$ specifies the application of a closure; it is processed by switching to an $eval$ -state to process the term t with an extension of the environment e ;
- a context `fun` (C'_1, C_1) specifies the application of a captured context; it is processed by pushing C_1 on top of the meta-context and switching to a $cont_1$ -state to process C'_1 .

The transitions of a $cont_2$ -state are defined by cases on its first component:

- an empty meta-context `•` specifies that the overall computation is completed; it is processed as a final transition;
- a non-empty meta-context specifies that the overall computation is not complete; $C_1 \cdot C_2$ is processed by switching to an $cont_1$ -state to process C_1 .

All in all, this abstract machine is a straight defunctionalized continuation-passing evaluator [6, 15].

2.2 An abstract machine for control and prompt

Unlike `shift` and `reset`, whose definition is based on CPS, `control` and `prompt` are specified by representing delimited continuations as a list of stack frames and their composition as the concatenation of these representations [23]. Such a concatenation function \star is defined as follows:

$$\begin{aligned} \text{end} \star C'_1 &= C'_1 \\ (\text{arg}((t, e), C_1)) \star C'_1 &= \text{arg}((t, e), C_1 \star C'_1) \\ (\text{fun}(v, C_1)) \star C'_1 &= \text{fun}(v, C_1 \star C'_1) \end{aligned}$$

It is then simple to modify the abstract machine to compose delimited continuations by concatenating their representation: in Figure 1, one merely replaces the transition applying a captured context C'_1 by pushing the current context C_1 onto the meta-context C_2 , i.e.,

$$\boxed{\langle \text{fun}(C'_1, C_1), v, C_2 \rangle_{cont_1} \Rightarrow \langle C'_1, v, C_1 \cdot C_2 \rangle_{cont_1}}$$

with a transition that applies a captured context C'_1 by concatenating it with the current context C_1 :

$$\boxed{\langle \text{fun}(C'_1, C_1), v, C_2 \rangle_{cont_1} \Rightarrow \langle C'_1 \star C_1, v, C_2 \rangle_{cont_1}}$$

This change gives \mathcal{S} (alias `shift`) the behavior of \mathcal{F} (alias `control`). In contrast, `reset` and `prompt` have the same definition. The rest of the machine does not change.

In our previous work [6, Section 4.5], we have pointed out that the dynamic behavior of `control` is incompatible with CPS because the modified abstract machine no longer corresponds to a defunctionalized continuation-passing evaluator [17]. Indeed `shift` is static, whereas `control` is dynamic in the following sense:

- `shift` captures a delimited continuation in a representation C_1 that remains distinct from the current context C'_1 , when it is applied. Consequently, the current context C'_1 *cannot* be accessed from C_1 by another use of `shift`.
- `control` captures a delimited continuation in a representation C_1 that grafts itself to the current context C'_1 , when it is applied. Consequently, the current context C'_1 *can* be accessed from C_1 by another use of `control`.

This difference of behavior can be observed with delimited continuations that, when applied, capture the current continuation [14, Section 6.1] [16, Section 5.3]. A `control`-given delimited continuation dynamically captures the current continuation, above and beyond its point of resumption, whereas a `shift`-given delimited continuation statically captures the current continuation up to its point of resumption.

2.3 Simulating shift in terms of control and prompt

It is simple to obtain the effect of `shift` using `control`: one should just replace every occurrence of a captured continuation k with $\lambda v. \#(k v)$, i.e., in ML, of `k` with `fn v => prompt (fn () => k v)`. This way, when k (i.e., some context C'_1) is applied, the context of its application is always `end` and it is a simple corollary that $C'_1 \star \text{end} = C'_1$.

2.4 Simulating control in terms of shift and reset

In his recent work [56], Shan presents a (curried) continuation+state-passing transformation of `control` and `prompt` where the state represents a prefix of a meta-continuation and is equipped with operators to compose the representation of delimited continuations. This continuation+state-passing transformation conservatively extends the usual call-by-value CPS transformation for the λ -calculus, though it hinges on the requirement that continuations are recursive (or more precisely, that their answer type is higher order and recursive).

As a byproduct of this continuation+state-passing transformation, Shan then expresses `control` and `prompt` in terms of `shift` and `reset`, which is unexpected since it shows that `shift` is actually as expressive as `control`. Shan's implementation is in Scheme. We present it in Standard ML in Appendix B. Our implementation uses Filinski's implementation of `shift` and `reset` in Standard ML of New Jersey [25], which we present in Appendix A. Filinski's implementation takes the form of a functor mapping the type of intermediate answers to a structure containing an instance of `shift` and `reset` at that type. Likewise, our implementation takes the form of a functor mapping the type of intermediate answers to a structure containing an instance of `control` and `prompt` at that type.

Whereas Filinski’s implementation of `shift` and `reset` has a correctness proof [25], we are not aware of one for Shan’s implementation of `control` and `prompt`, and it is not in the scope of the present article to provide it. We have, however, tested it both in Scheme and in ML, and compared it against our implementation of the abstract machine of Section 2.2 in a variety of situations.

2.5 Three examples in ML

Using the implementation of `shift` and `reset` (Appendix A), and of `control` and `prompt` (Appendix B), we can verify that the ML expression

```
reset
  (fn () => shift (fn k => 10 + (k 100))
    + shift (fn k' => 1))
```

evaluates to 11, whereas (replacing `reset` by `prompt` and `shift` by `control`)

```
prompt
  (fn () => control (fn k => 10 + (k 100))
    + control (fn k' => 1))
```

evaluates to 1 and (delimiting the application of `k` with `prompt`)

```
prompt
  (fn () => control (fn k => 10 + ((fn v => prompt (fn () => k v)) 100))
    + control (fn k' => 1))
```

evaluates to 11.

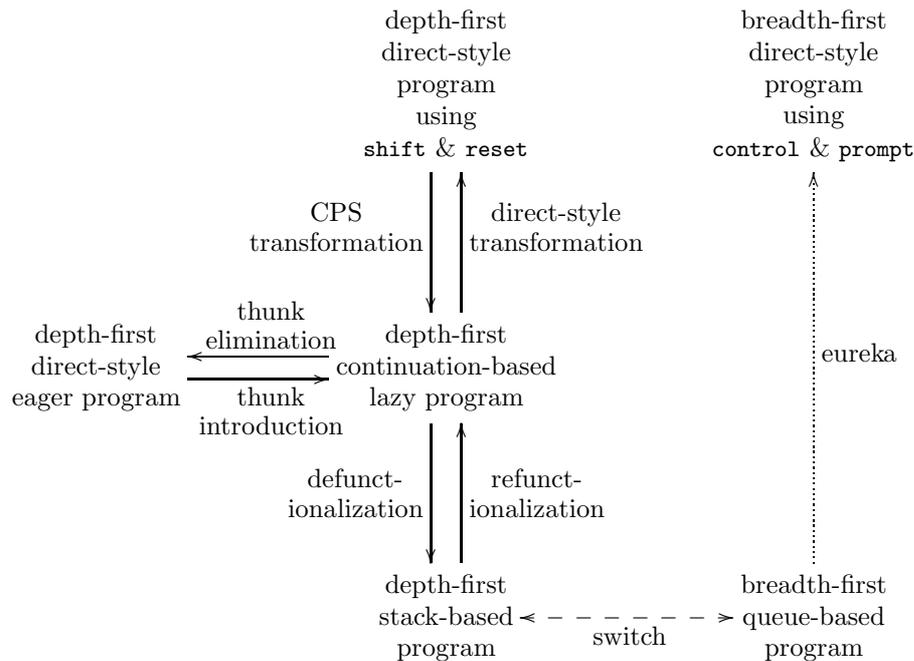
In the first case, when `k` is applied, the expression `shift (fn k' => 1)` is evaluated in a context that could be represented functionally as `fn v => 100 + v` and in a meta-context that could be represented as `(fn v => 10 + v) :: nil`; this context is captured and discarded and the intermediate answer is 1; this intermediate answer is plugged into the top context from the meta-context, i.e., `fn v => 10 + v` is applied to 1; the next intermediate answer is 11; and it is the final answer since the meta-context is empty.

In the second case, when `k` is applied, the expression `control (fn k' => 1)` is evaluated in a context that results from composing `fn v => 10 + v` and `fn v => 100 + v` (and therefore could be represented functionally as `fn v => 10 + (100 + v)`), and in a meta-context which is empty; this context is captured and discarded and the intermediate answer is 1; and it is the final answer since the meta-context is empty.

In the third case, when `k` is applied, the expression `control (fn k' => 1)` is evaluated in a context that results from composing `fn v => v` and `fn v => 100 + v` (and therefore could be represented functionally as `fn v => 100 + v`), and in a meta-context which could be represented as `(fn v => 10 + v) :: nil`; this context is captured and discarded and the intermediate answer is 1; this intermediate answer is plugged into the top context from the meta-context, i.e., `fn v => 10 + v` is applied to 1; the next intermediate answer is 11; and it is the final answer since the meta-context is empty.

3 Programming with delimited continuations

In Section 4 and 5, we present an array of solutions to the traditional samefringe example and to its breadth-first counterpart, and to Okasaki's breadth-first numbering pearl and to its depth-first counterpart. The presentation is structured according to the following diagram:



- Our starting point here is a direct-style eager program (left side of the diagram). We can make this program lazy by using thunks, i.e., functions of type `unit -> 'a` (center of the diagram).
- We can then defunctionalize the thunks in the lazy program, obtaining a stack-based program (bottom center of the diagram).
- Alternatively, we can view the type `unit -> 'a` not as a functional device to implement laziness but as a delimited continuation. The lazy program is then, in actuality, a continuation-based one, and one that is the CPS counterpart of a direct-style program using `shift` and `reset` (top center of the diagram).
- The stack-based program (bottom center of the diagram) implements a depth-first traversal. Replacing the stack with a queue yields a program implementing a breadth-first traversal (bottom right of the diagram).
- By analogy with the rest of the diagram, we infer the direct-style program using `control` and `prompt` (top right of the diagram) from this queue-based program.

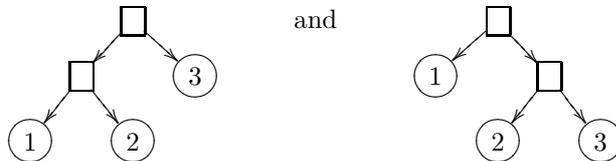
The three nodes in the center of the diagram—the CPS program, its direct-style counterpart, and its defunctionalized counterpart—follow the transformational tradition established in Reynolds’s and Wand’s seminal articles about continuations [53, 63]. In particular the ‘data-structure continuation’ [63, page 178] of the depth-first program is a stack. By analogy, the data-structure continuation of the breadth-first program is a queue. We conjecture that the queue-based program could be mechanically obtained from the direct-style one by some kind of ‘abstract CPS transformation’ [23, 51], but fleshing out this conjecture falls out of the scope of this article.

4 The samefringe problem

We present a spectrum of solutions to the samefringe problem, both in its traditional depth-first form and in its breadth-first counterpart. We work on Lisp-like binary trees of integers (S-expressions):

```
datatype tree = LEAF of int
              | NODE of tree * tree
```

The samefringe problem is traditionally stated as follows. Given two trees of integers, one wants to know whether they have the same sequence of leaves when read from left to right. For example, the two trees `NODE (NODE (LEAF 1, LEAF 2), LEAF 3)` and `NODE (LEAF 1, NODE (LEAF 2, LEAF 3))` have the same fringe [1, 2, 3] (representing it as a list) even though they are shaped differently:



Computing a fringe is done by traversing a tree depth-first and from left to right.

By analogy, we also address the breadth-first counterpart of the samefringe problem. Given two trees of integers, we want to know whether they have the same fringe when traversed in left-to-right breadth-first order. For example, the breadth-first fringe of the left tree just above is [3, 1, 2] and that of the right tree just above is [1, 2, 3].

We express the samefringe function generically by abstracting the representation of sequences of leaves with a data type `sequence` and a notion of computation (to compute the next element in a sequence):

```
signature GENERATOR
= sig
  type 'a computation
  datatype sequence = END
                    | NEXT of int * sequence computation
```

```

    val make_sequence : tree -> sequence
    val compute : sequence computation -> sequence
end

```

The following functor maps a representation of sequences of leaves to a structure containing the samefringe function. Given two trees, `same_fringe` maps them into two sequences of integers (with `make_sequence`) and iteratively traverses these sequences with an auxiliary `loop` function. This function stops as soon as one of the two sequences is exhausted or two differing leaves are found:

```

functor make_Same_Fringe (structure G : GENERATOR)
= struct
  (* same_fringe : tree * tree -> bool *)
  fun same_fringe (t1, t2)
    = let (* loop : G.sequence * G.sequence -> bool *)
        fun loop (G.END, G.END)
          = true
          | loop (G.NEXT (i1, a1), G.NEXT (i2, a2))
          = i1 = i2 andalso loop (G.compute a1, G.compute a2)
          | loop _
          = false
        in loop (G.make_sequence t1, G.make_sequence t2)
      end
end

```

In the remainder of this section, we review a variety of generators.

4.1 Depth first

4.1.1 An eager traversal

The simplest solution is to represent sequences as a data type isomorphic to that of lists. To this end, we define `make_sequence` as an accumulator-based flatten function:

```

structure Eager_generator : GENERATOR
= struct
  datatype sequence = END
                    | NEXT of int * sequence computation
  withtype 'a computation = 'a

  (* visit : tree * sequence computation -> sequence *)
  fun visit (LEAF i, a)
    = NEXT (i, a)
    | visit (NODE (t1, t2), a)
    = visit (t1, visit (t2, a))

  fun make_sequence t
    = visit (t, END)

  fun compute value
    = value
end

```

In this solution, the sequence of leaves is built eagerly and therefore completely before any comparison takes place. This choice is known to be inefficient because if two leaves differ, the remaining two sequences are not used and therefore did not need to be built.

4.1.2 A lazy traversal

A more efficient solution—and indeed a traditional motivation for lazy evaluation [27, 32]—is to construct the sequences lazily and to traverse them on demand. In the following generator, the data type `sequence` implements finite lazy sequences; the construction of the rest of the lazy sequence is delayed with a thunk of type `unit -> sequence`; and `make_sequence` is defined as an accumulator-based flatten function:

```
structure Lazy_generator : GENERATOR
= struct
  datatype sequence = END
                    | NEXT of int * sequence computation
  withtype 'a computation = unit -> 'a

  (* visit : tree * sequence computation -> sequence *)
  fun visit (LEAF i, a)
    = NEXT (i, a)
    | visit (NODE (t1, t2), a)
      = visit (t1, fn () => visit (t2, a))

  fun make_sequence t
    = visit (t, fn () => END)

  fun compute thunk
    = thunk ()
end
```

Unlike in the eager solution, the construction of the sequence in `Lazy_generator` and the comparisons in `same_fringe` are interleaved. This choice is known to be more efficient because if two leaves differ, the remaining two sequences are not built at all.

4.1.3 A continuation-based traversal

Alternatively to viewing the thunk of type `unit -> sequence`, in the lazy traversal of Section 4.1.2, as a functional device to implement laziness, we can view it as a delimited continuation that is initialized in the initial call to `visit` in `make_sequence`, extended in the inductive case of `visit`, captured in the base case of `visit`, and resumed in `compute`. From that viewpoint, the lazy traversal is also a continuation-based one.

4.1.4 A direct-style traversal with shift and reset

In direct style, the initialization of the delimited continuation `a` of Section 4.1.3 is obtained with the control delimiter `reset`, its extension is obtained by functional

sequencing, its capture is obtained with the delimited-control operator `shift`, and its resumption is obtained by function application.

Using Filinski's functor `Shift_and_Reset` defined in Appendix A, one can therefore define the lazy generator in direct style as follows:

```

structure Lazy_generator_with_shift_and_reset : GENERATOR
= struct
  datatype sequence = END
                    | NEXT of int * sequence computation
  withtype 'a computation = unit -> 'a

  local structure SR = Shift_and_Reset
                    (type intermediate_answer = sequence)
  in val shift = SR.shift
      val reset = SR.reset
  end

  (* visit : tree -> unit *)
  fun visit (LEAF i)
    = shift (fn a => NEXT (i, a))
  | visit (NODE (t1, t2))
    = let val () = visit t1
      in visit t2
      end

  fun make_sequence t
    = reset (fn () => let val () = visit t
                    in END
                    end)

  fun compute thunk
    = thunk ()
end

```

CPS-transforming `visit` and `make_sequence` yields the definitions of Section 4.1.2. The key point of this CPS transformation is that given a continuation `k`, the expression

```

let val () = visit t1
in visit t2
end

```

is transformed into:

```

visit (t1, fn () => visit (t2, k))

```

4.1.5 A stack-based traversal

Alternatively to writing the lazy solution in direct style, we can defunctionalize its computation (which has type `sequence computation`, i.e., `unit -> sequence`) and

obtain a first-order solution [17,53]. The inhabitants of the function space `unit -> sequence` are instances of the function abstractions in the initial call to `visit` (i.e., `fn () => END`) and in the inductive case of `visit` (i.e., `fn () => visit (t2, a)`). We therefore represent this function space by (1) a sum corresponding to these two possibilities, and (2) the corresponding apply function, `continue`, to interpret each of the summands. We represent this sum with an ML data type, which is recursive because of the recursive call to `visit`. This data type is isomorphic to that of a list of subtrees, which we use for simplicity in the code below. The result is essentially McCarthy's solution [40]:

```

structure Lazy_generator_stack_based : GENERATOR
= struct
  datatype sequence = END
                    | NEXT of int * sequence computation
  withtype 'a computation = tree list

  (* visit : tree * tree list -> sequence *)
  fun visit (LEAF i, a)
    = NEXT (i, a)
    | visit (NODE (t1, t2), a)
    = visit (t1, t2 :: a)
  (* continue : tree list * unit -> sequence *)
  and continue (nil, ())
    = END
    | continue (t :: a, ())
    = visit (t, a)

  fun make_sequence t
    = visit (t, nil)

  fun compute a
    = continue (a, ())
end

```

This solution traverses a given tree incrementally by keeping a stack of its subtrees. To make this point more explicit, and as a stepping stone towards breadth-first traversal, let us fold the definition of `continue` in the induction case of `visit` so that `visit` always calls `continue`:

```

| visit (NODE (t1, t2), a)
  = continue (t1 :: t2 :: a, ())

```

(Unfolding the call to `continue` gives back the definition above.)

We now clearly have a stack-based definition of depth-first traversal, and furthermore we have shown that this stack corresponds to the continuation of a function implementing a recursive descent. (Such a stack is referred to as a 'data-structure continuation' in the literature [63, page 179].)

4.2 Breadth first

4.2.1 A queue-based traversal

Replacing the (last-in, first-out) stack, in the definition of Section 4.1.5, by a (first-in, first-out) queue yields a definition that implements breadth-first, rather than depth-first, traversal:

```
structure Lazy_generator_queue_based : GENERATOR
= struct
  datatype sequence = END
                    | NEXT of int * sequence computation
  withtype 'a computation = tree list

  (* visit : tree * tree list -> sequence *)
  fun visit (LEAF i, a)
    = NEXT (i, a)
    | visit (NODE (t1, t2), a)
    = continue (a @ [t1, t2], ())

  (* continue : tree list * unit -> sequence *)
  and continue (nil, ())
    = END
    | continue (t :: a, ())
    = visit (t, a)

  fun make_sequence t
    = visit (t, nil)

  fun compute a
    = continue (a, ())
end
```

In contrast to Section 4.1.5, where the clause for nodes was (essentially) concatenating the two subtrees in front of the list of subtrees:

```
| visit (NODE (t1, t2), a)
  = continue ([t1, t2] @ a, ())  (* then *)
```

the clause for nodes is concatenating the two subtrees in the back of the list of subtrees:

```
| visit (NODE (t1, t2), a)
  = continue (a @ [t1, t2], ())  (* now *)
```

Nothing else changes in the definition of the generator. In particular, subtrees are still removed from the front of the list of subtrees by `continue`. With this last-in, first-out policy, the generator yields a sequence in breadth-first order.

Because the `::`-constructors of the list of subtrees are not solely consumed by `continue` but also by `@`, this definition *is not in the range of defunctionalization* [17]. Therefore, even though `visit` is tail-recursive and constructs a data structure that is interpreted in `continue`, it does not correspond to a continuation-passing function. And indeed, it is well-known that traversing an inductive data structure breadth-first does not mesh with the visitor pattern of functional programming, i.e., compositional recursive descent (catamorphism).

4.2.2 A direct-style traversal with control and prompt

The critical operation in the definition of `visit`, in Section 4.1.5, is the enqueueing of the subtrees `t1` and `t2` to the current queue `a`, which is achieved by the list concatenation `a @ [t1, t2]`. We observe that this concatenation matches the concatenation of stack frames in the specification of `control` in Section 2.2.

Therefore—and this is a eureka step—one can write `visit` in direct style using `control` and `prompt`. To this end, we represent the current queue `a` as a dynamic delimited continuation and we apply it before visiting `t1` and `t2`.

In order to be able to compose (concatenate) delimited continuations in the code below, we have to inject the type `unit` into the type `sequence`, so that all the delimited continuations of type `sequence computation` are endofunctions, i.e., functions with the same domain and co-domain. To this end, we represent the only value of type `unit`, i.e., `()`, by the value `END`. The direct-style traversal reads as follows:

```
structure Lazy_generator_with_control_and_prompt : GENERATOR
= struct
  datatype sequence = END
                    | NEXT of int * sequence computation
  withtype 'a computation = sequence -> 'a

  local structure CP = Control_and_Prompt
                    (type intermediate_answer = sequence)
  in val control = CP.control
      val prompt = CP.prompt
  end

  (* visit : tree -> sequence *)
  fun visit (LEAF i)
    = control (fn a => NEXT (i, a))
  | visit (NODE (t1, t2))
    = control (fn a => let val END = a END
                      val END = visit t1
                      in visit t2
                      end)

  fun make_sequence t
    = prompt (fn () => visit t)

  fun compute a
    = prompt (fn () => a END)
end
```

In the inductive case, the current delimited continuation (representing the current control queue) is captured, bound to `a`, and applied to `END`. The implicit continuation of this application visits `t1` and then `t2`, and therefore represents the queue `[t1, t2]`. Applying `a` seals it to the implicit continuation so that any continuation captured by a subsequent recursive call to `visit` in `a` captures both the rest of `a` and the traversal of `t1` and `t2`, i.e., the rest of the new control queue.

The above program could not be written with static control operators such as `shift` and `reset` without resorting to a non-trivial encoding of `control` and `prompt` in terms of `shift` and `reset` such as Shan’s [56]. In our experience, such an encoding obscures the code without shedding new light on dynamic delimited continuations.

4.3 Summary and conclusion

We first have presented a spectrum of solutions to the traditional depth-first samefringe problem. The one using `shift` and `reset` is new. We believe that connecting the lazy solution with McCarthy’s stack-based solution by defunctionalization is new as well.

By replacing the stack with a queue in the stack-based program, we have then obtained a solution to the breadth-first counterpart of the samefringe problem. Viewing this queue as a ‘data-structure continuation,’ we have observed that the operations upon it correspond to the operations induced by the composition of a dynamic delimited continuation and the current (delimited) continuation. We have then written this program in direct style using `control` and `prompt`.

5 Labeling a tree

We now turn to Okasaki’s problem of labeling a tree in breadth-first order with successive labels [48]. We express it in direct style with `control` and `prompt`, and we then outline its depth-first counterpart. Okasaki considers fully-labeled binary trees:

```
datatype tree = LEAF of int
              | NODE of tree * int * tree
```

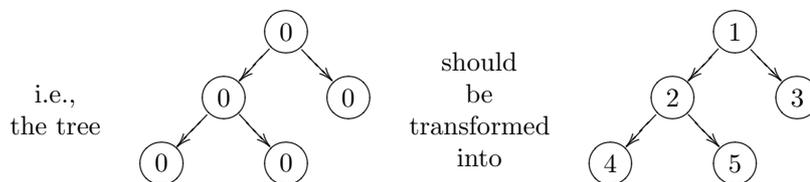
5.1 Breadth-first numbering

Given a tree T containing $|T|$ labels, we want to create a new tree of the same shape, but with the values in the nodes and leaves replaced by the numbers $1 \dots |T|$ in breadth-first order. For example, the tree

```
NODE (NODE (LEAF 0, 0, LEAF 0), 0, LEAF 0)
```

contains 5 labels and should be transformed into

```
NODE (NODE (LEAF 4, 2, LEAF 5), 1, LEAF 3)
```



5.1.1 A queue-based traversal

In his solution [48], Okasaki relabels a tree by mapping it recursively into a first-in, first-out list of subtrees (i.e., a forest) at call time and constructing the result at return time by reading this queue. To this end, he needs an auxiliary function

```
last_two_and_before : int list -> int list * int * int
```

such that applying it to the list $[x_n, \dots, x_3, x_2, x_1]$ yields the triple $([x_n, \dots, x_3], x_2, x_1)$.

Okasaki's solution reads as follows:

```
(* breadth_first_label : tree -> tree *)
fun breadth_first_label t
  = let (* visit : tree * int * tree list -> tree list *)
      fun visit (LEAF _, i, k)
        = (LEAF i) :: (continue (k, i+1))
      | visit (NODE (t1, _, t2), i, k)
        = let val (rest, t1', t2')
              = last_two_and_before
                (continue (k @ [t1, t2], i+1))
            in (NODE (t1', i, t2')) :: rest
            end
      (* continue : tree list * int -> tree list *)
      and continue (nil, _)
        = nil
      | continue (t :: k, i)
        = visit (t, i, k)
    in last (visit (t, 1, nil))
    end
```

where `last` is a function mapping a non-empty list to its last element.

The above algorithm uses two queues of trees (both representing a forest):

- the input queue, with function `visit` processing its front element, and with function `continue` processing its tail,
- the output backwards queue, which is enqueued in both clauses of function `visit`, and which is dequeued by functions `last_two_and_before` and `last`.

5.1.2 A direct-style traversal with control and prompt

As in Section 4.2.2, we observe that the concatenation, in the definition of `visit` just above, matches the concatenation of stack frames in the specification of `control` in Section 2.2. One can therefore write the function above in direct style, using `control` and `prompt`.

Again, as in Section 4.2.2, to be able to compose (concatenate) a captured delimited continuation with the current delimited continuation, we need to inject the type `int` into the type `list`, so that all the delimited continuations are endofunctions. To this end, we use a somewhat unusual data type of lists that end with an integer:

```

datatype 'a list_int = END of int
                    | CONS of 'a * 'a list_int

```

This unusual data type comes with otherwise usual list operations to test for the end of a integer-ended list and to fetch the head, the tail, or the last element of a non-empty integer-ended list:

```

fun last_int (CONS (x, xs))
  = let fun walk (candidate, END i)
        = candidate
        | walk (candidate, CONS (x, xs))
        = walk (x, xs)
    in walk (x, xs)
  end

```

We also need to adapt `last_two_and_before` to this data type.

The direct-style breadth-first numbering program reads as follows:

```

local structure CP = Control_and_Prompt
                    (type intermediate_answer = tree list_int)
in val control = CP.control
    val prompt = CP.prompt
end

(* breadth_first_label' : tree -> tree *)
fun breadth_first_label' t
  = let (* visit : tree * int -> tree list_int *)
        fun visit (LEAF _, i)
          = control
            (fn k =>
              CONS (LEAF i, k (END (i+1))))
        | visit (NODE (t1, _, t2), i)
          = control
            (fn k =>
              let val (rest, t1', t2')
                = last_two_and_before_int
                  (prompt
                   (fn () =>
                     let val (END i1) = k (END (i+1))
                         val (END i2) = visit (t1, i1)
                     in visit (t2, i2)
                     end))
                in CONS (NODE (t1', i, t2'), rest)
                end)
            in last_int (prompt (fn () => visit (t, 1)))
          end

```

In the inductive case, the current delimited continuation (of `visit`) is captured, bound to `k`, and applied to the integer end of a list. The implicit continuation of this application visits `t1` and then `t2`. Applying `k` seals it to the implicit continuation so that any continuation captured by an ulterior recursive call to `visit` in `k` captures both the rest of `k` and the visit of `t1` and `t2`.

The above solution required resorting to the change of representation of the output queue from `tree list` to `tree list_int` in order to unify the type `int` of the threaded index and the type `tree list` of the computation. Since all the captured continuations are endofunctions of type `tree list_int -> tree list_int`, we can compose them in the inductive case of `visit`. When the last leaf in the tree is visited, the captured continuation is applied to the number $|T| + 1$ (where T is the input tree and $|T|$ is the number of its labels), and this number becomes the tail of the output queue, as all the leaves are prepended to it in breadth-first order.

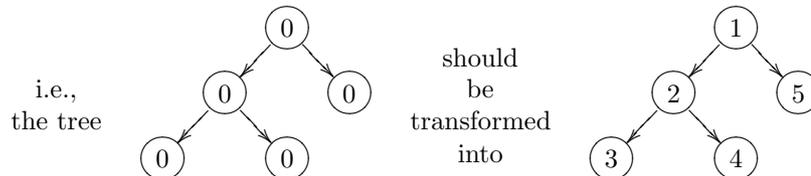
5.2 Depth-first numbering

We now turn to the depth-first counterpart of Okasaki's pearl, and present a spectrum of solutions to the problem of depth-first tree numbering. Given a tree T containing $|T|$ labels, we want to create a new tree of the same shape, but with the values in the nodes and leaves replaced by the numbers $1 \dots |T|$ in depth-first order. For example, the tree

```
NODE (NODE (LEAF 0, 0, LEAF 0), 0, LEAF 0)
```

should be transformed into

```
NODE (NODE (LEAF 3, 2, LEAF 4), 1, LEAF 5)
```



5.2.1 A stack-based traversal

It is trivial to write the depth-first counterpart of Okasaki's solution: one should just replace the queue with a stack, and instead of using `last_two_and_before`, use the auxiliary function

```
first_two_and_after : int list -> int * int * int list
```

such that applying it to the list $[x_1, x_2, x_3, \dots, x_n]$ yields the triple $(x_1, x_2, [x_3, \dots, x_n])$.

The depth-first solution reads as follows:

```

(* depth_first_label : tree -> tree *)
fun depth_first_label t
  = let (* visit : tree * int * tree list -> tree list *)
      fun visit (LEAF _, i, ts)
        = (LEAF i) :: (continue (ts, i+1))
      | visit (NODE (t1, _, t2), i, ts)
        = let val (t1', t2', rest)
              = first_two_and_after
                (continue (t1 :: t2 :: ts, i+1))
            in (NODE (t1', i, t2')) :: rest
            end
      (* continue : tree list * int -> tree list *)
      and continue (nil, _)
        = nil
      | continue (t :: k, i)
        = visit (t, i, k)
    in hd (visit (t, 1, nil))
    end

```

In contrast to Section 5.1.1, where the clause for nodes was concatenating the two subtrees in the back of the list of subtrees, in a first-in, first-out fashion,

```

      last_two_and_before
        (continue (k @ [t1, t2], i+1))    (* then *)

```

the clause for nodes is (essentially) concatenating the two subtrees in front of the list of subtrees, in a last-in, first-out fashion:

```

      first_two_and_after
        (continue ([t1, t2] @ ts, i+1))    (* now *)

```

We can see that the algorithm uses two stacks of trees (both representing a forest):

- the input stack, with function `visit` processing its top element, and with function `continue` processing its tail,
- the output stack, which is pushed on in both clauses of function `visit`, and which is popped off by functions `first_two_and_after` and `hd`.

5.2.2 A continuation-based traversal

In the induction case of `visit`, let us unfold the call to `continue` to obtain the following clause:

```

      | visit (NODE (t1, _, t2), i, ts)
        = let val (t1', t2', rest)
              = first_two_and_after
                (visit (t1, i+1, t2 :: ts))
            in (NODE (t1', i, t2')) :: rest
            end

```

The modified definition is in defunctionalized form: the data type is that of lists and `continue` is the corresponding apply function. The higher-order counterpart of this defunctionalized definition reads as follows:

```
(* depth_first_label' : tree -> tree *)
fun depth_first_label' t
  = let (* visit : tree * int * (int -> tree list) -> tree list *)
      fun visit (LEAF _, i, k)
        = (LEAF i) :: (k (i+1))
      | visit (NODE (t1, _, t2), i, k)
        = let val (t1', t2', rest)
            = first_two_and_after
              (visit (t1, i+1, fn j => visit (t2, j, k)))
          in (NODE (t1', i, t2')) :: rest
        end
      in hd (visit (t, 1, fn i => nil))
    end
```

5.2.3 A direct-style traversal with shift and reset

We view the function of type `int -> tree list`, in the definition just above, as a delimited continuation. This delimited continuation is initialized in the initial call to `visit`, extended in the inductive case, and captured and resumed in both clauses of `visit`. In direct style, the initialization is obtained with `reset`, the extension is obtained by functional sequencing, the capture is obtained with `shift`, and the resumption is obtained by function application. The result is another new example of programming with static delimited-control operators:

```
local structure SR = Shift_and_Reset
  (type intermediate_answer = tree list)
in val shift = SR.shift
    val reset = SR.reset
end
```

```

(* depth_first_label'' : tree -> tree *)
fun depth_first_label'' t
  = let (* visit : tree * int -> tree list *)
      fun visit (LEAF _, i)
        = shift
          (fn k =>
            (LEAF i) :: (k (i+1)))
        | visit (NODE (t1, _, t2), i)
          = shift
            (fn k =>
              let val (t1', t2', rest)
                = first_two_and_after
                  (reset
                    (fn () => k (let val j = visit (t1, i+1)
                                in visit (t2, j)
                                end)))
                in (NODE (t1', i, t2')) :: rest
              end)
            in hd (reset (fn () => let val i = visit (t, 1)
                                in nil
                                end))
          end)
    end
end

```

CPS-transforming `visit` yields the definition of Section 5.2.2.

5.3 Summary and conclusion

Okasaki's solution relabels its input tree in breadth-first order and uses a queue. We have expressed it in direct style using `control` and `prompt`. In so doing, we have internalized the explicit data operations on the queue into implicit control operations. These control operations crucially involve delimited continuations whose extent is dynamic.

The stack-based counterpart of Okasaki's solution relabels its input tree in depth-first order. We have mechanically refunctionalized this program into another one, which is continuation-based, and we have expressed this continuation-based program in direct style using `shift` and `reset`. These control operators crucially involve delimited continuations whose extent is static.

6 Conclusion and issues

Over the last 15 years, it has been repeatedly claimed that `control` has more expressive power than `shift`. Even though Shan has now disproved this claim [56], his simulation does not shed new light on how to program with `control`-like dynamic delimited continuations. In fact, in 15 years, only toy examples have been advanced to illustrate the difference between static and dynamic delimited continuations, such as the one in Section 2.5.

In this article, we have filled this vacuum by using dynamic delimited continuations to program breadth-first traversal. We have accounted for the dynamic queuing mechanism inherent to breadth-first traversal with the dynamic concatenation of stack frames that is specific to `control` and that makes it go beyond what is traditionally agreed upon as being continuation-passing style (CPS). We have presented two examples of breadth-first traversal: the breadth-first counterpart of the traditional `samefringe` function and Okasaki’s breadth-first numbering pearl.

One lesson we have learned here is how helpless one can feel when going beyond CPS. Unlike with `shift` and `reset`, there is no infrastructure for transforming programs that use `control` and `prompt`. We have therefore relied on CPS and on defunctionalization as guidelines, and we have built on the vision of data-structure continuations (stacks for depth-first traversals and queues for breadth-first traversals) proposed by Friedman 25 years ago [63, page 179] to infer the breadth-first traversals. We would have been hard pressed to come up with these examples only by groping for delimited continuations in direct style.¹

Since `control`, even more dynamic delimited-control operators have been proposed [30, 33, 44, 46, 52], all of which go beyond CPS but only two of which, to the best of our knowledge, come with motivating examples illustrating their dynamicity:

- In his PhD thesis [2], Balat uses the extra expressive power of Gunter, Rémy, and Riecke’s control operators `set` and `cupto` over that of `shift` and `reset` to prototype a type-directed partial evaluator for the lambda-calculus with sums [3, 4].
- In his PhD thesis [46], Nanevski introduces two new dynamic delimited-control operators, `mark` and `recall`, and illustrates them with a function partitioning a natural number into the lists of natural numbers that add to it. He considers both depth-first and breadth-first generation strategies, and conjectures that the latter cannot be written using `shift` and `reset`. As such, his is our closest related work.

These applications are rare and they tend to be daunting. Dynamic delimited continuations need simpler examples, more reasoning tools, and more meaning-preserving program transformations.

¹“You are not Superman.” – Aunt May (2002)

Acknowledgments: We are grateful to Mads Sig Ager, Małgorzata Biernacka, Andrzej Filinski, Julia Lawall, and Kevin Millikin for their comments.

Ken Shan pointed out that our initial writing of

```
control (fn a' => a' (a END))
```

in Section 4.2.2, could be simplified to

```
a END
```

The second author would also like to thank Andrzej Filinski, Mayer Goldberg, Julia Lawall, and Olin Shivers for participating to a one-week brain storm about continuations in August 2004, and to BRICS and DAIMI for hosting us during that week. Thanks are also due to Bernd Grobauer and Julia Lawall for sharing brain cycles about breadth-first numbering in Summer 2000.

This work is partially supported by the ESPRIT Working Group APPSEM II (<http://www.appsem.org>) and by the Danish Natural Science Research Council, Grant no. 21-03-0545.

A An implementation of shift and reset

Filinski's functor `Shift_and_Reset` maps a type of intermediate answers into a structure providing instances of `shift` and `reset` at that type [25]:

```
signature ESCAPE
= sig
  type void
  val coerce : void -> 'a
  val escape : (('a -> void) -> 'a) -> 'a
end

structure Escape : ESCAPE
= struct
  datatype void = VOID of void
  fun coerce (VOID v) = coerce v
  local open SMLofNJ.Cont
  in fun escape f
      = callcc (fn k => f (fn x => throw k x))
  end
end

signature SHIFT_AND_RESET
= sig
  type intermediate_answer
  val shift : (('a -> intermediate_answer) -> intermediate_answer)
            -> 'a
  val reset : (unit -> intermediate_answer) -> intermediate_answer
end
```

```

functor Shift_and_Reset (type intermediate_answer) : SHIFT_AND_RESET
= struct
  open Escape

  exception MISSING_RESET

  val mk : (intermediate_answer -> void) ref
      = ref (fn _ => raise MISSING_RESET)

  fun abort x
      = coerce (!mk x)

  type intermediate_answer = intermediate_answer

  fun reset thunk
      = escape (fn k => let val m = !mk
                        in mk := (fn r => (mk := m; k r));
                        abort (thunk ())
                        end)

  fun shift function
      = escape
        (fn k => abort (function (fn v => reset
                                (fn () => coerce (k v))))))

end

```

B An implementation of control and prompt

The functor `Control_and_Prompt` maps a type of intermediate answers into a structure providing instances of control and prompt at that type:

```

signature CONTROL_AND_PROMPT
= sig
  type intermediate_answer
  val control : ((intermediate_answer -> intermediate_answer)
                -> intermediate_answer)
                -> intermediate_answer
  val prompt : (unit -> intermediate_answer) -> intermediate_answer
end

functor Control_and_Prompt (type intermediate_answer)
: CONTROL_AND_PROMPT
= struct
  datatype ('t, 'w) context'
      = CONTEXT of 't -> ('w, 'w) context' option -> 'w

  fun send v NONE
      = v
    | send v (SOME (CONTEXT mc))
      = mc v NONE

```

```

fun compose (CONTEXT c, NONE)
  = CONTEXT c
  | compose (CONTEXT c, SOME mc1)
    = CONTEXT (fn v => fn mc2 => c v (SOME (compose (mc1, mc2))))

structure SR
= Shift_and_Reset
  (type intermediate_answer
   = (intermediate_answer, intermediate_answer) context' option
   -> intermediate_answer)
val shift = SR.shift
val reset = SR.reset

type intermediate_answer = intermediate_answer

fun prompt thunk
  = reset (fn () => send (thunk ())) NONE

exception MISSING_PROMPT

fun control function
  = shift
    (fn c1 =>
     fn mc1 =>
     let val k
       = fn x =>
         shift
           (fn c2 =>
            fn mc2 =>
              let val (CONTEXT c1') = compose (CONTEXT c1, mc1)
                in c1' x (SOME (compose (CONTEXT c2, mc2)))
              end)
            in reset (fn () => send (function k)) NONE
            end) handle MISSING_RESET => raise MISSING_PROMPT
end)
end

```

References

- [1] Harold Abelson and Gerald Jay Sussman with Julie Sussman. *Structure and Interpretation of Computer Programs*. The MIT Press, Cambridge, Massachusetts, 1985.
- [2] Vincent Balat. *Une étude des sommes fortes: isomorphismes et formes normales*. PhD thesis, PPS, Université Denis Diderot (Paris VII), Paris, France, December 2002.
- [3] Vincent Balat, Roberto Di Cosmo, and Marcelo P. Fiore. Extensional normalisation and type-directed partial evaluation for typed lambda calculus with sums. In Xavier Leroy, editor, *Proceedings of the Thirty-First Annual ACM Symposium on Principles of Programming Languages*, pages 64–76, Venice, Italy, January 2004. ACM Press.

- [4] Vincent Balat and Olivier Danvy. Memoization in type-directed partial evaluation. In Don Batory, Charles Consel, and Walid Taha, editors, *Proceedings of the 2002 ACM SIGPLAN/SIGSOFT Conference on Generative Programming and Component Engineering*, number 2487 in Lecture Notes in Computer Science, pages 78–92, Pittsburgh, Pennsylvania, October 2002. Springer-Verlag.
- [5] Josh Berdine. *Linear and Affine Typing of Continuation-Passing Style*. PhD thesis, Queen Mary, University of London, 2004.
- [6] Małgorzata Biernacka, Dariusz Biernacki, and Olivier Danvy. An operational foundation for delimited continuations in the CPS hierarchy. Technical Report BRICS RS-04-29, DAIMI, Department of Computer Science, University of Aarhus, Aarhus, Denmark, December 2004. A preliminary version was presented at the the Fourth ACM SIGPLAN Workshop on Continuations (CW 2004).
- [7] Mats Carlsson. On implementing Prolog in functional programming. *New Generation Computing*, 2(4):347–359, 1984.
- [8] Robert (Corky) Cartwright, editor. *Proceedings of the Eighteenth Annual ACM Symposium on Principles of Programming Languages*, Orlando, Florida, January 1991. ACM Press.
- [9] Eugene Charniak, Christopher Riesbeck, and Drew McDermott. *Artificial Intelligence Programming*. Lawrence Earlbaum Associates, 1980.
- [10] William Clinger, Daniel P. Friedman, and Mitchell Wand. A scheme for a higher-level semantic algebra. In John Reynolds and Maurice Nivat, editors, *Algebraic Methods in Semantics*, pages 237–250. Cambridge University Press, 1985.
- [11] Olivier Danvy. On some functional aspects of control. In Thomas Johnsson, Simon Peyton Jones, and Kent Karlsson, editors, *Proceedings of the Workshop on Implementation of Lazy Functional Languages*, pages 445–449. Program Methodology Group, University of Göteborg and Chalmers University of Technology, September 1988. Report 53.
- [12] Olivier Danvy. Programming with tighter control. *Special issue of the Bigre journal: Putting Scheme to Work*, 65:10–29, July 1989.
- [13] Olivier Danvy. On evaluation contexts, continuations, and the rest of the computation. In Hayo Thielecke, editor, *Proceedings of the Fourth ACM SIGPLAN Workshop on Continuations*, Technical report CSR-04-1, Department of Computer Science, Queen Mary’s College, pages 13–23, Venice, Italy, January 2004. Invited talk.
- [14] Olivier Danvy and Andrzej Filinski. A functional abstraction of typed contexts. DIKU Rapport 89/12, DIKU, Computer Science Department, University of Copenhagen, Copenhagen, Denmark, July 1989.

- [15] Olivier Danvy and Andrzej Filinski. Abstracting control. In Wand [64], pages 151–160.
- [16] Olivier Danvy and Andrzej Filinski. Representing control, a study of the CPS transformation. *Mathematical Structures in Computer Science*, 2(4):361–391, 1992.
- [17] Olivier Danvy and Lasse R. Nielsen. Defunctionalization at work. In Harald Søndergaard, editor, *Proceedings of the Third International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP'01)*, pages 162–174, Firenze, Italy, September 2001. ACM Press. Extended version available as the technical report BRICS RS-01-23.
- [18] Bruce F. Duba, Robert Harper, and David B. MacQueen. Typing first-class continuations in ML. In Cartwright [8], pages 163–173.
- [19] Matthias Felleisen. *The Calculi of λ -v-CS Conversion: A Syntactic Theory of Control and State in Imperative Higher-Order Programming Languages*. PhD thesis, Department of Computer Science, Indiana University, Bloomington, Indiana, August 1987.
- [20] Matthias Felleisen. The theory and practice of first-class prompts. In Ferrante and Mager [24], pages 180–190.
- [21] Matthias Felleisen and Daniel P. Friedman. Control operators, the SECD machine, and the λ -calculus. In Martin Wirsing, editor, *Formal Description of Programming Concepts III*, pages 193–217. Elsevier Science Publishers B.V. (North-Holland), Amsterdam, 1986.
- [22] Matthias Felleisen, Daniel P. Friedman, Bruce Duba, and John Merrill. Beyond continuations. Technical Report 216, Computer Science Department, Indiana University, Bloomington, Indiana, February 1987.
- [23] Matthias Felleisen, Mitchell Wand, Daniel P. Friedman, and Bruce F. Duba. Abstract continuations: A mathematical semantics for handling full functional jumps. In Robert (Corky) Cartwright, editor, *Proceedings of the 1988 ACM Conference on Lisp and Functional Programming*, pages 52–62, Snowbird, Utah, July 1988. ACM Press.
- [24] Jeanne Ferrante and Peter Mager, editors. *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Programming Languages*, San Diego, California, January 1988. ACM Press.
- [25] Andrzej Filinski. Representing monads. In Hans-J. Boehm, editor, *Proceedings of the Twenty-First Annual ACM Symposium on Principles of Programming Languages*, pages 446–457, Portland, Oregon, January 1994. ACM Press.
- [26] Andrzej Filinski. *Controlling Effects*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, May 1996. Technical Report CMU-CS-96-119.

- [27] Daniel P. Friedman and David S. Wise. CONS should not evaluate its arguments. In S. Michaelson and Robin Milner, editors, *Third International Colloquium on Automata, Languages, and Programming*, pages 257–284. Edinburgh University Press, Edinburgh, Scotland, July 1976.
- [28] Carsten Führmann. *The Structure of Call-by-Value*. PhD thesis, University of Edinburgh, Edinburgh, Scotland, 2000.
- [29] Timothy G. Griffin. A formulae-as-types notion of control. In Paul Hudak, editor, *Proceedings of the Seventeenth Annual ACM Symposium on Principles of Programming Languages*, pages 47–58, San Francisco, California, January 1990. ACM Press.
- [30] Carl Gunter, Didier Rémy, and Jon G. Riecke. A generalization of exceptions and control in ML-like languages. In Simon Peyton Jones, editor, *Proceedings of the Seventh ACM Conference on Functional Programming and Computer Architecture*, pages 12–23, La Jolla, California, June 1995. ACM Press.
- [31] John Hatcliff. *The Structure of Continuation-Passing Styles*. PhD thesis, Department of Computing and Information Sciences, Kansas State University, Manhattan, Kansas, June 1994.
- [32] Peter Henderson and James H. Morris Jr. A lazy evaluator. In Susan L. Graham, editor, *Proceedings of the Third Annual ACM Symposium on Principles of Programming Languages*, pages 95–103. ACM Press, January 1976.
- [33] Robert Hieb, R. Kent Dybvig, and Claude W. Anderson, III. Subcontinuations. *Lisp and Symbolic Computation*, 5(4):295–326, December 1993.
- [34] Gregory F. Johnson. GL – a denotational testbed with continuations and partial continuations as first-class objects. In Mark Scott Johnson, editor, *Proceedings of the ACM SIGPLAN’87 Symposium on Interpreters and Interpretive Techniques*, SIGPLAN Notices, Vol. 22, No 7, pages 154–176, Saint-Paul, Minnesota, June 1987. ACM Press.
- [35] Gregory F. Johnson and Dominic Duggan. Stores and partial continuations as first-class objects in a language and its environment. In Ferrante and Mager [24], pages 158–168.
- [36] Richard Kelsey, William Clinger, and Jonathan Rees, editors. Revised⁵ report on the algorithmic language Scheme. *Higher-Order and Symbolic Computation*, 11(1):7–105, 1998.
- [37] Jim Laird. *A semantic analysis of control*. PhD thesis, University of Edinburgh, Edinburgh, Scotland, 1998.
- [38] Peter Landin. A generalization of jumps and labels. *Higher-Order and Symbolic Computation*, 11(1):125–143, 1998.
- [39] Julia L. Lawall. *Continuation Introduction and Elimination in Higher-Order Programming Languages*. PhD thesis, Computer Science Department, Indiana University, Bloomington, Indiana, July 1994.

- [40] John McCarthy. Another samefringe. *SIGART Newsletter*, 61, February 1977.
- [41] Chris Mellish and Steve Hardy. Integrating Prolog in the POPLOG environment. In John A. Campbell, editor, *Implementations of PROLOG*, pages 147–162. Ellis Horwood, 1984.
- [42] Albert R. Meyer and Mitchell Wand. Continuation semantics in typed lambda-calculi (summary). In Rohit Parikh, editor, *Logics of Programs – Proceedings*, number 193 in Lecture Notes in Computer Science, pages 219–224, Brooklyn, New York, June 1985. Springer-Verlag.
- [43] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML (Revised)*. The MIT Press, 1997.
- [44] Luc Moreau and Christian Queinnec. Partial continuations as the difference of continuations, a duumvirate of control operators. In Manuel Hermenegildo and Jaan Penjam, editors, *Sixth International Symposium on Programming Language Implementation and Logic Programming*, number 844 in Lecture Notes in Computer Science, pages 182–197, Madrid, Spain, September 1994. Springer-Verlag.
- [45] Chetan R. Murthy. *Extracting Constructive Content from Classical Proofs*. PhD thesis, Department of Computer Science, Cornell University, Ithaca, New York, 1990.
- [46] Aleksandar Nanevski. *Functional Programming with Names and Necessity*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, June 2004. Technical Report CMU-CS-04-151.
- [47] Lasse R. Nielsen. *A study of defunctionalization and continuation-passing style*. PhD thesis, BRICS PhD School, University of Aarhus, Aarhus, Denmark, July 2001. BRICS DS-01-7.
- [48] Chris Okasaki. Breadth-first numbering: lessons from a small exercise in algorithm design. In Philip Wadler, editor, *Proceedings of the 2000 ACM SIGPLAN International Conference on Functional Programming*, SIGPLAN Notices, Vol. 35, No. 9, pages 131–136, Montréal, Canada, September 2000. ACM Press.
- [49] Gordon D. Plotkin. Call-by-name, call-by-value and the λ -calculus. *Theoretical Computer Science*, 1:125–159, 1975.
- [50] Jeff Polakow. *Ordered Linear Logic and Applications*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, August 2001. Technical Report CMU-CS-01-152.
- [51] Christian Queinnec. Value transforming style. In *Proceedings of the Second International Workshop on Static Analysis WSA '92*, volume 81-82 of *Bigre Journal*, pages 20–28, Bordeaux, France, September 1992. IRISA, Rennes, France.

- [52] Christian Queinnec and Bernard Serpette. A dynamic extent control operator for partial continuations. In Cartwright [8], pages 174–184.
- [53] John C. Reynolds. Definitional interpreters for higher-order programming languages. *Higher-Order and Symbolic Computation*, 11(4):363–397, 1998. Reprinted from the proceedings of the 25th ACM National Conference (1972), with a foreword.
- [54] Amr Sabry. *The Formal Relationship between Direct and Continuation-Passing Style Optimizing Compilers: A Synthesis of Two Paradigms*. PhD thesis, Computer Science Department, Rice University, Houston, Texas, August 1994. Technical report 94-242.
- [55] Erik Sandewall. An early use of continuations and partial evaluation for compiling rules written in FOPC. *Higher-Order and Symbolic Computation*, 12(1):105–113, 1999.
- [56] Chung-chieh Shan. Shift to control. In Olin Shivers and Oscar Waddell, editors, *Proceedings of the 2004 ACM SIGPLAN Workshop on Scheme and Functional Programming*, Snowbird, Utah, September 2004.
- [57] Dorai Sitaram. *Models of Control and their Implications for Programming Language Design*. PhD thesis, Computer Science Department, Rice University, Houston, Texas, April 1994.
- [58] Dorai Sitaram and Matthias Felleisen. Reasoning with continuations II: Full abstraction for models of control. In Wand [64], pages 161–175.
- [59] Guy L. Steele Jr. Rabbit: A compiler for Scheme. Master’s thesis, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, May 1978. Technical report AI-TR-474.
- [60] Christopher Strachey and Christopher P. Wadsworth. Continuations: A mathematical semantics for handling full jumps. *Higher-Order and Symbolic Computation*, 13(1/2):135–152, 2000. Reprint of the technical monograph PRG-11, Oxford University Computing Laboratory (1974), with a foreword.
- [61] Gerald J. Sussman and Guy L. Steele Jr. Scheme: An interpreter for extended lambda calculus. *Higher-Order and Symbolic Computation*, 11(4):405–439, 1998. Reprinted from the AI Memo 349, MIT (1975), with a foreword.
- [62] Hayo Thielecke. *Categorical Structure of Continuation Passing Style*. PhD thesis, University of Edinburgh, Edinburgh, Scotland, 1997. ECS-LFCS-97-376.
- [63] Mitchell Wand. Continuation-based program transformation strategies. *Journal of the ACM*, 27(1):164–180, January 1980.
- [64] Mitchell Wand, editor. *Proceedings of the 1990 ACM Conference on Lisp and Functional Programming*, Nice, France, June 1990. ACM Press.

Recent BRICS Report Series Publications

- RS-05-2 Dariusz Biernacki and Olivier Danvy. *On the Dynamic Extent of Delimited Continuations*. January 2005. ii+30 pp.
- RS-05-1 Mayer Goldberg. *On the Recursive Enumerability of Fixed-Point Combinators*. January 2005. 7 pp. Superseeds BRICS report RS-04-25.
- RS-04-41 Olivier Danvy. *Sur un Exemple de Patrick Greussay*. December 2004. 14 pp.
- RS-04-40 Mads Sig Ager, Olivier Danvy, and Henning Korsholm Rohde. *Fast Partial Evaluation of Pattern Matching in Strings*. December 2005. 22 pp. To appear in TOPLAS. Supersedes BRICS report RS-03-20.
- RS-04-39 Olivier Danvy and Lasse R. Nielsen. *CPS Transformation of Beta-Redexes*. December 2004. ii+11 pp. Superseeds an article to appear in *Information Processing Letters* and BRICS report RS-00-35.
- RS-04-38 Olin Shivers and Mitchell Wand. *Bottom-Up β -Substitution: Uplinks and λ -DAGs*. December 2004.
- RS-04-37 Jørgen Iversen and Peter D. Mosses. *Constructive Action Semantics for Core ML*. December 2004. 68 pp. To appear in a special *Language Definitions and Tool Generation* issue of the journal *IEE Proceedings Software*.
- RS-04-36 Mark van den Brand, Jørgen Iversen, and Peter D. Mosses. *An Action Environment*. December 2004. 27 pp. Appears in Hedin and Van Wyk, editors, *Fourth ACM SIGPLAN Workshop on Language Descriptions, Tools and Applications, LDTA '04, 2004*, pages 149–168.
- RS-04-35 Jørgen Iversen. *Type Checking Semantic Functions in ASDF*. December 2004.
- RS-04-34 Anders Møller and Michael I. Schwartzbach. *The Design Space of Type Checkers for XML Transformation Languages*. December 2004. 21 pp. Appears in Eiter and Libkin, editors, *Database Theory: 10th International Conference, ICDT '05 Proceedings, LNCS 3363, 2005*, pages 17–36.