Basic Research in Computer Science

BRICS RS-01-23 Danvy & Nielsen: Defunctionalization at Work

# Defunctionalization at Work

**Olivier Danvy**
**Lasse R. Nielsen**

**BRICS Report Series** **RS-01-23**

**ISSN 0909-0878** **June 2001**

**Copyright © 2001, Olivier Danvy & Lasse R. Nielsen.
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/01/23/`

# Defunctionalization at Work *

Olivier Danvy and Lasse R. Nielsen

BRICS †
Department of Computer Science
University of Aarhus ‡

June, 2001

**Abstract**

Reynolds's defunctionalization technique is a whole-program transformation from higher-order to first-order functional programs. We study practical applications of this transformation and uncover new connections between seemingly unrelated higher-order and first-order specifications and between their correctness proofs. Defunctionalization therefore appears both as a springboard for revealing new connections and as a bridge for transferring existing results between the first-order world and the higher-order world.

*Extended version of an article to appear in the proceedings of PPDP 2001, Firenze, Italy.
†Basic Research in Computer Science (`www.brics.dk`),
funded by the Danish National Research Foundation.
‡Ny Munkegade, Building 540, DK-8000 Aarhus C, Denmark
E-mail: {`danvy,lrn`}`@brics.dk`

## Contents


**1 Background and Introduction** **4**
- 1.1 A sample higher-order program with a static number of closures 4
- 1.2 A sample higher-order program with a dynamic number of closures 5
- 1.3 Defunctionalization in a nutshell . . . . . 7
- 1.4 Related work . . . . . 7
- 1.5 This work . . . . . 7

**2 Defunctionalization of List- and of Tree-Processing Programs** **9**
- 2.1 Flattening a binary tree into a list . . . . . 9
- 2.2 Higher-order representations of lists . . . . . 10
- 2.3 Defunctionalizing Church-encoded non-recursive data structures . 13
- 2.4 Defunctionalizing Church-encoded recursive data structures . . . 15
- 2.5 Church-encoding the result of defunctionalization . . . . . 16
- 2.6 Summary and conclusion . . . . . 17

**3 Defunctionalization of CPS-Transformed First-Order Programs** **17**
- 3.1 String parsing . . . . . 17
- 3.2 Continuation-based program transformation strategies . . . . . 20
- 3.3 Summary and conclusion . . . . . 20

**4 Two Syntactic Theories in the Light of Defunctionalization** **20**
- 4.1 Arithmetic expressions . . . . . 20
  - 4.1.1 A syntactic theory for arithmetic expressions . . . . . 21
  - 4.1.2 Implementation . . . . . 21
  - 4.1.3 Refunctionalization . . . . . 22
  - 4.1.4 Back to direct style . . . . . 23
- 4.2 The call-by-value $\lambda$-calculus . . . . . 24
  - 4.2.1 A syntactic theory for the call-by-value $\lambda$-calculus . . . . 24
  - 4.2.2 Implementation . . . . . 24
  - 4.2.3 Refunctionalization . . . . . 25
  - 4.2.4 Back to direct style . . . . . 26
- 4.3 Summary and conclusion . . . . . 27

**5 A Comparison between Correctness Proofs before and after Defunctionalization: Matching Regular Expressions** **28**
- 5.1 Regular expressions . . . . . 28
- 5.2 The two matchers . . . . . 29
  - 5.2.1 The higher-order matcher . . . . . 29
  - 5.2.2 The first-order matcher . . . . . 30
- 5.3 The two correctness proofs . . . . . 30
  - 5.3.1 Correctness proof of the higher-order matcher . . . . . 32
  - 5.3.2 Correctness proof of the first-order matcher . . . . . 32
- 5.4 Comparison between the two correctness proofs . . . . . 33
- 5.5 Summary and conclusion . . . . . 34

**6 Conclusions and Issues** **35**

**A Correctness Proof of the Higher-Order Matcher** **36**

**B Correctness Proof of the First-Order Matcher** **39**

## List of Figures

1 Higher-order, continuation-based matcher for regular expressions 29
2 First-order, stack-based matcher for regular expressions . . . . . 31

# 1 Background and Introduction

In first-order programs, all functions are named and each call refers to the callee by its name. In higher-order programs, functions may be anonymous, passed as arguments, and returned as results. As Strachey put it [50], functions are *second-class* denotable values in a first-order program, and *first-class* expressible values in a higher-order program. One may then wonder how first-class functions are represented at run time.

- First-class functions are often represented with *closures*, i.e., expressible values pairing a code pointer and the denotable values of the variables occurring free in that code, as proposed by Landin in the mid-1960's [30]. Today, closures are the most common representation of first-class functions in the world of eager functional programming [1, 17, 32], as well as a standard representation for implementing object-oriented programs [23]. They are also used to implement higher-order logic programming [8].

- Alternatively, higher-order programs can be *defunctionalized* into first-order programs, as proposed by Reynolds in the early 1970's [43]. In a defunctionalized program, first-class functions are represented with first-order data types: a first-class function is introduced with a constructor holding the values of the free variables of a function abstraction, and it is eliminated with a case expression dispatching over the corresponding constructors.

- First-class functions can also be dealt with by translating functional programs into *combinators* and using graph reduction, as proposed by Turner in the mid-1970's [52]. This implementation technique has been investigated extensively in the world of lazy functional programming [27, 29, 37, 38].

Compared to closure conversion and to combinator conversion, defunctionalization has been used very little. The goal of this article is to study practical applications of it.

We first illustrate defunctionalization with two concrete examples (Sections 1.1 and 1.2). In the first program, two function abstractions are instantiated once, and in the second program, one function abstraction is instantiated repeatedly. We then characterize defunctionalization in a nutshell (Section 1.3) before reviewing related work (Section 1.4). Finally, we raise questions to which defunctionalization provides answers (Section 1.5).

## 1.1 A sample higher-order program with a static number of closures

In the following ML program, `aux` is passed a first-class function, applies it to 1 and 10, and sums the results. The `main` function calls `aux` twice and multiplies the results. All in all, two function abstractions occur in this program, in `main`.

```
(*  aux : (int -> int) -> int  *)
fun aux f
    = f 1 + f 10

(*  main : int * int * bool -> int  *)
fun main (x, y, b)
    = aux (fn z => x + z) *
      aux (fn z => if b then y + z else y - z)
```

Defunctionalizing this program amounts to defining a data type with two constructors, one for each function abstraction, and its associated apply function. The first function abstraction contains one free variable (`x`, of type integer), and therefore the first data-type constructor requires an integer. The second function abstraction contains two free variables (`y`, of type integer, and `b`, of type boolean), and therefore the second data-type constructor requires an integer and a boolean.

In `main`, the first first-class function is thus introduced with the first constructor and the value of `x`, and the second with the second constructor and the values of `y` and `b`.

In `aux`, the functional argument is passed to a second-class function `apply` that eliminates it with a case expression dispatching over the two constructors.

```
datatype lam = LAM1 of int
             | LAM2 of int * bool

(*  apply : lam * int -> int  *)
fun apply (LAM1 x, z)
    = x + z
  | apply (LAM2 (y, b), z)
    = if b then y + z else y - z

(*  aux_def : lam -> int  *)
fun aux_def f
    = apply (f, 1) + apply (f, 10)

(*  main_def : int * int * bool -> int  *)
fun main_def (x, y, b)
    = aux_def (LAM1 x) * aux_def (LAM2 (y, b))
```

### 1.2 A sample higher-order program with a dynamic number of closures

A PPDP reviewer wondered what happens for programs that "dynamically generate" new closures, and whether such programs lead to new constants and thus require extensible case expressions. The following example illustrates such a situation and shows that no new constants and no extensible case expressions are needed.

In the following ML program, `aux` is passed two arguments and applies one to the other. The `main` function is given a number `i` and a list of numbers `[j1, j2, ...]` and returns the list of numbers `[i+j1, i+j2, ...]`. One function abstraction, `fn i => i + j`, occurs in this program, in `main`, as the second argument of `aux`. Given an input list of length $n$, the function abstraction is instantiated $n$ times in the course of the computation.

```
(*  aux : int * (int -> int) -> int  *)
fun aux (i, f)
    = f i

(*  main = fn : int * int list -> int list  *)
fun main (i, js)
    = let fun walk nil
              = nil
            | walk (j :: js)
              = (aux (i, fn i => i + j)) :: (walk js)
      in walk js
      end
```

Defunctionalizing this program amounts to defining a data type with only one constructor, since there is only one function abstraction, and its associated apply function. The function abstraction contains one free variable (`j`, of type integer), and therefore the data-type constructor requires an integer.

In `main`, the first-class function is introduced with the constructor and the value of `j`.

In `aux`, the functional argument is passed to a second-class function `apply` that eliminates it with a case expression dispatching over the constructor.

```
datatype lam = LAM of int

(*  apply : lam * int -> int  *)
fun apply (LAM j, i)
    = i + j

(*  aux_def : int * lam -> int  *)
fun aux_def (i, f)
    = apply (f, i)

(*  main_def : int * int list -> int list  *)
fun main_def (i, js)
    = let fun walk nil
              = nil
            | walk (j :: js)
              = (aux_def (i, LAM j)) :: (walk js)
      in walk js
      end
```

Given an input list of length $n$, the constructor `LAM` is used $n$ times in the course of the computation.

### 1.3 Defunctionalization in a nutshell

In a higher-order program, first-class functions arise as instances of function abstractions. All these function abstractions can be enumerated in a whole program. Defunctionalization is thus *a whole-program transformation where function types are replaced by an enumeration of the function abstractions in this program.*

Defunctionalization therefore takes its roots in type theory. Indeed, a function type hides typing assumptions from the context, and, as pointed out by Minamide, Morrisett, and Harper in their work on typed closure conversion [32], making these assumptions explicit requires an existential type. For a whole program, this existential type can be represented with a finite sum together with the corresponding injections and case dispatch, and this representation is precisely what defunctionalization achieves.

These type-theoretical roots do not make defunctionalization a straitjacket, to paraphrase Reynolds about Algol [42]. For example, one can use several apply functions, e.g., grouped by types, as in Bell, Bellegarde, and Hook's work [4]. One can also defunctionalize a program selectively, e.g., only its continuations, as in Section 3. One can even envision a lightweight defunctionalization similar to Steckler and Wand's lightweight closure conversion [48], as in Banerjee, Heintze, and Riecke's recent work [2].

### 1.4 Related work

Originally, Reynolds devised defunctionalization to transform a higher-order interpreter into a first-order one [43]. He presented it as a programming technique, and never used it again [44], save for deriving a first-order semantics in his textbook on programming language [45, Section 12.4].

Since then, defunctionalization has not been used much, though when it has, it was as a full-fledged implementation technique: Bondorf uses it to make higher-order programs amenable to first-order partial evaluation [5]; Tolmach and Oliva use it to compile ML programs into Ada [51]; Fegaras uses it in his object-oriented database management system, lambda-DB [18]; Wang and Appel use it in type-safe garbage collectors [55]; and defunctionalization is an integral part of MLton [7] and of Boquist's Haskell compiler [6].

Only lately has defunctionalization been formalized: Bell, Bellegarde, and Hook showed that it preserves types [4]; Nielsen proved its partial correctness using denotational semantics [35, 36]; and Banerjee, Heintze and Riecke proved its total correctness using operational semantics [2].

### 1.5 This work

Functional programming encourages fold-like recursive descents, typically using auxiliary recursive functions. Often, these auxiliary functions are higher order in that their co-domain is a function space. For example, if an auxiliary function has an accumulator of type $\alpha$, its co-domain is $\alpha \to \beta$, for some $\beta$. For another example, if an auxiliary function has a continuation of type $\alpha \to \beta$, for some $\beta$,

its co-domain is $(\alpha \to \beta) \to \beta$. How do these functional programs compare to programs written using a first-order, data-structure oriented approach?

Wand's classical work on continuation-based program transformations [54] was motivated by the question "What is a data-structure continuation?". Each of the examples considered in Wand's paper required a eureka step to design a data structure for representing a continuation. Are such eureka steps always necessary?

Continuations are variously presented as a functional representation of the rest of the computation and as a functional representation of the context of a computation [20]. Wand's work addressed the former view, so let us consider the latter one. For example, in his PhD thesis [19], Felleisen developed a syntactic approach to semantics relying on the first-order notions of 'evaluation context' and of 'plugging expressions into contexts'. How do these first-order notions compare to the notion of continuation?

In the rest of this article, we show that defunctionalization provides a single answer to all these questions. All the programs we consider perform a recursive descent and use an auxiliary function. When this auxiliary function is higher-order, defunctionalization yields a first-order version with an accumulator (e.g., tree flattening in Section 2.1 and list reversal in Section 2.2). When this auxiliary function is first-order, we transform it into continuation-passing style; defunctionalization then yields an iterative first-order version with an accumulator in the form of a data structure (e.g., string parsing in Section 3.1 and regular-expression matching in Section 5). We also consider interpreters for two syntactic theories and we identify that they are written in a defunctionalized form. We then "refunctionalize" them and obtain continuation-passing interpreters whose continuations represent the evaluation contexts of the corresponding syntactic theory (Section 4).

In addition, we observe that defunctionalization and Church encoding have dual purposes, since Church encoding is a classical way to represent data structures with higher-order functions. What is the result of defunctionalizing a Church-encoded data structure? And what does one obtain when Church-encoding the result of defunctionalization?

Similarly, we observe that backtracking is variously implemented in a first-order setting with one or two stacks, and in a higher-order setting with one or two continuations. It is natural enough to wonder what is the result of Church-encoding the stacks and of defunctionalizing the continuations. One can wonder as well about the correctness proofs of these programs—how do they compare?

In the rest of this article, we also answer these questions. We defunctionalize two programs using Hughes's higher-order representation of intermediate lists and obtain two efficient and traditional first-order programs (Section 2.2). We also clarify the extent to which Church encoding and defunctionalization can be considered as inverses of each other (Sections 2.3, 2.4, and 2.5). Finally, we compare and contrast a regular-expression matcher and its proof before and after defunctionalization (Section 5).

## 2 Defunctionalization of List- and of Tree-Processing Programs

We consider several canonical higher-order programs over lists and trees and we defunctionalize them. In each case, defunctionalization yields a known, but unrelated solution. We then turn to Church encoding, which provides a uniform higher-order representation of data structures. We consider the result of defunctionalizing Church-encoded data structures, and we consider the result of Church-encoding the result of defunctionalization.

### 2.1 Flattening a binary tree into a list

To flatten a binary tree into a list of its leaves, we choose to map a leaf to a curried list constructor and a node to function composition, homomorphically. In other words, we map a list into the monoid of functions from lists to lists. This definition hinges on the built-in associativity of function composition.

```
datatype 'a bt = LEAF of 'a
               | NODE of 'a bt * 'a bt

(*  cons : 'a -> 'a list -> 'a list  *)
fun cons x xs
    = x :: xs

(*  flatten : 'a bt -> 'a list             *)
(*     walk : 'a bt -> 'a list -> 'a list  *)
fun flatten t
    = let fun walk (LEAF x)
              = cons x
            | walk (NODE (t1, t2))
              = (walk t1) o (walk t2)
      in walk t nil
      end
```

Eta-expanding the result of `walk` and inlining `cons` and `o` yields a curried version of the fast flatten function with an accumulator.

```
(*  flatten_ee : 'a bt -> 'a list             *)
(*        walk : 'a bt -> 'a list -> 'a list  *)

fun flatten_ee t
    = let fun walk (LEAF x) a
              = x :: a
            | walk (NODE (t1, t2)) a
              = walk t1 (walk t2 a)
      in walk t nil
      end
```

It is also instructive to defunctionalize `flatten`. Two functional values occur—one for the leaves and one for the nodes—and therefore they give rise to a data

type with two constructors. Since `flatten` is homomorphic, the new data type is isomorphic to the data type of binary trees, and therefore the associated apply function could be made to work directly on the input tree, e.g., using deforestation [53]. At any rate, we recognize this apply function as an uncurried version of the fast flatten function with an accumulator.

```
datatype 'a lam = LAM1 of 'a | LAM2 of 'a lam * 'a lam

(*  apply : 'a lam * 'a list -> 'a list  *)
fun apply (LAM1 x, xs)
    = x :: xs
  | apply (LAM2 (f1, f2), xs)
    = apply (f1, apply (f2, xs))

(*  cons_def : 'a -> 'a lam  *)
fun cons_def x
    = LAM1 x

(*  o_def : 'a lam * 'a lam -> 'a lam  *)
fun o_def (f1, f2)
    = LAM2 (f1, f2)

(*  flatten_def : 'a bt -> 'a list  *)
(*         walk : 'a bt -> 'a lam   *)
fun flatten_def t
    = let fun walk (LEAF x)
              = cons_def x
            | walk (NODE (t1, t2))
              = o_def (walk t1, walk t2)
      in apply (walk t, nil)
      end
```

The monoid of functions from lists to lists corresponds to Hughes's novel representations of lists [28], which we treat next.

### 2.2 Higher-order representations of lists

In the mid-1980's, Hughes proposed to represent intermediate lists as partially applied concatenation functions [28], so that instead of constructing a list `xs`, one instantiates the function abstraction `fn ys => xs @ ys`. The key property of this higher-order representation is that lists can be concatenated in constant time. Therefore, the following naive version of `reverse` operates in linear time instead of in quadratic time, as with the usual linked representation of lists, where lists are concatenated in linear time.

```
(*  append : 'a list -> 'a list -> 'a list  *)
fun append xs ys
    = xs @ ys
```

```
(*  reverse : 'a list -> 'a list  *)
fun reverse xs
    = let fun walk nil
              = append nil
            | walk (x :: xs)
              = (walk xs) o (append [x])
      in walk xs nil
      end
```

Let us defunctionalize this program. First, like Hughes, we recognize that appending the empty list is the identity function and that appending a single element amounts to consing it.

```
(*  id : 'a list -> 'a list  *)
fun id ys
    = ys

(*  cons : 'a -> 'a list -> 'a list *)
fun cons x xs
    = x :: xs

(* reverse : 'a list -> 'a list             *)
(*    walk : 'a list -> 'a list -> 'a list  *)
fun reverse xs
    = let fun walk nil
              = id
            | walk (x :: xs)
              = (walk xs) o (cons x)
      in walk xs nil
      end
```

The function space `'a list -> 'a list` arises because of three functional values: `id`, in one conditional branch; and, in the other, the results of consing an element and of calling `walk`.

We thus defunctionalize the program using a data type with three constructors and its associated apply function.

```
datatype 'a lam = LAM0
                | LAM1 of 'a
                | LAM2 of 'a lam * 'a lam

(*  apply : 'a lam * 'a list -> 'a list  *)
fun apply (LAM0, ys)
    = ys
  | apply (LAM1 x, ys)
    = x :: ys
  | apply (LAM2 (f, g), ys)
    = apply (f, apply (g, ys))
```

This data type makes it plain that in Hughes's monoid of intermediate lists, concatenation is performed in constant time (here with `LAM2`).

The rest of the defunctionalized program reads as follows.

```
(*  id_def : 'a lam  *)
val id_def = LAM0

(*  cons_def : 'a -> 'a lam  *)
fun cons_def x = LAM1 x

(*  o_def : 'a lam * 'a lam -> 'a lam  *)
fun o_def (f, g) = LAM2 (f, g)

(*  reverse_def : 'a list -> 'a list  *)
(*         walk : 'a list -> 'a lam   *)
fun reverse_def xs
    = let fun walk nil
              = id_def
            | walk (x :: xs)
              = o_def (walk xs, cons_def x)
      in apply (walk xs, nil)
      end
```

The auxiliary functions are only aliases for the data-type constructors. We also observe that `LAM1` and `LAM2` are always used in connection with each other. Therefore, they can be fused in a single constructor `LAM3` and so can their treatment in `apply_lam`. The result reads as follows.

```
datatype 'a lam_alt = LAM0
                    | LAM3 of 'a lam_alt * 'a

(*  apply_lam_alt : 'a lam_alt * 'a list -> 'a list  *)
fun apply_lam_alt (LAM0, ys)
    = ys
  | apply_lam_alt (LAM3 (f, x), ys)
    = apply_lam_alt (f, x :: ys)

(*  reverse_def_alt : 'a list -> 'a list     *)
(*            walk : 'a list -> 'a lam_alt  *)
fun reverse_def_alt xs
    = let fun walk nil
              = LAM0
            | walk (x :: xs)
              = LAM3 (walk xs, x)
      in apply_lam_alt (walk xs, nil)
      end
```

As in Section 2.1, we can see that `reverse_def_alt` embeds its input list into the data type `lam_alt`, homomorphically. The associated apply function could therefore be made to work directly on the input list. We also recognize `apply_lam_alt` as an uncurried version of the fast reverse function with an accumulator.

Hughes also uses his representation of intermediate lists to define a 'fields' function that extracts words from strings. His representation gives rise to an efficient implementation of the fields function. And indeed, as for reverse above, defunctionalizing this implementation gives the fast implementation that accumulates words in reverse order and reverses them using a fast reverse function once the whole word has been found. Defunctionalization thus confirms the effectiveness of Hughes's representation.

### 2.3 Defunctionalizing Church-encoded non-recursive data structures

Church-encoding a value amounts to representing it by a $\lambda$-term in such a way that operations on this value are carried out by applying the representation to specific $\lambda$-terms [3, 9, 24, 33].

A data structure is a sum in a domain. (When the data structure is inductive, the domain is recursive.) A sum is defined by its corresponding injection functions and a case dispatch [56, page 133]. Church-encoding a data structure consists in (1) combining injection functions and case dispatch into $\lambda$-terms and (2) operating by function application.

In the rest of this section, for simplicity, we only consider Church-encoded data structures that are uncurried. This way, we can defunctionalize them as a whole.

For example, monotyped Church pairs and their selectors are defined as follows.

```
(*  Church_pair : 'a * 'a -> ('a * 'a -> 'a) -> 'a  *)
fun Church_pair (x1, x2)
    = fn s : 'a * 'a -> 'a => s (x1, x2)

(*  Church_fst : (('a * 'a -> 'a) -> 'b) -> 'b  *)
fun Church_fst p
    = p (fn (x1, x2) => x1)

(*  Church_snd : (('a * 'a -> 'a) -> 'b) -> 'b  *)
fun Church_snd p
    = p (fn (x1, x2) => x2)
```

A pair is represented as a $\lambda$-term expecting one argument. This argument is a selector corresponding to the first or the second projection.

In general, each of the injection functions defining a data structure has the following form.

$$\mathrm{inj}_i = \lambda(x_1, ..., x_n).\underline{\lambda(s_1, ..., s_m).s_i\ (x_1, ..., x_n)}$$

So what happens if one defunctionalizes a Church-encoded data structure, i.e., the result of the injection functions? Each injection function gives rise to a data-type constructor whose arguments correspond to the free variables in the term underlined just above. These free variables are precisely the parameters

of the injection functions, which are themselves the parameters of the original constructors that were Church encoded.

Therefore defunctionalizing Church-encoded data structures (i.e., the result of their injection functions) gives rise to the same data structures, prior to Church encoding. These data structures are accessed through the auxiliary apply functions introduced by defunctionalization.

For example, monotyped Church pairs and their selectors are defunctionalized as follows.

- The selectors are closed terms and therefore the corresponding constructors are parameterless. By definition, a selector is passed a tuple of arguments and returns one of them.

```
datatype sel = FST
             | SND

(*  apply_sel : sel * ('a * 'a) -> 'a  *)
fun apply_sel (FST, (x1, x2))
    = x1
  | apply_sel (SND, (x1, x2))
    = x2
```

- There is one injection function for pairs, and therefore it gives rise to a data type with one constructor for the values of the two free variables of the result of `Church_pair`. The corresponding apply function performs a selection. (N.B: `apply_pair` calls `apply_sel`, reflecting the curried type of `Church_pair`.)

```
datatype 'a pair = PAIR of 'a * 'a

(*  apply_pair : 'a pair * sel -> 'a  *)
fun apply_pair (PAIR (x1, x2), s)
    = apply_sel (s, (x1, x2))
```

- Finally, constructing a pair amounts to constructing a pair, à la Tarski one could say [21], and selecting a component of a pair is achieved by calling `apply_pair`, which in turns calls `apply_sel`.

```
(*  Church_pair_def : 'a * 'a -> 'a pair  *)
fun Church_pair_def (x1, x2)
    = PAIR (x1, x2)

(*  Church_fst_def : 'a pair -> 'a  *)
fun Church_fst_def p
    = apply_pair (p, FST)

(*  Church_snd_def : 'a pair -> 'a  *)
fun Church_snd_def p
    = apply_pair (p, SND)
```

An optimizing compiler would inline both apply functions. The resulting selectors, together with the defunctionalized pair constructor, would then coincide with the original definition of pairs, prior to Church encoding.

### 2.4 Defunctionalizing Church-encoded recursive data structures

Let us briefly consider Church-encoded binary trees. Two injection functions occur: one for the leaves, and one for the nodes. A Church-encoded tree is a $\lambda$-term expecting two arguments. These arguments are the selectors corresponding to whether the tree is a leaf or whether it is a node.

```
fun Church_leaf x
    = fn (s1, s2) => s1 x

fun Church_node (t1, t2)
    = fn (s1, s2) => s2 (t1 (s1, s2), t2 (s1, s2))
```

Due to the inductive nature of binary trees, `Church-node` propagates the selectors to the subtrees.

In general, each of the injection functions defining a data structure has the following form.

$$\mathrm{inj}_i = \lambda(x_1, ..., x_n).\underline{\lambda(s_1, ..., s_m).s_i\ (x_1, ..., x_j\ (s_1, ..., s_m), ..., x_n)}$$

where $x_j\ (s_1, ..., s_m)$ occurs for each $x_j$ that is in the data type.

So what happens if one defunctionalizes a Church-encoded recursive data structure, i.e., the result of the injection functions? Again, each injection function gives rise to a data-type constructor whose arguments correspond to the free variables in the term underlined just above. These free variables are precisely the parameters of the injection functions, which are themselves the parameters of the original constructors that were Church encoded.

Therefore defunctionalizing Church-encoded recursive data structures (i.e., the result of their injection functions) also gives rise to the same data structures, prior to Church encoding. These data structures are accessed through the auxiliary apply functions introduced by defunctionalization.

Let us get back to Church-encoded binary trees. Since defunctionalization is a whole-program transformation, we consider a whole program. Let us consider the function computing the depth of a Church-encoded binary tree. This function passes two selectors to its argument. The first is the constant function returning 0, and accounting for the depth of a leaf. The second is a function that will be applied to the depth of the subtrees of each node, and computes the depth of the node by taking the max of the depths of the two subtrees and adding one.

```
fun Church_depth t
    = t (fn x => 0,
         fn (d1, d2) => 1 + Int.max (d1, d2))
```

This whole program is defunctionalized as follows.

- The selectors give rise to two constructors, `SEL_LEAF` and `SEL_NODE`, and the corresponding two apply functions, `apply_sel_leaf` and `apply_sel_node`.

```
datatype sel_leaf = SEL_LEAF

fun apply_sel_leaf (SEL_LEAF, x)
    = 0

datatype sel_node = SEL_NODE

fun apply_sel_node (SEL_NODE, (d1, d2))
    = Int.max (d1, d2) + 1
```

- As for the injection functions, as noted above, they give rise to two constructors, `LEAF` and `NODE`, and the corresponding apply function.

```
datatype 'a tree = LEAF of 'a
                 | NODE of 'a tree * 'a tree

fun Church_leaf_def x
    = LEAF x

fun Church_node_def (t1, t2)
    = NODE (t1, t2)
```

- Finally, the defunctionalized main function applies its argument to the two selectors.

```
(*   depth_def : 'a tree -> int                          *)
(*  apply_tree : 'a tree * (sel_leaf * sel_node) -> int  *)
fun depth_def t
    = apply_tree (t, (SEL_LEAF, SEL_NODE))
and apply_tree (LEAF x, (sel_leaf, sel_node))
    = apply_sel_leaf (sel_leaf, x)
  | apply_tree (NODE (t1, t2), (sel_leaf, sel_node))
    = apply_sel_node (sel_node, (apply_tree (t1, (sel_leaf, sel_node)),
                                 apply_tree (t2, (sel_leaf, sel_node))))
```

Again, an optimizing compiler would inline both apply functions and `SEL_LEAF` and `SEL_NODE` would then disappear. The result would thus coincide with the original definition of binary trees, prior to Church encoding.

### 2.5 Church-encoding the result of defunctionalization

As can be easily verified with the Church pairs and the Church trees above, Church-encoding the result of defunctionalizing a Church-encoded data structure gives back this Church-encoded data structure: the apply functions revert

to simple applications, the main data-structure constructors become injection functions, and the auxiliary data-structure constructors become selectors.

In practice, however, one often inlines selectors during Church encoding if they only occur once—which Shivers refers to as "super-beta" [46]. Doing so yields an actual inverse to defunctionalization, as illustrated in Sections 4.1.3 and 4.2.3. This "refunctionalization" is also used, e.g., in Danvy, Grobauer, and Rhiger's work on goal-directed evaluation [13].

In Church-encoded data structures, selectors have the flavor of a continuation. In Section 3, we consider how to defunctionalize continuations.

### 2.6 Summary and conclusion

We have considered a variety of typical higher-order programs and have defunctionalized them. The resulting programs provide a first-order view that reveals the effect of higher-orderness in functional programming. For example, returning a function of type $\alpha \to \beta$ often naturally provides a way to write this function with an $\alpha$-typed accumulator. For another example, defunctionalizing uncurried Church-encoded data structures leads one back to these data structures prior Church encoding, illustrating that Church encoding and defunctionalization transform data flow into control flow and vice-versa.

## 3 Defunctionalization of CPS-Transformed First-Order Programs

As functional representations of control, continuations provide a natural target for defunctionalization. In this section, we investigate the process of transforming direct-style programs into continuation-passing style (CPS) programs [12, 49] and defunctionalizing their continuation. We then compare this process with Wand's continuation-based program-transformation strategies [54].

### 3.1 String parsing

We consider a recognizer for the language $\{0^n1^n \mid n \in \mathbb{N}\}$. We write it as a function of type `int list -> bool`. The input is a list of integers, and the recognizer checks whether it is the concatenation of a list of $n$ 0's and of a list of $n$ 1's, for some $n$.

We start with a recursive-descent parser traversing the input list.

```
(*  rec0 : int list -> bool      *)
(*  walk : int list -> int list  *)
fun rec0 xs
    = let exception NOT
          fun walk (0 :: xs')
              = (case walk xs'
                   of 1 :: xs''
                      => xs''
                    | _
                      => raise NOT)
            | walk xs
              = xs
      in (walk xs = nil) handle NOT => false
      end
```

The auxiliary function `walk` traverses the input list. Every time it encounters `0`, it calls itself recursively. When it meets something else than `0`, it returns the rest of the list, and expects to find `1` at every return. In case of mismatch (i.e., a list element other than `1` for returns, or a list that is too short or too long), an exception is raised.

Let us write `walk` in continuation-passing style [12, 49].

```
(*  rec1 : int list -> bool                        *)
(*  walk : int list * (int list -> bool) -> bool  *)
fun rec1 xs
    = let fun walk (0 :: xs', k)
              = walk (xs', fn (1 :: xs'')
                              => k xs''
                            | _
                              => false)
            | walk (xs, k)
              = k xs
      in walk (xs, fn xs' => xs' = nil)
      end
```

The auxiliary function `walk` traverses the input list tail-recursively (and thus does not need any exception). If it meets something else than `0`, it sends the current list to the current continuation. If it encounters `0`, it iterates down the list with a new continuation. If the new continuation is sent a list starting with `1`, it sends the tail of that list to the current continuation; otherwise, it returns `false`. The initial continuation tests whether it is sent the empty list.

Let us defunctionalize `rec1`. Two function abstractions occur: one for the initial continuation and one for intermediate continuations.

```
datatype cont = CONT0
              | CONT1 of cont
```

```
(*  apply2 : (cont * int list) -> bool  *)
fun apply2 (CONT0, xs')
    = xs' = nil
  | apply2 (CONT1 k, 1 :: xs'')
    = apply2 (k, xs'')
  | apply2 (CONT1 k, _)
    = false

(*  rec2 : int list -> bool         *)
(*  walk : int list * cont -> bool  *)
fun rec2 xs
    = let fun walk (0 :: xs', k)
              = walk (xs', CONT1 k)
            | walk (xs, k)
              = apply2 (k, xs)
      in walk (xs, CONT0)
      end
```

We identify the result as implementing a push-down automaton [26]. This automaton has two states and one element in the stack alphabet. The two states are represented by the two functions `walk` and `apply2`. The stack is implemented by the data type `cont`. The transitions are the tail-recursive calls. This automaton accepts an input if processing this input ends with an empty stack.

We also observe that `cont` implements Peano numbers. Let us replace them with ML integers.

```
(*  apply3 : (int * int list) -> bool  *)
fun apply3 (0, xs')
    = xs' = nil
  | apply3 (k, 1 :: xs'')
    = apply3 (k-1, xs'')
  | apply3 (k, _)
    = false

(*  rec3 : int list -> bool        *)
(*  walk : int list * int -> bool  *)
fun rec3 xs
    = let fun walk (0 :: xs', k)
              = walk (xs', k+1)
            | walk (xs, k)
              = apply3 (k, xs)
      in walk (xs, 0)
      end
```

The result is the usual iterative two-state recognizer with a counter.

In summary, we started from a first-order recursive version and we CPS-transformed it, making it higher-order and thus defunctionalizable. We identified that the defunctionalized program implements a push-down automaton. Noticing that the defunctionalized continuation implements Peano arithmetic, we changed its representation to built-in integers and we identified that the result is the usual iterative two-state recognizer with a counter.

### 3.2 Continuation-based program transformation strategies

Wand's classical work on continuation-based program transformation [54] suggests one (1) to CPS-transform a program; (2) to design a data-structure representing the continuation; and (3) to use this representation to improve the initial program. We observe that in each of the examples mentioned in Wand's article, defunctionalization answers the challenge of finding a data structure representing the continuation—which is significant because finding such "data-structure continuations" was one of the motivations of the work. Yet defunctionalization is not considered in the textbooks and articles that refer to Wand's article and that we are aware of, which includes those found in the Research Index at `http://citeseer.nj.nec.com/`.

Wand's work is seminal in that it shows how detouring via CPS yields iterative programs with accumulators. In addition, Reynolds's work shows how defunctionalizing the continuation of CPS-transformed programs gives rise to traditional, first-order accumulators.

We also observe that defunctionalized continuations account for the call/return patterns of recursively defined functions. Therefore, as pointed out by Dijkstra in the late 1950's [16], they evolve in a stack-like fashion. A corollary of this remark is that before defunctionalization, continuations are also used LIFO when they result from the CPS transformation of a program that does not use control operators [10, 11, 12, 14, 15, 39, 40, 49].

### 3.3 Summary and conclusion

Defunctionalizing a CPS-transformed first-order program provides a systematic way to construct an iterative version of this program that uses a push-down accumulator. One can then freely change the representation of this accumulator.

## 4 Two Syntactic Theories in the Light of Defunctionalization

In this section, we present interpreters for two syntactic theories [19, 57]. One is for simple arithmetic expressions and the other for the call-by-value $\lambda$-calculus. We observe that both of these interpreters correspond to the output of defunctionalization. We then present the corresponding higher-order interpreters, which are in continuation-passing style. In each interpreter, the continuation represents the evaluation context of the corresponding syntactic theory.

### 4.1 Arithmetic expressions

We consider a simplified language of arithmetic expressions. An arithmetic expression is either a value (a literal) or a computation. A computation is either an addition or a conditional expression testing whether its first argument is zero.

$$e \quad ::= \quad n \mid e + e \mid \text{IFZ } e\ e\ e$$

### 4.1.1 A syntactic theory for arithmetic expressions

A syntactic theory provides a reduction relation on expressions by defining values, evaluation contexts, and redexes [19].

The values are literals, and the evaluation contexts are defined as follows.

$$\begin{array}{rcl} v & ::= & n \\ E & ::= & [\,] \mid E[[\,] + e] \mid E[v + [\,]] \mid E[\textsc{ifz}\ [\,]\ e\ e] \end{array}$$

Plugging an expression $e$ into a context $E$, written $E[e]$, is defined as follows.

$$\begin{array}{rcl} ([\,])[e] & = & e \\ (E[[\,] + e'])[e] & = & E[e + e'] \\ (E[v + [\,]])[e] & = & E[v + e] \\ (E[\textsc{ifz}\ [\,]\ e_1\ e_2])[e] & = & E[\textsc{ifz}\ e\ e_1\ e_2] \end{array}$$

The reduction relation is then defined by the following rules, where the expressions plugged into the context on the left-hand side are called *redexes*.

$$\begin{array}{rcll} E[n_1 + n_2] & \rightarrow & E[n_3] & \text{where } n_3 \text{ is the sum of } n_1 \text{ and } n_2 \\ E[\textsc{ifz}\ 0\ e_1\ e_2] & \rightarrow & E[e_1] & \\ E[\textsc{ifz}\ n\ e_1\ e_2] & \rightarrow & E[e_2] & \text{if } n \neq 0 \end{array}$$

These definitions satisfy a "unique decomposition" lemma [19, 57]: any expression, $e$, that is not a value can be uniquely decomposed into an evaluation context, $E$, and a redex, $r$, such that $e = E[r]$.

### 4.1.2 Implementation

Arithmetic expressions are defined with the following data type.

```
datatype aexp = VAL of int    (* trivial terms *)
              | COMP of comp  (* serious terms *)
     and comp = ADD of aexp * aexp
              | IFZ of aexp * aexp * aexp
```

In `aexp`, we distinguish between trivial terms, i.e., values (literals) and serious terms, i.e., computations (additions and conditional expressions), as traditional.

Evaluation contexts are defined with the following data type.

```
datatype evalcont = EMPTY
                  | ADD1 of evalcont * aexp
                  | ADD2 of evalcont * int
                  | IFZ0 of evalcont * aexp * aexp
```

The corresponding plugging function reads as follows.

```
(*  plug : evalcont * aexp -> aexp  *)
fun plug (EMPTY, ae)
    = ae
  | plug (ADD1 (ec, ae2), ae1)
    = plug (ec, COMP (ADD (ae1, ae2)))
  | plug (ADD2 (ec, i1), ae2)
    = plug (ec, COMP (ADD (VAL i1, ae2)))
  | plug (IFZ0 (ec, ae1, ae2), ae0)
    = plug (ec, COMP (IFZ (ae0, ae1, ae2)))
```

A computation undergoes a reduction step when (1) it is decomposed into a redex and its context, (2) the redex is contracted, and (3) the result is plugged into the context.

```
(*  reduce1 : comp * evalcont -> aexp  *)
fun reduce1 (ADD (VAL i1, VAL i2), ec)
    = plug (ec, VAL (i1+i2))
  | reduce1 (ADD (VAL i1, COMP s2), ec)
    = reduce1 (s2, ADD2 (ec, i1))
  | reduce1 (ADD (COMP s1, ae2), ec)
    = reduce1 (s1, ADD1 (ec, ae2))
  | reduce1 (IFZ (VAL 0, ae1, ae2), ec)
    = plug (ec, ae1)
  | reduce1 (IFZ (VAL i, ae1, ae2), ec)
    = plug (ec, ae2)
  | reduce1 (IFZ (COMP s0, ae1, ae2), ec)
    = reduce1 (s0, IFZ0 (ec, ae1, ae2))
```

Evaluation is specified by repeatedly performing a reduction until a value is obtained.

```
(*  eval : ae -> int  *)
fun eval (VAL i)
    = i
  | eval (COMP s)
    = eval (reduce1 (s, EMPTY))
```

### 4.1.3 Refunctionalization

We observe that the interpreter of Section 4.1.2 precisely corresponds to the output of defunctionalization: `plug` is the apply function of `ec`. The corresponding input to defunctionalization thus reads as follows.

```
(*  reduce1 : comp * (aexp -> int) -> int  *)
fun reduce1 (ADD (VAL i1, VAL i2), ec)
    = ec (VAL (i1+i2))
  | reduce1 (ADD (VAL i1, COMP s2), ec)
    = reduce1 (s2, fn ae2 => ec (COMP (ADD (VAL i1, ae2))))
  | reduce1 (ADD (COMP s1, ae2), ec)
    = reduce1 (s1, fn ae1 => ec (COMP (ADD (ae1, ae2))))
  | reduce1 (IFZ (VAL 0, ae1, ae2), ec)
    = ec ae1
  | reduce1 (IFZ (VAL i, ae1, ae2), ec)
    = ec ae2
  | reduce1 (IFZ (COMP s0, ae1, ae2), ec)
    = reduce1 (s0, fn ae0 => ec (COMP (IFZ (ae0, ae1, ae2))))

(*  eval : ae -> int  *)
fun eval (VAL i)
    = i
  | eval (COMP s)
    = eval (reduce1 (s, fn e => e))
```

We observe that `reduce1` is written in continuation-passing style. Its continuation therefore represents the evaluation context of the syntactic theory.

### 4.1.4 Back to direct style

In Section 4.1.3, since `reduce1` uses its continuation canonically, it can be mapped back to direct style [10, 14]. The result reads as follows.

```
(*  reduce1 : comp -> aexp  *)
fun reduce1 (ADD (VAL i1, VAL i2))
    = VAL (i1+i2)
  | reduce1 (ADD (VAL i1, COMP s2))
    = COMP (ADD (VAL i1, reduce1 s2))
  | reduce1 (ADD (COMP s1, ae2))
    = COMP (ADD (reduce1 s1, ae2))
  | reduce1 (IFZ (VAL 0, ae1, ae2))
    = ae1
  | reduce1 (IFZ (VAL i, ae1, ae2))
    = ae2
  | reduce1 (IFZ (COMP s0, ae1, ae2))
    = COMP (IFZ (reduce1 s0, ae1, ae2))

(*  eval : ae -> int  *)
fun eval (VAL i)
    = i
  | eval (COMP s)
    = eval (reduce1 s)
```

The resulting direct-style interpreter represents evaluation contexts implicitly.

### 4.2 The call-by-value $\lambda$-calculus

We now present a syntactic theory for the call-by-value $\lambda$-calculus, and its implementation.

#### 4.2.1 A syntactic theory for the call-by-value $\lambda$-calculus

We consider the call-by-value $\lambda$-calculus. A $\lambda$-expression is either a value (an identifier or a $\lambda$-abstraction) or a computation. A computation is always an application.

$$\begin{array}{rcl} e & ::= & x \mid \lambda x.e \mid e\ e \\ \\ v & ::= & x \mid \lambda x.e \\ E & ::= & [\,] \mid E[[\,]\ e] \mid E[v\ [\,]] \end{array}$$

Plugging an expression $e$ into a context $E$ is defined as follows.

$$\begin{array}{rcl} ([\,])[e] & = & e \\ (E[[\,]\ e'])[e] & = & E[e\ e'] \\ (E[v\ [\,]])[e] & = & E[v\ e] \end{array}$$

Any application of values, $v_1\ v_2$, is called a *redex*, although not all of them actually reduce to anything. Only $\beta$-redexes do; the others are stuck.

The reduction relation is then defined by the following rule.

$$E[(\lambda x.e)\ v] \quad \rightarrow \quad E[e[v/x]]$$

These definitions also satisfy a "unique decomposition" lemma, allowing us to implement the one-step reduction relation as a function.

#### 4.2.2 Implementation

The $\lambda$-calculus is defined with the following data type, where we distinguish between values and computations.

```
datatype term = VAL of value  (* trivial terms *)
              | COMP of comp  (* serious terms *)
    and value = VAR of string
              | LAM of string * term
     and comp = APP of term * term
```

In `term`, we distinguish between trivial terms, i.e., values (variables and $\lambda$-abstractions) and serious terms, i.e., computations (applications), as traditional.

We also need a substitution function of the following type:

```
val subst : term * value * string -> term
```

Evaluation contexts are defined with the following data type.

```
datatype evalcont = EMPTY
                  | APP1 of evalcont * term
                  | APP2 of evalcont * value
```

The corresponding plugging function reads as follows.

```
(*  plug : evalcont * term -> term  *)
fun plug (EMPTY, e)
    = e
  | plug (APP1 (ec, e'), e)
    = plug (ec, COMP (APP (e, e')))
  | plug (APP2 (x, t), e)
    = plug (x, COMP (APP (VAL t, e)))
```

A computation undergoes a reduction step when (1) it is decomposed into a redex and its context, (2) the redex is contracted, if possible, and (3) the result is plugged into the context. We therefore use an option type to account for the possibility in (2).

```
datatype 'a option = NONE
                   | SOME of 'a

(*  reduce1 : comp * evalcont -> term option  *)
fun reduce1 (APP (VAL (LAM (x, e)), VAL t), ec)
    = SOME (plug (ec, subst (e, t, x)))
  | reduce1 (s as APP (VAL (VAR x), VAL t), ec)
    = NONE
  | reduce1 (APP (VAL t, COMP s), ec)
    = reduce1 (s, APP2 (ec, t))
  | reduce1 (APP (COMP s, e), ec)
    = reduce1 (s, APP1 (ec, e))
```

Again, evaluation is specified by repeatedly performing a reduction until a value, if any, is obtained. We use an option type to account for stuck terms.

```
(*  eval : term -> value option  *)
fun eval (VAL t)
    = SOME t
  | eval (COMP s)
    = (case reduce1 (s, EMPTY)
         of SOME e => eval e
          | NONE => NONE)
```

There are two ways of not obtaining a value: evaluating a stuck term and evaluating a diverging term (`eval` then does not terminate).

### 4.2.3 Refunctionalization

Again, we observe that the program above precisely corresponds to the output of defunctionalization with `plug` as the apply function of `ec`. The corresponding input to defunctionalization thus reads as follows.

```
(*  reduce1 : comp * (term -> term option) -> term option  *)
fun reduce1 (APP (VAL (LAM (x, e)), VAL t), ec)
    = ec (subst (e, t, x))
  | reduce1 (s as APP (VAL (VAR x), VAL t), ec)
    = NONE
  | reduce1 (APP (VAL t, COMP s), ec)
    = reduce1 (s, fn e => ec (COMP (APP (VAL t, e))))
  | reduce1 (APP (COMP s, e), ec)
    = reduce1 (s, fn e' => ec (COMP (APP (e', e))))

(*  eval : term -> value option  *)
fun eval (VAL t)
    = SOME t
  | eval (COMP s)
    = (case reduce1 (s, fn e => SOME e)
         of SOME e => eval e
          | NONE => NONE)
```

We observe that `reduce1` is again written in continuation-passing style. Its continuation therefore represents the evaluation context of the syntactic theory.

### 4.2.4 Back to direct style

Since `reduce1` uses its continuation canonically, it can be mapped back to direct style, using a local exception for stuck terms. The direct-style version of `reduce1` reads as follows.

```
exception STUCK

(*  reduce1 : comp -> term  *)
fun reduce1 (APP (VAL (LAM (x, e)), VAL t))
    = subst (e, t, x)
  | reduce1 (s as APP (VAL (VAR x), VAL t))
    = raise STUCK
  | reduce1 (APP (VAL t, COMP s))
    = COMP (APP (VAL t, reduce1 s))
  | reduce1 (APP (COMP s, e))
    = COMP (APP (reduce1 s, e))

(*  eval : term -> value option  *)
fun eval (VAL t)
    = SOME t
  | eval (COMP s)
    = eval (reduce1 s)
      handle STUCK => NONE
```

The result is a direct-style interpreter with an implicit representation of evaluation contexts, just as for arithmetic expressions.

### 4.3 Summary and conclusion

We have considered two interpreters for syntactic theories, and we have observed that the evaluation contexts and their plugging function are the defunctionalized counterparts of continuations. This observation has led us to implement both interpreters in direct style. (In that sense, Sections 3 and 4 are symmetric, since Section 3 starts with a direct-style program and ends with a defunctionalized CPS program.) This observation also suggests how to automatically obtain a grammar of evaluation contexts out of the BNF of a language: defunctionalize the CPS counterpart of a recursive descent over this BNF (e.g., a one-step reduction function). The data type representing the continuation is isomorphic to the desired grammar, and the apply function is the corresponding plug function.

In general, evaluation contexts are specified so that a context can easily be extended to the left or to the right with an elementary context. For the call-by-value $\lambda$-calculus, the two specifications read as follows:

$$\begin{array}{lcl} E & ::= & [\,] \mid E\ e \mid v\ E \\ E & ::= & [\,] \mid E[[\,]\ e] \mid E[v\ [\,]] \end{array}$$

Or again, written using an explicit composition operator $\circ$ (satisfying $(E_1 \circ E_2)[e] = E_1[E_2[e]]$):

$$\begin{array}{lcl} E & ::= & [\,] \mid [[\,]\ e] \circ E \mid [v\ [\,]] \circ E \\ E & ::= & [\,] \mid E \circ [[\,]\ e] \mid E \circ [v\ [\,]] \end{array}$$

Since all evaluation contexts can be constructed by composing elementary contexts and since composition is associative, the two specifications define effectively the same contexts. Only their representations differ. With the first representation, an evaluation context is isomorphic to a list of elementary contexts. With the second representation, the same evaluation context is isomorphic to the reversed list.

Therefore, with the first representation, plugging an expression in a context is recursively carried out by a right fold over the context,[1] and with the second representation, plugging an expression in a context is iteratively carried out by a left fold over the context.[2] The latter is implicitly what we have done in Sections 4.1.2 and 4.2.2. In that light, let us reconsider the observation above on how to automatically obtain a grammar of evaluation contexts out of the BNF of a language. We can see that the resulting grammar is of the second kind and that the associated plug/apply function is iterative, which is characteristic of a left fold.

Let us get back to the first representation for arithmetic expressions and for the call-by-value $\lambda$-calculus. The unstated BNF for arithmetic expressions reads as follows.

$$E ::= [\,] \mid E + e \mid v + E \mid \textsc{ifz}\ E\ e\ e$$

[1] Reminder: $\mathit{foldr}\ f\ b\ [x_1, x_2, ..., x_n] = f(x_1, f(x_2, ..., f(x_n, b)...))$.
[2] Reminder: $\mathit{foldl}\ f\ b\ [x_n, ..., x_2, x_1] = f(x_1, f(x_2, ..., f(x_n, b)...))$.

In the two interpreters, decomposition produces an evaluation context and plugging consumes it. We observe that deforesting this combination of decomposition and plugging yields the same direct-style interpreters as the ones obtained in Sections 4.1.4 and 4.2.4.

One may also wonder how the various representations of evaluation contexts in a syntactic theory—i.e., as data types and as functions—influence reasoning about programs. In the next section, we compare two correctness proofs of a program, before and after defunctionalization.

# 5 A Comparison between Correctness Proofs before and after Defunctionalization: Matching Regular Expressions

We consider a traditional continuation-based matcher for regular expressions [26], we defunctionalize it, and we compare and contrast its correctness proof before and after defunctionalization. To this end, Section 5.1 briefly reviews regular expressions and the languages they represent; Section 5.2 presents the continuation-based matcher, which is higher-order, and its defunctionalized counterpart; and Section 5.3 compares and contrasts their correctness proofs.

## 5.1 Regular expressions

The grammar for regular expressions, $r$, over the alphabet $\Sigma$ and the corresponding language, $\mathcal{L}(r)$, are as follows.

$$
\begin{array}{rcll}
r & ::= & \mathbf{0} & \mathcal{L}(\mathbf{0}) = \emptyset \\
 & | & \mathbf{1} & \mathcal{L}(\mathbf{1}) = \{\epsilon\} \\
 & | & \mathbf{c} & \mathcal{L}(\mathbf{c}) = \{\mathbf{c}\} \quad \text{where } \mathbf{c} \in \Sigma \\
 & | & r\ r & \mathcal{L}(r_1\ r_2) = \mathcal{L}(r_1)\mathcal{L}(r_2) \\
 & | & r{+}r & \mathcal{L}(r_1{+}r_2) = \mathcal{L}(r_1) \cup \mathcal{L}(r_2) \\
 & | & r^* & \mathcal{L}(r^*) = \mathcal{L}(r)^* = \bigcup_{i \in \omega} (\mathcal{L}(r))^i
\end{array}
$$

We represent strings as lists of ML characters, and regular expressions as elements of the following ML datatype.

```
datatype regexp = ZERO
                | ONE
                | CHAR of char
                | CAT of regexp * regexp
                | SUM of regexp * regexp
                | STAR of regexp
```

We define the corresponding notion of "the language of a regular expression" as follows.

$$
\begin{array}{rcl@{\qquad}rcl}
\mathcal{L}(\texttt{ZERO}) & = & \{\} & \mathcal{L}(\texttt{CAT(r1, r2)}) & = & \mathcal{L}(\texttt{r1})\mathcal{L}(\texttt{r2}) \\
\mathcal{L}(\texttt{ONE}) & = & \{\texttt{nil}\} & \mathcal{L}(\texttt{SUM(r1, r2)}) & = & \mathcal{L}(\texttt{r1}) \cup \mathcal{L}(\texttt{r2}) \\
\mathcal{L}(\texttt{CHAR c}) & = & \{[\texttt{c}]\} & \mathcal{L}(\texttt{STAR r}) & = & \bigcup_{i \in \omega} (\mathcal{L}(\texttt{r}))^i
\end{array}
$$

```
(*  accept      : regexp * char list * (char list -> bool) -> bool  *)
(*  accept_star : regexp * char list * (char list -> bool) -> bool  *)
fun accept (r, s, k)
    = (case r of ZERO
                  => false
                | ONE
                  => k s
                | CHAR c
                  => (case s of (c'::s')
                                  => c = c' andalso k s'
                                | nil
                                  => false)
                | CAT (r1, r2)
                  => accept (r1, s, fn s' => accept (r2, s', k))
                | SUM (r1, r2)
                  => accept (r1, s, k) orelse accept (r2, s, k)
                | STAR r'
                  => accept_star (r', s, k))
and accept_star (r, s, k)
    = k s
      orelse accept (r, s, fn s' => not (s = s')
                                    andalso accept_star (r, s', k))

(*  match : regexp * char list -> bool  *)
fun match (r, s)
    = accept (r, s, fn s' => s' = nil)
```

Figure 1: Higher-order, continuation-based matcher for regular expressions

The concatenation of languages is defined as $L_1L_2 = \{\texttt{x@y} \mid \texttt{x} \in L_1 \wedge \texttt{y} \in L_2\}$, where we use the append function (noted `@` as in ML) to concatenate strings.

### 5.2 The two matchers

Our reference matcher for regular expressions is higher-order (Figure 1). We then present its defunctionalized counterpart (Figure 2).

#### 5.2.1 The higher-order matcher

Figure 1 displays our reference matcher, which is compositional and continuation-based. Compositional: all recursive calls to `accept` operate on a proper subpart of the regular-expression under consideration. And continuation-based: the control flow of the matcher is driven by continuations.

The main function is `match`. It is given a regular expression and a list of characters, and calls `accept` with the regular expression, the list, and an initial continuation expecting a list of characters and testing whether this list is empty.

The `accept` function recursively descends its input regular expression, threading the list of characters.

The `accept_star` function is a lambda-lifted version of a recursive continuation defined locally in the `STAR` branch. (The situation is exactly the same as in a compositional interpreter for an imperative language with while loops, where one writes an auxiliary recursive function to interpret loops.) This recursive continuation checks that matching has progressed in the string.

Recently, Harper has published a similar matcher to illustrate "proof-directed debugging" [25]. Playfully, he considered a non-compositional matcher that does not check progress when matching a Kleene star. His article shows (1) how one stumbles on the non-compositional part when attempting a proof by structural induction; and (2) how one realizes that the matcher diverges for pathological regular expressions such as `STAR ONE`. Harper then (1) makes his matcher compositional and (2) normalizes regular expressions to exclude pathological regular expressions. Instead, we start from a compositional matcher and we include a progress test in `accept_star`, which lets us handle pathological regular expressions.

### 5.2.2 The first-order matcher

Defunctionalizing the matcher of Figure 1 yields a data type representing the continuations and its associated apply function.

The data type represents a stack of regular expressions (possibly with a side condition for the test in Kleene stars). The apply function merely pops the top element off this stack and tries to match it against the rest of the string. We thus name the data type "`regexp_stack`" and the apply function "`pop_and_accept`". We also give a meaningful name to the datatype constructors. Figure 2 displays the result.

## 5.3 The two correctness proofs

We give a correctness proof of both the higher-order version and the first-order version, and we investigate whether each proof can be converted to a proof for the other version.

The correctness criterion we choose is simply that for all regular expressions `r` and strings `s` (represented by a list of characters),

$$\begin{cases} \text{evaluating } \texttt{match (r, s)} \text{ terminates, and} \\ \texttt{match (r, s)} \rightsquigarrow \texttt{true} \;\Leftrightarrow\; \texttt{s} \in \mathcal{L}(\texttt{r}) \end{cases}$$

When writing `match (r, s)` $\rightsquigarrow$ `true`, we mean that evaluating `match (r, s)` terminates and yields the result `true`. We also reason about ML programs equationally, writing $e \equiv e'$ if $e$ and $e'$ are defined to be equal, e.g., by a function definition. If $e \equiv e'$ then $e$ and $e'$ evaluate to the same value, if any, so $e \rightsquigarrow v \Leftrightarrow e' \rightsquigarrow v$. We write $e = e'$ on ML expressions only if they represent the same *value*.

```
datatype regexp_stack = EMPTY
                      | ACCEPT of regexp * regexp_stack
                      | ACCEPT_STAR of char list * regexp * regexp_stack

(*  accept_def      : regexp * char list * regexp_stack -> bool  *)
(*  accept_star_def : regexp * char list * regexp_stack -> bool  *)
(*  pop_and_accept  : regexp_stack * char list                   *)
fun accept_def (r, s, k)
    = (case r of ZERO
                 => false
               | ONE
                 => pop_and_accept (k, s)
               | CHAR c
                 => (case s of (c'::s')
                               => c = c'
                                  andalso pop_and_accept (k, s')
                             | nil
                               => false)
               | CAT (r1, r2)
                 => accept_def (r1, s, ACCEPT (r2, k))
               | SUM (r1, r2)
                 => accept_def (r1, s, k) orelse accept_def (r2, s, k)
               | STAR r'
                 => accept_star_def (r', s, k))
and accept_star_def (r, s, k)
    = pop_and_accept (k, s)
      orelse accept_def (r, s, ACCEPT_STAR (s, r, k))
and pop_and_accept (EMPTY, s')
    = s' = nil
  | pop_and_accept (ACCEPT (r2, k), s')
    = accept_def (r2, s', k)
  | pop_and_accept (ACCEPT_STAR (s, r, k), s')
    = not (s = s')
      andalso accept_star_def (r, s', k)

(*  match : regexp * char list -> bool  *)
fun match (r, s)
    = accept_def (r, s, EMPTY)
```

Figure 2: First-order, stack-based matcher for regular expressions

### 5.3.1 Correctness proof of the higher-order matcher

Since `match (r, s)` $\equiv$ `accept (r, s, fn s' => s' = nil)`, by definition, it is sufficient to prove that for `s` and `r` as above, and for any function from lists of characters to booleans terminating on all suffixes of `s`, denoted by `k`,

$$\begin{cases} \text{evaluating } \texttt{accept (r, s, k)} \text{ terminates, and} \\ \texttt{accept (r, s, k)} \rightsquigarrow \texttt{true} \;\Leftrightarrow\; \texttt{s} \in \mathcal{L}(\texttt{r})\mathcal{L}(\texttt{k}) \end{cases}$$

where we define the language of a "string-acceptor" `k` as the set $\{\texttt{s} \mid \texttt{k s} \rightsquigarrow \texttt{true}\}$.

The proof is by structural induction on the regular expression. In the case where `r` = `STAR r'`, a subproof shows that the following holds for any string.

$$\begin{cases} \text{evaluating } \texttt{accept_star (r', s, k)} \text{ terminates, and} \\ \texttt{accept_star (r', s, k)} \rightsquigarrow \texttt{true} \;\Leftrightarrow\; \texttt{s} \in \mathcal{L}(\texttt{r'})^*\mathcal{L}(\texttt{k}) \end{cases}$$

The subproof is by well-founded induction on the structure of the string (suffixes are smaller) for the "$\Leftarrow$" direction, and by mathematical induction on the natural number $n$ such that $\texttt{s} \in \mathcal{L}(\texttt{r'})^n\mathcal{L}(\texttt{k})$ for the "$\Rightarrow$" direction. Both subproofs use the outer induction hypothesis for `accept (r', s, k)`. (See Appendix A.)

We can transfer this proof to the defunctionalized version. Since `k s` translates to `pop_and_accept (k, s)`, we define $\mathcal{L}(\texttt{k})$ to read $\{\texttt{s} \mid$ `pop_and_accept (k, s)` $\rightsquigarrow \texttt{true}\}$. The proof then goes through in exactly the same format.

### 5.3.2 Correctness proof of the first-order matcher

Alternatively, if we were to prove the correctness of the first-order matcher directly, we would be less inclined to recognize the stack `k` as representing a function. Instead, we could easily end up proving the following three propositions by mutual induction.

$$\begin{aligned} P_1(\texttt{r},\texttt{s},\texttt{k}) &\stackrel{\text{def}}{=} \texttt{accept (r, s, k)} \rightsquigarrow \texttt{true} \;\Leftrightarrow\; \texttt{s} \in \mathcal{L}(\texttt{r})\mathcal{L}(\texttt{k}) \\ P_2(\texttt{k},\texttt{s}) &\stackrel{\text{def}}{=} \texttt{pop_and_accept (k, s)} \rightsquigarrow \texttt{true} \;\Leftrightarrow\; \texttt{s} \in \mathcal{L}(\texttt{k}) \\ P_3(\texttt{r},\texttt{s},\texttt{k}) &\stackrel{\text{def}}{=} \texttt{accept_star (r, s, k)} \rightsquigarrow \texttt{true} \;\Leftrightarrow\; \texttt{s} \in \mathcal{L}(\texttt{r})^*\mathcal{L}(\texttt{k}) \end{aligned}$$

where we define the language of a stack of regular expressions as follows.

$$\begin{aligned} \mathcal{L}(\texttt{EMPTY}) &= \{\texttt{nil}\} \\ \mathcal{L}(\texttt{ACCEPT (r, k)}) &= \mathcal{L}(\texttt{r})\mathcal{L}(\texttt{k}) \\ \mathcal{L}(\texttt{ACCEPT_STAR (s, r, k)}) &= (\mathcal{L}(\texttt{r})^*\mathcal{L}(\texttt{k})) \setminus \{\texttt{s}\} \end{aligned}$$

For brevity we ignore the termination part of the proof and assume that all the functions are total. We prove, by well-founded induction on the propositions themselves, that $P_1$, $P_2$, and $P_3$ hold for any choices of `s`, `r`, and `k`. The ordering is an intricate mapping into $\omega \times \omega$, ordered lexicographically so that the proof of a proposition only depends on "smaller" propositions. (See Appendix B.)

This proof is more convoluted than the higher-order one for two reasons:

1. it separates the language of the continuation from the function that matches it, so one has to check whether the function really matches the correct language; and

2. it combines the two nested inductions of the higher-order proof into one well-founded induction.

Still this proof reveals a property of the continuations in the higher-order version, namely that there are at most three different kinds of continuations in use, something that cannot be seen from the type of the continuation—a full function space.

We could thus define a subset of this function space inductively, so that it only contains the functions that can be generated by the three abstractions. The first-order proof could then be extended to the higher-order program by assuming everywhere that the continuation `k` lies in this subset and showing that newly generated continuations do too. In effect, the set of continuations is partitioned into disjoint subsets, just as the first-order datatype represents a sum, and then we can prove something about elements in each part.

### 5.4 Comparison between the two correctness proofs

The proof of correctness of the first-order matcher directly uses the fact that the inductively defined data type representing continuations is a sum. Given a value of this type, we can reason by inversion and do a proof by cases. We show that a proposition holds for each of the three possible summands, and we conclude that the proposition must hold for any value of that type. This is reminiscent of defunctionalization, where in an entire program, the functions occupying a function space are exactly those originating from the function abstractions in the program, which can be represented by a finite sum.

We can translate the proof of the first-order matcher directly into a proof of the higher-order matcher. The resulting proof uses global reasoning, namely that all the used functions arise from only finitely many function abstractions,[3] and furthermore that these function abstractions inductively define a subset of the function space. We change the induction hypothesis to assume that the functions are taken from that subset. When instantiating a function abstraction, we must show that the result belongs to the subset—which follows from the assumptions on the free variables—allowing us to use the induction hypothesis on this new function. When using a function, we can then reason by inversion and do a proof by cases to show a property of this function. Each case corresponds to an abstraction that can have created this function. We show that a proposition holds for each of the three possible function abstractions, and we conclude that the proposition must hold for the function.

The proof of correctness of the higher-order matcher uses local reasoning instead. We do not assume that the functions lie in the exact subset of the

[3]Such global reasoning is the one enabled by control-flow analysis [46].

function space that is generated by the function abstractions. Instead, we assume the weaker property that given a string, the function terminates on all suffixes of that string. This assumption is sufficient to complete the proof.

We can translate the proof of the higher-order matcher directly into a proof of the first-order matcher by replacing function abstractions with datatype constructors and applications with calls to the apply function. The resulting proof uses local reasoning, namely that all the defined functions satisfy a property. The property that the function `k` terminates on all suffixes of the current string is translated to the property that `pop_and_accept (k, s')` terminates if `s'` is a suffix of the current string. This assumption is then propagated to the calls to `pop_and_accept`.

Therefore, the two proofs differ where the inductive reasoning occurs:

- in the first-order case, the inductive reasoning about a data-type value is carried out at the use point; in contrast, no reasoning takes place where the data-type value is defined; and
- in the higher-order case, the inductive reasoning about a function value is carried out at the definition point; in contrast, no reasoning takes place where the function value is used.

This difference is emphasized by the translations between the proofs. It originates from the different handlings of recursion in the two matchers. In the first-order case, recursion is handled in the case dispatch where the values are used. In the higher-order case, recursion is handled in a function abstraction, which is only available to reason upon where the function is defined.

### 5.5 Summary and conclusion

We have considered a matcher for regular expressions, both in higher-order form and in first-order form, and we have compared them and their correctness proof. The difference between the function-based and the datatype-based representation of continuations is reminiscent of the concept of 'junk' in algebraic semantics [22]. One representation is a full function space where many elements do not correspond to an actual continuation, and the other representation only contains elements corresponding to actual continuations. This difference finds an echo in the correctness proofs of the two matchers, as analyzed in Section 5.4.

More generally, this section also illustrates that defunctionalizing a functional interpreter for a backtracking language that uses success continuations yields a recursive interpreter with one stack [13, 31, 41]. Similarly, defunctionalizing a functional interpreter that uses success and failure continuations yields an iterative interpreter with two stacks [13, 34].

## 6 Conclusions and Issues

Reynolds's defunctionalization technique connects the world of higher-order programs and the world of first-order programs. In this article, we have illustrated this connection by considering a variety of situations where defunctionalization proves fruitful in a declarative setting.

Higher-order functions provide a convenient support for specifying and for transforming programs. As we have shown, defunctionalization can lead to more concrete specifications, e.g., that use first-order accumulators. And as we have seen with Wand's continuation-based program-transformation strategies, defunctionalization can automate eureka steps to represent data-structure continuations.

Conversely, defunctionalization also increases one's awareness that some first-order programs naturally correspond to other, higher-order, programs. For example, we have seen that the evaluation contexts of a syntactic theory correspond to continuations. For another example, functional interpreters for backtracking languages are variously specified with one or two control stacks and with one or two continuations, but these specifications are not disconnected, since defunctionalizing the continuation-based interpreters yields the corresponding stack-based ones. On a related note, CPS-transforming an interpreter with one continuation is already known to automatically yield an interpreter with two continuations [12]. We are, however, not aware of a similar transformation for their stack-based counterparts.

We also have compared and contrasted the correctness proofs of a program, before and after defunctionalization. We have found that while the first-order and the higher-order programming methods suggest different proof methods, each of the proofs can be adapted to fit the other version of the program.

Finally, we have pointed out at the type-theoretical foundations of defunctionalization.

**Acknowledgments:** Andrzej Filinski and David Toman provided most timely comments on an earlier version of this article. This article has also benefited from Daniel Damian, Julia Lawall, Karoline Malmkjær, and Morten Rhiger's comments, from John Reynolds's encouraging words, and from the attention of the PPDP'01 anonymous reviewers.

This work is supported by the ESPRIT Working Group APPSEM (`http://www.md.chalmers.se/Cs/Research/Semantics/APPSEM/`).

# A Correctness Proof of the Higher-Order Matcher

To prove the correctness criterion given in Section 5.3.1, it suffices to prove the following property for all regular expressions, `r`, strings, `s`, and functions from strings (i.e., lists of characters) to booleans, `k`, that terminate on all suffixes of `s`:

> evaluating `accept(r,s,k)` terminates, and
> `accept(r,s,k)` $\leadsto$ `true` $\Leftrightarrow$ `s` $\in \mathcal{L}$(`r`)$\mathcal{L}$(`k`)

where we define the language of a string-acceptor, `k`, by $\mathcal{L}$(`k`) = $\{$`s` | `k(s)` $\leadsto$ `true`$\}$.

The proof is by structural induction on the regular expression `r`. In each case let `k` and `s` be given such that evaluating `k` terminates on all suffixes of `s`.

**case** `r` = `ZERO` : Evaluating `accept(r,s,k)` terminates, since it evaluates immediately to `false`.

The equivalence holds since both sides of the bi-implication are false, one because `accept(r,s,k)` $\leadsto$ `false` and the other because $\mathcal{L}$(`r`) = $\emptyset$.

**case** `r` = `ONE` : Evaluating `accept(r,s,k)` terminates, since `k` terminates on any suffix of `s`, including `s` itself.

Also, `accept(ONE,s,k)` $\equiv$ `k(s)` $\leadsto$ `true` is, by definition, equivalent to `s` $\in$ $\mathcal{L}$(`k`), and $\mathcal{L}$(`r`)$\mathcal{L}$(`k`) = $\mathcal{L}$(`k`).

**case** `r` = `CHAR c` : Evaluating `accept(r,s,k)` terminates since either it yields `false` immediately, or it calls `k` with a suffix of `s`.

If (and only if) `s` $\in \mathcal{L}$(`CHAR c`)$\mathcal{L}$(`k`) = $\{[c]\}\,\mathcal{L}(k)$ then there is a `s'` such that `s` = `c` :: `s'` and `s'` $\in \mathcal{L}$(`k`).

By definition, `s'` $\in \mathcal{L}$(`k`) is equivalent to `k(s')` $\leadsto$ `true`, and `s` = `c` :: `s'` and `k(s')` $\leadsto$ `true` are exactly the conditions under which `accept(r,s,k)` $\leadsto$ `true`.

**case** `r` = `SUM(r1,r2)` : Evaluating `accept(r,s,k)` terminates since `accept(r1,s,k)` and `accept(r2,s,k)` are both known to terminate from the induction hypothesis.

The condition `s` $\in \mathcal{L}$(`r`)$\mathcal{L}$(`k`) is equivalent to `s` $\in \mathcal{L}$(`r1`)$\mathcal{L}$(`k`)$\vee$`s` $\in \mathcal{L}$(`r2`)$\mathcal{L}$(`k`).

By induction hypothesis, `s` $\in \mathcal{L}$(`r1`)$\mathcal{L}$(`k`) $\Leftrightarrow$ `accept(r1,s,k)` $\leadsto$ `true` and `s` $\in \mathcal{L}$(`r2`)$\mathcal{L}$(`k`) $\Leftrightarrow$ `accept(r2,s,k)` $\leadsto$ `true`, and both applications of `accept` terminate. Together these conditions imply

```
accept(r1,s,k) orelse accept(r2,s,k) ⇝ true
```

which is equivalent to `accept(r,s,k)` $\leadsto$ `true`.

**case** `r` = `CAT(r1,r2)` : Evaluating `accept(r,s,k)` terminates, since when `s'` is a suffix of `s`, the induction hypothesis tells us that evaluating `accept(r2,s',k)` terminates (`k` terminates for all suffixes of `s'` since they are also suffixes of `s`). Therefore the function `fn s'=>accept(r2,s',k)` terminates on all

suffixes of `s`. In this case, the induction hypothesis tells us that evaluating `accept(r1,s,fn s=>...)` terminates.

Also, $\texttt{s} \in \mathcal{L}(\texttt{r})\mathcal{L}(\texttt{k})$ is equivalent to $\texttt{s} \in \mathcal{L}(\texttt{r1})\mathcal{L}(\texttt{r2})\mathcal{L}(\texttt{k})$ (language concatenation associates).

Let $k' =$ `fn s'=>accept(r2,s',k)`. Then, by induction hypothesis, $k'$ terminates for arguments that are suffixes of `s` since `k` does, and `k(s')` $\rightsquigarrow$ `true` iff `accept(r2,s',k)` $\rightsquigarrow$ `true`, which by induction hypothesis is exactly if `s'` $\in \mathcal{L}(\texttt{r2})\mathcal{L}(k)$ holds. Also by induction hypothesis, $\texttt{s} \in \mathcal{L}(\texttt{r1})\mathcal{L}(k') \rightsquigarrow \mathcal{L}(\texttt{r1})\mathcal{L}(\texttt{r2})\mathcal{L}(\texttt{k}) \Leftrightarrow$ `accept(r1,s,`$k'$`)` $\rightsquigarrow$ `true` which holds exactly when `accept(r,s,k)` $\rightsquigarrow$ `true` holds.

**case** `r = STAR(r1)` : In this case $\texttt{s} \in \mathcal{L}(\texttt{r})\mathcal{L}(\texttt{k})$ is equivalent to $\texttt{s} \in \mathcal{L}(\texttt{r1})^*\mathcal{L}(\texttt{k})$, and `accept(r,s,k)` $\rightsquigarrow$ `true` is equivalent to `accept_star(r1,s,k)` $\rightsquigarrow$ `true`. Thus, we want to show that

> evaluating `accept_star(r1,s,k)` terminates, and
> $\texttt{s} \in \mathcal{L}(\texttt{r1})^*\mathcal{L}(\texttt{k}) \Leftrightarrow$ `accept_star(r1,s,k)` $\rightsquigarrow$ `true`

We show this by showing the implication both ways.

**The $\Rightarrow$ direction:** We want to show that (for all strings `s`)

$$\texttt{s} \in \mathcal{L}(\texttt{r1})^*\mathcal{L}(\texttt{k}) \Rightarrow \texttt{accept_star(r1,s,k)} \rightsquigarrow \texttt{true}$$

To do this we show the equivalent

$$(\exists n.\texttt{s} \in \mathcal{L}(\texttt{r1})^n\mathcal{L}(\texttt{k})) \Rightarrow \texttt{accept_star(r1,s,k)} \rightsquigarrow \texttt{true}$$

by course-of-values induction on $n$.

- If $n = 0$ then $\texttt{s} \in \mathcal{L}(\texttt{r1})^0\mathcal{L}(\texttt{k}) = \mathcal{L}(\texttt{k})$. In that case `k(s)` $\rightsquigarrow$ `true`, so `accept_star(r1,s,k)` $\rightsquigarrow$ `true`.
- If $n > 0$ and $\texttt{s} \in \mathcal{L}(\texttt{r1})^n\mathcal{L}(\texttt{k})$, then either $n$ is minimal such that the property holds, or there is a $m < n$ such that $\texttt{s} \in \mathcal{L}(\texttt{r1})^m\mathcal{L}(\texttt{k})$.
  If there is such an $m$ then the induction hypothesis tells us immediately that `accept_star(r1,s,k)` $\rightsquigarrow$ `true`.
  If $n$ is minimal, then we know that $s \notin \mathcal{L}(\texttt{k})$, so `k(s)` $\rightsquigarrow$ `false` since it is assumed to terminate. Also, there exists $x$ and $y$ such that `s` $\rightsquigarrow x@y$, $x \in \mathcal{L}(\texttt{r1})^n$, and $y \in \mathcal{L}(k)$. Since $\mathcal{L}(\texttt{r1})^n = \mathcal{L}(\texttt{r1})\mathcal{L}(\texttt{r1})^{n-1}$, we can again split $x$ into $x = x_1@x_2$ where $x_1 \in \mathcal{L}(\texttt{r1})$ and $x_2 \in \mathcal{L}(\texttt{r1})^{n-1}$. Since $n$ was minimal, we know that $x \notin \mathcal{L}(r1)^{n-1}$, so $x_1$ is not the empty string.
  Then $x_2@y \in \mathcal{L}(\texttt{r1})^{n-1}\mathcal{L}(\texttt{k})$, so from the induction hypothesis we derive `accept_star(r1,`$x_2@y$`,k)` $\rightsquigarrow$ `true`, and since $x_1$ was not the empty string, $x_2@y \neq s$. If we let

$$k' = \texttt{fn s'=>not (s=s') andalso accept_star(r1,s',k)}$$

then $x_2@y \in \mathcal{L}(k')$, and so by the original induction hypothesis, `accept(r1,s,`$k'$`)` $\rightsquigarrow$ `true`, which means that `accept_star(r1,s,k)` $\rightsquigarrow$ `true`.

**Termination and the $\Leftarrow$ direction:** The proof for this case is by well-founded induction on the structure of `s`, in the ordering where a string is greater than or equal to its suffixes (and strictly greater than a proper suffix).

Evaluating `accept_star(r1,s,k)` terminates, since both `k(s)` terminates by assumption about `k`, and for any proper suffix of `s'`, the induction hypothesis tells us that evaluating `accept_star(r1,s',k)` terminates, so the function `(fn s'=> not (s=s') andalso accept_star(r1,s',k))` terminates on all suffixes of `s`, so the outer induction hypothesis tells us that evaluating `accept(r1,s,fn s'=>...)` terminates.

Assume `accept_star(r1,s,k)` $\rightsquigarrow$ `true`. This means that either

- `k(s)` $\rightsquigarrow$ `true`, in which case `s` $\in \mathcal{L}($`k`$) \subseteq \mathcal{L}($`STAR(r1)`$)\mathcal{L}($`k`$)$, or
- `k(s)` $\rightsquigarrow$ `false` and `accept(r1,s,fn s'=>...)` $=$ `true`. Let $k'$ denote

$$k' = \texttt{fn s'=>not (s=s') andalso accept_star(r1,s',k)}$$

  Then $k'$ terminates on all suffixes of `s`. On `s` itself it evaluates to false (comparing lists always terminates), and on all proper suffixes, `s'`, the induction hypothesis tells us that evaluating `accept_star(r1,s',k)` terminates.

  In this case the structural induction hypothesis tells us that `s` $\in \mathcal{L}($`r1`$)\mathcal{L}(k')$,

  By definition this means that there exists $x$ and $y$ such that `s` $= x@y$, $x \in \mathcal{L}($`r1`$)$, and $y \in \mathcal{L}(k')$ (i.e., both `not(s=`$y$`)` $\rightsquigarrow$ `true` and also `accept_star(r1,`$y$`,k)` $\rightsquigarrow$ `true`).

  Therefore $y \neq s$, so $y$ is a proper suffix of `s` (and $x$ is not the empty string). In that case the well-founded induction hypothesis gives us that $y \in \mathcal{L}($`r1`$)^*\mathcal{L}($`k`$)$, so we have shown

$$s = x@y \in \mathcal{L}(\texttt{r1})\mathcal{L}(\texttt{r1})^*\mathcal{L}(\texttt{k}) \subseteq \mathcal{L}(\texttt{r1})^*\mathcal{L}(\texttt{k}).$$

QED

## B Correctness Proof of the First-Order Matcher

To prove the correctness criterion given in Section 5.3.2,

$$\begin{array}{rcrcl} P_1(\mathtt{r},\mathtt{s},\mathtt{k}) & \stackrel{\text{def}}{=} & \mathtt{accept(r,s,k)} \rightsquigarrow \mathtt{true} & \Leftrightarrow & \mathtt{s} \in \mathcal{L}(\mathtt{r})\mathcal{L}(\mathtt{k}) \\ P_2(\mathtt{k},\mathtt{s}) & \stackrel{\text{def}}{=} & \mathtt{pop_and_accept(k,s)} \rightsquigarrow \mathtt{true} & \Leftrightarrow & \mathtt{s} \in \mathcal{L}(\mathtt{k}) \\ P_3(\mathtt{r},\mathtt{s},\mathtt{k}) & \stackrel{\text{def}}{=} & \mathtt{accept_star(r,s,k)} \rightsquigarrow \mathtt{true} & \Leftrightarrow & \mathtt{s} \in \mathcal{L}(\mathtt{r})^*\mathcal{L}(\mathtt{k}) \end{array}$$

by mutual well-founded induction, we give a measure on each proposition such that each case of the proof only depends on strictly smaller propositions. The measure of the propositions are given as pairs of natural numbers, ordered lexicographically.

$$\begin{array}{rcl} |P_1(\mathtt{r},\mathtt{s},\mathtt{k})| & = & (|\mathtt{s}|, |\mathtt{r}|_{|\mathtt{s}|} + |\mathtt{k}|_{|\mathtt{s}|}) \\ |P_2(\mathtt{k},\mathtt{s})| & = & (|\mathtt{s}|, |\mathtt{k}|_{|\mathtt{s}|}) \\ |P_3(\mathtt{r},\mathtt{s},\mathtt{k})| & = & (|\mathtt{s}|, |\mathtt{r}|_{|\mathtt{s}|}(3\,|\mathtt{s}|+2) + |\mathtt{k}|_{|\mathtt{s}|}) \\ & & \text{where } |\mathtt{s}| = \mathit{length}(\mathtt{s}) \end{array}$$

$$\begin{array}{rcl} |\mathtt{ZERO}|_n & = & 1 \\ |\mathtt{ONE}|_n & = & 1 \\ |\mathtt{CHAR(c)}|_n & = & 1 \\ |\mathtt{SUM(r1,r2)}|_n & = & |\mathtt{r1}|_n + |\mathtt{r2}|_n \\ |\mathtt{CAT(r1,r2)}|_n & = & |\mathtt{r1}|_n + |\mathtt{r2}|_n + 2 \\ |\mathtt{STAR(r)}|_n & = & (3n+2)\,|\mathtt{r}|_n \end{array}$$

$$\begin{array}{rcl} |\mathtt{EMPTY}|_n & = & 0 \\ |\mathtt{ACCEPT(r,k)}|_n & = & |\mathtt{r}|_n + |\mathtt{k}|_n + 1 \\ |\mathtt{ACCEPT_STAR(s,r,k)}|_n & = & |r|_{\min(3(n+1),3|\mathtt{s}|)} + |\mathtt{k}|_n \end{array}$$

The cases of the proof are:

$P_1(\mathtt{r},\mathtt{s},\mathtt{k})$**:**

- If $\mathtt{r} = \mathtt{ZERO}$ then $\mathtt{accept(r,s,k)} \rightsquigarrow \mathtt{true}$ does not hold, and neither does $\mathtt{s} \in \mathcal{L}(\mathtt{ZERO})\mathcal{L}(\mathtt{k}) = \emptyset$, so the bi-implication holds.
- If $\mathtt{r} = \mathtt{ONE}$ then $\mathtt{accept(r,s,k)} \rightsquigarrow \mathtt{true}$ if and only if $\mathtt{pop_and_accept(k,s)} \rightsquigarrow \mathtt{true}$. By well-ordered induction hypothesis, $\mathtt{pop_and_accept(k,s)} \rightsquigarrow \mathtt{true}$ is equivalent to $\mathtt{s} \in \mathcal{L}(\mathtt{k})$, and $\mathcal{L}(\mathtt{ONE}) = \{[]\}$, so $\mathtt{s} \in \mathcal{L}(\mathtt{ONE})\mathcal{L}(\mathtt{k})$.
- If $\mathtt{r} = \mathtt{CHAR(c)}$ then $\mathtt{accept(r,s,k)} \rightsquigarrow \mathtt{true}$ if and only if $\mathtt{s} = \mathtt{c::s'}$ and $\mathtt{pop_and_accept(k,s')} \rightsquigarrow \mathtt{true}$. In that case, $|s'| < |s|$ and then $|P_2(\mathtt{k},\mathtt{s'})| < |P_1(\mathtt{CHAR(c)},\mathtt{s},\mathtt{k})|$. By induction hypothesis $\mathtt{s'} \in \mathcal{L}(\mathtt{k})$ so $s = \mathtt{[c]@s'} \in \mathcal{L}(\mathtt{CHAR(c)})\mathcal{L}(\mathtt{k})$.
- If $\mathtt{r} = \mathtt{SUM(r1,r2)}$ then $\mathtt{accept(r,s,k)} \rightsquigarrow \mathtt{true}$ holds if and only if either $\mathtt{accept(r1,s,k)} \rightsquigarrow \mathtt{true}$ or $\mathtt{accept(r2,s,k)} \rightsquigarrow \mathtt{true}$. Since $|r|_n$ is always

positive, both of these are "smaller", so the induction hypothesis tells us that this is equivalent to `s` $\in \mathcal{L}($`r1`$)\mathcal{L}($`k`$)$ or `s` $\in \mathcal{L}($`r2`$)\mathcal{L}($`k`$)$, which is the same as `s` $\in (\mathcal{L}($`r1`$)\mathcal{L}($`k`$)) \cup (\mathcal{L}($`r2`$)\mathcal{L}($`k`$)) = (\mathcal{L}($`r1`$) \cup \mathcal{L}($`r2`$))\mathcal{L}($`k`$) = \mathcal{L}($`r`$)\mathcal{L}($`k`$)$.

- If `r = CAT(r1,r2)` then `accept(r,s,k)` $\rightsquigarrow$ `true` holds if and only if the equivalent `accept(r1,s,ACCEPT(r2,k))` $\rightsquigarrow$ `true` holds.

  From $P_1($`r1, s, ACCEPT(r2,k)`$)$, which is smaller in our ordering, we know that this is equivalent to `s` $\in \mathcal{L}($`r1`$)\mathcal{L}($`ACCEPT(r2,k)`$) = \mathcal{L}($`r1`$)\mathcal{L}($`r2`$)\mathcal{L}($`k`$) = \mathcal{L}($`r`$)\mathcal{L}($`k`$)$.

- If `r = STAR(r1)`, then assuming `accept(r,s,k)` $\rightsquigarrow$ `true` is equivalent to assuming `accept_star(r1,s,k)` $\rightsquigarrow$ `true`. Thus, we can just show $P_3($`r1, s, k`$)$ instead.

  Now, `accept_star(r1,s,k)` $\rightsquigarrow$ `true` if either `pop_and_accept(k,s)` $\rightsquigarrow$ `true` or `accept(r1,s,ACCEPT_STAR(s,r1,k))` $\rightsquigarrow$ `true`.

  The first case is equivalent to `s` $\in \mathcal{L}($`k`$) \subseteq \mathcal{L}($`r1`$)^*\mathcal{L}($`k`$)$, by induction hypothesis ($P_2($`k, s`$)$ is "smaller").

  The second case is equivalent to `s` $\in \mathcal{L}($`r1`$)\mathcal{L}($`ACCEPT_STAR(s,r1,k)`$) = \mathcal{L}($`r1`$)((\mathcal{L}($`r1`$)^*\mathcal{L}($`k`$)) \setminus \{$`s`$\})$, also by induction hypothesis.

  Using basic operations on sets, the disjunction of these two predicates can be seen to be equivalent to `s` being in the union of the sets, or equivalently, `s` $\in (\mathcal{L}($`r1`$)^0 \cup (\mathcal{L}($`r1`$) \setminus \{$`[]`$\})\mathcal{L}($`r1`$)^*)\mathcal{L}($`k`$) = \mathcal{L}($`r1`$)^*\mathcal{L}($`k`$)$.

  Notice that while $P_3($`r1, s, k`$)$ is not strictly smaller than $P_1($`STAR(r1), s, k`$)$, that is not a problem, since we are not assuming $P_3($`r1, s, k`$)$ here.

$P_2($`k, s`$)$**:**

- If `k = EMPTY` then `s` $\in \mathcal{L}($`k`$)$ if and only if `s` is the empty string, which is exactly when `pop_and_accept(k,s)` $\rightsquigarrow$ `true` holds.

- If `k = ACCEPT(r,k')` then `s` $\in \mathcal{L}($`k`$) = \mathcal{L}($`r`$)\mathcal{L}($`k'`$)$ which is equivalent to `accept(r,s,k')` $\rightsquigarrow$ `true` by induction hypothesis ($P_1($`r, s, k'`$)$ is smaller than $P_2($`ACCEPT(r,k'), s`$)$ by just one). In this case, `pop_and_accept(k,s)` $\equiv$ `accept(r,s,k')` $\rightsquigarrow$ `true`.

- If `k = ACCEPT_STAR(s',r,k')` then `s` $\in \mathcal{L}($`k`$)$ exactly when `s` $\neq$ `s'` and `s` $\in \mathcal{L}($`r`$)^*\mathcal{L}($`k'`$)$. By induction hypothesis, `s` $\in \mathcal{L}($`r`$)^*\mathcal{L}($`k'`$)$ is equivalent to `accept_star(r,s,k)` $\rightsquigarrow$ `true`, and the inequality is reflected in the ML program, so `not(s=s')` $\rightsquigarrow$ `true`. Taken together, `s` $\in \mathcal{L}($`k`$)$ holds exactly when `pop_and_accept(k,s)` $\rightsquigarrow$ `true` does.

$P_3($`r, s, k`$)$**:** This case is similar to the sub-induction in $P_1($`STAR(r), s, k`$)$, and indeed, $|P_3($`r, s, k`$)| = |P_1($`STAR(r), s, k`$)|$.

QED

## References


[1] Andrew W. Appel and Trevor Jim. Continuation-passing, closure-passing style. In Michael J. O'Donnell and Stuart Feldman, editors, *Proceedings of the Sixteenth Annual ACM Symposium on Principles of Programming Languages*, pages 293–302, Austin, Texas, January 1989. ACM Press.

[2] Anindya Banerjee, Nevin Heintze, and Jon G. Riecke. Semantics-based design and correctness of control-flow analysis-based program transformations. Unpublished, March 2001.

[3] Henk Barendregt. *The Lambda Calculus: Its Syntax and Semantics*, volume 103 of *Studies in Logic and the Foundation of Mathematics*. North-Holland, 1984. Revised edition.

[4] Jeffrey M. Bell, Françoise Bellegarde, and James Hook. Type-driven defunctionalization. In Mads Tofte, editor, *Proceedings of the 1997 ACM SIGPLAN International Conference on Functional Programming*, pages 25–37, Amsterdam, The Netherlands, June 1997. ACM Press.

[5] Anders Bondorf. *Self-Applicable Partial Evaluation*. PhD thesis, DIKU, Computer Science Department, University of Copenhagen, Copenhagen, Denmark, 1990. DIKU Rapport 90/17.

[6] Urban Boquist. *Code Optimization Techniques for Lazy Functional Languages*. PhD thesis, Department of Computing Science, Chalmers University of Technology, Göteborg University, Göteborg, Sweden, April 1999.

[7] Henry Cejtin, Suresh Jagannathan, and Stephen Weeks. Flow-directed closure conversion for typed languages. In Smolka [47], pages 56–71.

[8] Weidong Chen, Michael Kifer, and David S. Warren. Hilog: A foundation for higher-order logic programming. *The Journal of Logic Programming*, 15(3):187–230, February 1993.

[9] Alonzo Church. *The Calculi of Lambda-Conversion*. Princeton University Press, 1941.

[10] Olivier Danvy. Back to direct style. *Science of Computer Programming*, 22(3):183–195, 1994.

[11] Olivier Danvy. Formalizing implementation strategies for first-class continuations. In Smolka [47], pages 88–103.

[12] Olivier Danvy and Andrzej Filinski. Representing control, a study of the CPS transformation. *Mathematical Structures in Computer Science*, 2(4):361–391, 1992.

[13] Olivier Danvy, Bernd Grobauer, and Morten Rhiger. A unifying approach to goal-directed evaluation. *New Generation Computing*, 20(1), 2001. To appear. A preliminary version is available in the proceedings of SAIG 2001.

[14] Olivier Danvy and Julia L. Lawall. Back to direct style II: First-class continuations. In William Clinger, editor, *Proceedings of the 1992 ACM Conference on Lisp and Functional Programming*, LISP Pointers, Vol. V, No. 1, pages 299–310, San Francisco, California, June 1992. ACM Press.

[15] Olivier Danvy and Frank Pfenning. The occurrence of continuation parameters in CPS terms. Technical report CMU-CS-95-121, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, February 1995.

[16] Edsger W. Dijkstra. Recursive programming. In Saul Rosen, editor, *Programming Systems and Languages*, chapter 3C, pages 221–227. McGraw-Hill, New York, 1960.

[17] R. Kent Dybvig. *Three Implementation Models for Scheme*. PhD thesis, Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, April 1987. Technical Report #87-011.

[18] Leonidas Fegaras. lambda-DB. Available online at `http://lambda.uta.edu/lambda-DB/manual/`, 1999-2001.

[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] Daniel P. Friedman, Mitchell Wand, and Christopher T. Haynes. *Essentials of Programming Languages, second edition*. The MIT Press, 2001.

[21] Jean-Yves Girard. Locus solum. *Mathematical Structures in Computer Science*, 11(3), 2001. To appear.

[22] Joseph A. Goguen, James W. Thatcher, and Eric G. Wagner. An initial algebra approach to the specification, correctness and implementation of abstract data types. In *Current Trends in Programming Methodology*, volume IV, pages 80–149. Prentice-Hall, 1978.

[23] Adele Goldberg and David Robson. *Smalltalk-80: The Language and its Implementation*. Addison-Wesley, 1983.

[24] Mayer Goldberg. *Recursive Application Survival in the λ-Calculus*. PhD thesis, Computer Science Department, Indiana University, Bloomington, Indiana, May 1996.

[25] Robert Harper. Proof-directed debugging. *Journal of Functional Programming*, 9(4):463–469, July 1999.

[26] John E. Hopcroft and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.

[27] John Hughes. Super combinators: A new implementation method for applicative languages. In Daniel P. Friedman and David S. Wise, editors, *Conference Record of the 1982 ACM Symposium on Lisp and Functional Programming*, pages 1–10, Pittsburgh, Pennsylvania, August 1982. ACM Press.

[28] John Hughes. A novel representation of lists and its application to the function "reverse". *Information Processing Letters*, 22(3):141–144, 1986.

[29] Thomas Johnsson. Lambda lifting: Transforming programs to recursive equations. In Jean-Pierre Jouannaud, editor, *Functional Programming Languages and Computer Architecture*, number 201 in Lecture Notes in Computer Science, pages 190–203, Nancy, France, September 1985. Springer-Verlag.

[30] Peter J. Landin. The mechanical evaluation of expressions. *Computer Journal*, 6:308–320, 1964.

[31] 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.

[32] Yasuhiko Minamide, Greg Morrisett, and Robert Harper. Typed closure conversion. In Guy L. Steele Jr., editor, *Proceedings of the Twenty-Third Annual ACM Symposium on Principles of Programming Languages*, pages 271–283, St. Petersburg Beach, Florida, January 1996. ACM Press.

[33] Torben Æ. Mogensen. Efficient self-interpretation in lambda calculus. *Journal of Functional Programming*, 2(3):345–363, 1992.

[34] Tim Nicholson and Norman Y. Foo. A denotational semantics for Prolog. *ACM Transactions on Programming Languages and Systems*, 11(4):650–665, 1989.

[35] Lasse R. Nielsen. A denotational investigation of defunctionalization. Progress report (superseded by [36]), BRICS PhD School, University of Aarhus, June 1999.

[36] Lasse R. Nielsen. A denotational investigation of defunctionalization. Technical Report BRICS RS-00-47, DAIMI, Department of Computer Science, University of Aarhus, Aarhus, Denmark, December 2000.

[37] Simon L. Peyton Jones. *The Implementation of Functional Programming Languages*. Prentice Hall International Series in Computer Science. Prentice-Hall International, 1987.

[38] Simon L. Peyton Jones and David R. Lester. *Implementing Functional Languages*. Prentice Hall International Series in Computer Science. Prentice-Hall, 1992.

[39] Jeff Polakow. Linear logic programming with ordered contexts. In Maurizio Gabbrielli and Frank Pfenning, editors, *Proceedings of the Second International Conference on Principles and Practice of Declarative Programming*, pages 68–79, Montréal, Canada, September 2000. ACM Press.

[40] Jeff Polakow and Kwangkeun Yi. Proving syntactic properties of exceptions in an ordered logical framework. In Herbert Kuchen and Kazunori Ueda, editors, *Fifth International Symposium on Functional and Logic Programming*, number 2024 in Lecture Notes in Computer Science, pages 61–77, Tokyo, Japan, March 2001. Springer-Verlag.

[41] Todd A. Proebsting. Simple translation of goal-directed evaluation. In Ron K. Cytron, editor, *Proceedings of the ACM SIGPLAN'97 Conference on Programming Languages Design and Implementation*, SIGPLAN Notices, Vol. 32, No 5, pages 1–6, Las Vegas, Nevada, June 1997. ACM Press.

[42] John C. Reynolds. The essence of Algol. In van Vliet, editor, *International Symposium on Algorithmic Languages*, pages 345–372, Amsterdam, The Netherlands, 1982. North-Holland.

[43] 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).

[44] John C. Reynolds. Definitional interpreters revisited. *Higher-Order and Symbolic Computation*, 11(4):355–361, 1998.

[45] John C. Reynolds. *Theories of Programming Languages*. Cambridge University Press, 1998.

[46] Olin Shivers. *Control-Flow Analysis of Higher-Order Languages or Taming Lambda*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, May 1991. Technical Report CMU-CS-91-145.

[47] Gert Smolka, editor. *Proceedings of the Ninth European Symposium on Programming*, number 1782 in Lecture Notes in Computer Science, Berlin, Germany, March 2000. Springer-Verlag.

[48] Paul A. Steckler and Mitchell Wand. Lightweight closure conversion. *ACM Transactions on Programming Languages and Systems*, 19(1):48–86, 1997.

[49] Guy L. Steele Jr. Rabbit: A compiler for Scheme. Technical Report AI-TR-474, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, May 1978.

[50] Christopher Strachey. Fundamental concepts in programming languages. *Higher-Order and Symbolic Computation*, 13(1/2):1–49, 2000.

[51] Andrew Tolmach and Dino P. Oliva. From ML to Ada: Strongly-typed language interoperability via source translation. *Journal of Functional Programming*, 8(4):367–412, 1998.

[52] David A. Turner. A new implementation technique for applicative languages. *Software—Practice and Experience*, 9(1):31–49, 1979.

[53] Philip Wadler. Deforestation: Transforming programs to eliminate trees. *Theoretical Computer Science*, 73(2):231–248, 1989.

[54] Mitchell Wand. Continuation-based program transformation strategies. *Journal of the ACM*, 27(1):164–180, January 1980.

[55] Daniel C. Wang and Andrew W. Appel. Type-safe garbage collectors. In Hanne Riis Nielson, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on Principles of Programming Languages*, pages 166–178, London, United Kingdom, January 2001. ACM Press.

[56] Glynn Winskel. *The Formal Semantics of Programming Languages*. Foundation of Computing Series. The MIT Press, 1993.

[57] Yong Xiao, Amr Sabry, and Zena M. Ariola. From syntactic theories to interpreters: Automating proofs of decomposition lemma. *Higher-Order and Symbolic Computation*, 14(4), 2001. To appear.

# Recent BRICS Report Series Publications

**RS-01-23** Olivier Danvy and Lasse R. Nielsen. *Defunctionalization at Work*. June 2001. Extended version of an article to appear in Søndergaard, editor, *3rd International Conference on Principles and Practice of Declarative Programming*, PPDP '01 Proceedings, 2001.

**RS-01-22** Zoltán Ésik. *The Equational Theory of Fixed Points with Applications to Generalized Language Theory*. June 2001. 21 pp. To appear in Kuich, editor, *5th International Conference*, Developments in Language Theory DLT '01 Proceedings, LNCS, 2001.

**RS-01-21** Luca Aceto, Zoltán Ésik, and Anna Ingólfsdóttir. *Equational Theories of Tropical Semirings*. June 2001. 52 pp. Extended abstracts of parts of this paper have appeared in Honsell and Miculan, editors, *Foundations of Software Science and Computation Structures*, FoSSaCS '01 Proceedings, LNCS 2030, 2000, pages 42–56 and in Gaubert and Loiseau, editors, *Workshop on Max-plus Algebras and their Applications to Discrete-event Systems, Theoretical Computer Science, and Optimization*, MAX-PLUS '01 Proceedings, IFAC (International Federation of Automatic Control) IFAC Publications, 2001.

**RS-01-20** Catuscia Palamidessi and Frank D. Valencia. *A Temporal Concurrent Constraint Programming Calculus*. June 2001. 31 pp.

**RS-01-19** Jiří Srba. *On the Power of Labels in Transition Systems*. June 2001. 23 pp. Full and extended version of Larsen and Nielsen, editors, *Concurrency Theory: 12th International Conference*, CONCUR '01 Proceedings, LNCS, 2001.

**RS-01-18** Katalin M. Hangos, Zsolt Tuza, and Anders Yeo. *Some Complexity Problems on Single Input Double Output Controllers*. May 2001. 27 pp.

**RS-01-17** Claus Brabrand, Anders Møller, Steffan Olesen, and Michael I. Schwartzbach. *Language-Based Caching of Dynamically Generated HTML*. May 2001. 18 pp.

**RS-01-16** Olivier Danvy, Morten Rhiger, and Kristoffer H. Rose. *Normalization by Evaluation with Typed Abstract Syntax*. May 2001. 9 pp. To appear in *Journal of Functional Programming*.