BRICS RS-08-4 Danvy & Millikin: Refunctionalization at Work

Basic Research in Computer Science

Refunctionalization at Work

Olivier Danvy
Kevin Millikin

BRICS Report Series **RS-08-4**

ISSN 0909-0878 **June 2008**

**Copyright © 2008, Olivier Danvy & Kevin Millikin.
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
IT-parken, Aabogade 34
DK–8200 Aarhus N
Denmark
Telephone: +45 8942 9300
Telefax: +45 8942 5601
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/08/4/`

Refunctionalization at Work*

Olivier Danvy and Kevin Millikin
BRICS
Department of Computer Science
University of Aarhus†

Time-stamp: <2007-08-03 21:57:20 danvy>

Abstract

We present the left inverse of Reynolds's defunctionalization and we show its relevance to programming and to programming languages. We propose two methods to transform a program that is almost in defunctionalized form into one that is actually in defunctionalized form, and we illustrate them with a recognizer for Dyck words and with Dijkstra's shunting-yard algorithm.

Revised version of BRICS RS-07-7.

*This tutorial documents an invited talk at MPC '06 [26].

†IT-parken, Aabogade 34, DK-8200 Aarhus N, Denmark.
Email: <danvy@brics.dk>, <kmillikin@brics.dk>

Contents

1 Introduction **1**
1.1 Defunctionalization: origin, motivation, and applications 1
1.2 Refunctionalization 1
1.3 Overview 2

2 Refunctionalization: the left inverse of defunctionalization **2**
2.1 Defunctionalization 2
2.2 Example: from higher-order to first-order environments 3
2.3 Refunctionalization 5
2.4 Example: from first-order to higher-order environments 6

3 Towards putting programs in defunctionalized form **7**
3.1 Disentangling 8
3.2 Merging apply functions 8

4 A worked-out example: recognizing Dyck words **8**

5 A worked-out example: Dijkstra's shunting-yard algorithm **11**

6 Perspectives **15**
6.1 Other applications 15
6.2 Limitations and alternatives 16

7 Conclusion **16**
7.1 Wirth 16
7.2 Dijkstra 17

List of Figures

1 An environment signature 4
2 A Dyck-words recognizer in Standard ML 9
3 Dyck-words recognizer: disentangled version 10
4 Dyck-words recognizer: merged version 10
5 Dyck-words recognizer: refunctionalized version 10
6 Dyck-words recognizer: refunctionalized version expressed in direct style . . 10
7 Dijkstra's shunting-yard algorithm in ML 13
8 The shunting-yard algorithm: disentangled version 13
9 The shunting-yard algorithm: refunctionalized version expressed in direct style 14

1 Introduction

This article is a continuation of Danvy and Nielsen's earlier article "Defunctionalization at Work" [33], that outlined the extent to which Reynolds's defunctionalization [59] is pervasive in writing and transforming programs, and in specifying and implementing programming languages. Our goal here is to show that the left inverse of *de*functionalization, i.e., *re*functionalization, is also relevant to programming and to programming languages.

1.1 Defunctionalization: origin, motivation, and applications

Reynolds introduced defunctionalization to specify higher-order definitional interpreters using first-order means [59]. He presented it as a mere programming technique, and, except for deriving a first-order semantics in his textbook on programming languages [62, Section 12.4], he never used it again [61]. Since then, defunctionalization has been chiefly used in compilers [15, 17], program transformers [68], and partial evaluators [14, 20]. It has also been independently discovered in the context of logic programming [70], and subsequently it has been formalized in a typed setting [5, 6, 56, 58]. Today it is being used for programming the web [44] and for modeling aspect-oriented programming [39].

Over the last few years [1,10,25,53,54], Danvy and his students have retraced Reynolds's steps from higher-order to first-order interpreter by closure conversion (to make the data flow first order), transformation into Continuation-Passing Style (to sequentialize the control flow), and defunctionalization (to make the control flow first order). They have identified not only that a defunctionalized, CPS-transformed, and closure-converted interpreter has the structure of an abstract machine, but also that a large number of independently designed abstract machines for variants of the λ-calculus are the defunctionalized, CPS-transformed, and closure-converted counterpart of a compositional interpreter [2–4, 7, 11, 24, 32, 64], including Felleisen et al.'s CEK machine, which was actually designed as such [42, page 196]. A practical byproduct of this observation is that evaluation contexts—which are notoriously non-trivial to get right—can be obtained mechanically by defunctionalizing the continuations of a compositional interpreter.

1.2 Refunctionalization

Not all abstract machines, however, are in defunctionalized form. It is our thesis here [25, 54] that a number of them can be restated to be so. The goal of this article is to propose two techniques—disentangling and merging apply functions—to restate programs that are almost in defunctionalized form into ones that actually are in defunctionalized form (see Section 3). These programs can then be refunctionalized into a higher-order counterpart. In particular, it is our observation that refunctionalizing abstract machines always gives rise to a continuation-passing program.

Of course, not all defunctionalized programs and not all refunctionalized programs are interesting independently of each other, but still some of them are. We mentioned above the case of abstract machines for the λ-calculus. Another example is the samefringe problem: McCarthy's famous simple solution [52] and, e.g., Henderson and Morris's iterative solution based on lazy lists [47] are the defunctionalized / refunctionalized counterparts of each other [13, Section 5]. Another example is the reverse function: Hughes's efficient reverse function based on curried list constructors [48] and the usual accumulator-based fast reverse function are the defunctionalized / refunctionalized counterparts of each other [33]. More

examples are available elsewhere [25, 33, 43, 53, 54] as well as in Sections 4 and 5, but let us mention a last one that pertains to refunctionalization. On the basis that continuations are a *functional* representation of control in a continuation-passing evaluation function, Landin is not listed among the co-discoverers of continuations [60]. Yet his SECD machine is in defunctionalized form and the defunctionalized counterpart of a compositional continuation-passing evaluation function [24]. Furthermore, the J operator captures the current dump, i.e., in refunctionalized form, the continuation of the caller [32]. This observation has led us to suggest that Landin's name be added to the list of co-discoverers of continuations [32, Section 8].

1.3 Overview

We first review defunctionalization and its left inverse, refunctionalization (Section 2). We then propose two steps for restating programs from almost being in defunctionalized form to actually being in defunctionalized form: disentangling and merging apply functions (Section 3), and we illustrate these techniques with two parsing examples: recognizing Dyck words (Section 4) and Dijkstra's shunting-yard algorithm (Section 5). After reviewing other applications of refunctionalization (Section 6), we conclude (Section 7).

Prerequisites: We expect a basic familiarity with Standard ML,[1] with continuation-passing style (CPS) [29, 57, 67], and with the left inverse of the CPS transformation [31]. The presentation of defunctionalization below should be self-contained, but the reader should not deny himself the pleasure of re-reading "Definitional Interpreters" [59].

2 Refunctionalization: the left inverse of defunctionalization

In his study of definitional interpreters for programming languages [59], Reynolds drew a distinction between those that use higher-order functions, and thus possibly depend on the scoping rules of their metalanguage, and those that use only first-order data structures, and thus are independent of the scoping rules of their metalanguage. He introduced defunctionalization to transform programs that use higher-order functions into ones that use first-order data structures. Refunctionalization is the inverse transformation. It transforms programs that use first-order data structures into ones that use higher-order functions.

2.1 Defunctionalization

Operationally, defunctionalization replaces a function type with a polynomial (sum of products) data type, and introduces an apply function that interprets the data type given the argument values of the original function type. It replaces abstractions creating elements of the function type with applications of the data-type constructors. It replaces applications of values of the function type to arguments with an application of the apply function to a value of the first-order type and the same arguments. The defunctionalization algorithm can be informally sketched as follows.

The algorithm assumes that one has identified a target set of function abstractions. In a typed language, these abstractions minimally all have the same type. Reynolds identified sets of abstractions by their "types" according to the semantic domains of a definitional

[1] And with the option data type: `datatype 'a option = NONE | SOME of 'a`.

interpreter. Later work [5, 6, 33, 56, 58] considered all functions of a given type, or sets of functions determined by a control-flow analysis [66]. In any case, the set has to be closed under data flow to call sites: if a function in the set can be called from a call site in the program, then *all* functions that can be called from that call site are in the set.

We assume the identification of a set of n abstractions $\{\lambda x.M_i \mid 1 \leq i \leq n\}$ all of type $\tau \to \tau'$, and closed under data flow to call sites.[2] Each abstraction has a (possibly empty) set of m_i typed free variables $\{x^i_j : \tau^i_j \mid 1 \leq j \leq m_i\}$, not including variables that are bound in the top-level environment. Defunctionalization consists of performing all of the following steps:

1. **Introduce a first-order data type:** There are n data-type constructors C_i for $1 \leq i \leq n$. The constructors uniquely identify the abstractions being defunctionalized, and each variant of the data type represents one of the abstractions. There are thus n variants and variant i represents the values of the variables occurring free in abstraction i:

$$\begin{aligned} \textit{datatype } \tau_\textit{arrow}_\tau' = {} & C_1 \textit{ of } \tau^1_1 \times \ldots \times \tau^1_{m_1} \\ \mid {} & \ldots \\ \mid {} & C_n \textit{ of } \tau^n_1 \times \ldots \times \tau^n_{m_n} \end{aligned}$$

   Introduce an apply function: The apply function dispatches on the data-type constructor to determine which abstraction is being applied and to bind its free variables:

$$\begin{aligned} \textit{fun apply}(f, x) = \textit{case } f \textit{ of} & \\ & C_1(x^1_1, \ldots, x^1_{m_1}) \Rightarrow M_1 \\ \mid {} & \ldots \\ \mid {} & C_n(x^n_1, \ldots, x^n_{m_n}) \Rightarrow M_n \end{aligned}$$

   The apply function may be curried or not; here, it is not.

2. **Replace abstractions:** An abstraction $\lambda x.M_i$ with free variables $\{x^i_j : \tau^i_j \mid 1 \leq j \leq m_i\}$ is replaced with an application of the data-type constructor, $C_i(x^i_1, \ldots, x^i_{m_i})$.

   Replace applications: Any call site $f\ x$ that may be an application of one of the abstractions is replaced with a call to the apply function, $\textit{apply}(f, x)$.

2.2 Example: from higher-order to first-order environments

An environment is a mapping from identifiers to values that minimally satisfies the ML signature displayed in Figure 1. To satisfy this signature, a module should define a representation of identifiers (e.g., as strings), a type of environments that is polymorphic with respect to the values denoted by the identifiers, a representation of an empty environment, a function `extend` to extend a given environment with a given identifier and a given value, a function `lookup` to retrieve the value associated with a given identifier in a given environment, and an exception `UNBOUND` to be raised by `lookup` if the given identifier is not in the domain of the given environment.

Environments are traditionally implemented either as functions from identifiers to values or as association lists. In this section we show that defunctionalizing the former implementation yields the latter one.

Here is a functional implementation:

[2] For simplicity we consider single-argument functions. Without loss of generality, defunctionalization and refunctionalization can be applied to programs with multiple-argument, curried or uncurried, functions as well.

```
signature ENVIRONMENT
= sig
    type ide
    type 'a env
    exception UNBOUND of ide
    val empty  : 'a env
    val extend : ide * 'a * 'a env -> 'a env
    val lookup : ide * 'a env -> 'a
  end
```

Figure 1: An environment signature

```
structure Higher_Order_Environment :> ENVIRONMENT
= struct
    type ide = string
    type 'a env = ide -> 'a  (* the environment as a function *)
    exception UNBOUND of ide

    val empty = fn y => raise (UNBOUND y)

    fun extend (x, v, e) = fn y => if x = y then v else e y

    fun lookup (y, e) = e y
  end
```

Here, an environment is a function mapping an identifier to a value or raising the exception `UNBOUND`. Looking up an identifier in an environment is achieved by applying this environment to this identifier.

Let us defunctionalize this representation of environments. In Standard ML, the opaque signature ascription (`:>`) prevents this function type from being used outside of this structure, and therefore we can proceed independently of the use of this structure elsewhere in a program.

Only two abstractions give rise to inhabitants of the function type `ide -> 'a`:

- `fn y => raise (UNBOUND y)`,
  which has no free variables (we do not consider the exception name `UNBOUND` to be free because it will be lexically visible in the apply function), and
- `fn y => if x = y then v else e y`,
  which has three free variables: `x` of type `ide`, `v` of type `'a`, and `e` of type `ide -> 'a`.

Let us follow the steps outlined in Section 2.1:

1. **Introduce a first-order data type:** The function type `ide -> 'a` is replaced by the following new data type:

   ```
   datatype 'a ide_arrow_alpha = C1
                               | C2 of ide * 'a * 'a ide_arrow_alpha
   ```

   This data type is isomorphic to `(ide * 'a) list`, which we use in the sequel.

Introduce an apply function: The apply function dispatches on the constructors of the first-order data type (`[]` and `::`), and interprets each of them as the body of the corresponding higher-order function:

```
fun apply ([], y)
    = raise UNBOUND y
  | apply ((x, v) :: e, y)
    = if x = y then v else e y
```

2. **Replace abstractions:** The two abstractions are replaced with `[]` and `(x, v) :: e` respectively.

   Replace applications: The two applications, `e x` in the original `lookup` function and `e y` in the apply function just above, are respectively replaced with `apply (e, x)` and `apply (e, y)`.

The result of defunctionalization is the traditional first-order representation of environments as an association list:

```
structure First_Order_Environment :> ENVIRONMENT
= struct
    type ide = string
    type 'a env = (ide * 'a) list  (* the environment as an association list *)
    exception UNBOUND of ide

    val empty = []

    fun extend (x, v, e) = (x, v) :: e

    fun apply ([], y)
        = raise (UNBOUND y)
      | apply ((x, v) :: e, y)
        = if x = y then v else apply (e, y)

    fun lookup (y, e) = apply (e, y)
  end
```

2.3 Refunctionalization

Refunctionalization is the left inverse of defunctionalization, mapping data types and apply functions in the image of defunctionalization back to higher-order functions. Danvy and Nielsen were the first to consider an inverse of defunctionalization [33].

A program with a first-order data structure and an apply function dispatching on it is in the image of Reynolds's defunctionalization algorithm if this apply function is the sole point of consumption of values of the data type. More generally, if there is only one case dispatch on the data type, then the dispatch can be abstracted into an apply function whose arguments are the free variables of the entire case expression and whose return type is the type of the case expression, as described in Section 3.1.

Once a data type and an apply function are recognized as being in defunctionalized form, refunctionalization proceeds by reversing the steps of defunctionalization. A data type δ with an apply function of type $\delta \times \tau \to \tau'$ can be refunctionalized into the functional type $\tau \to \tau'$. Refunctionalization consists of performing all of the following steps:

1. **Replace applications of the apply function:** A call to the apply function, $apply(f, x)$, is replaced with the call to the applied function, $f\ x$.

   Replace data-type constructor applications: Data-type constructor applications are replaced with abstractions based on the apply function. We assume an apply function:

$$\begin{aligned} \mathit{fun}\ apply(f, x) = \mathit{case}\ f\ \mathit{of} \\ \quad C_1(x_1^1, \ldots, x_{m_1}^1) &\Rightarrow M_1 \\ \mid \ \ldots \\ \mid \ C_n(x_1^n, \ldots, x_{m_n}^n) &\Rightarrow M_n \end{aligned}$$

   Each constructor application $C_i(v_1, \ldots, v_{m_i})$ of C_i applied to values v_j for $1 \leq j \leq m_i$ is replaced by $(\lambda x.M_i)\{v_1/x_1^i, \ldots, v_{m_i}/x_{m_i}^i\}$, the capture-avoiding substitution of the constructor arguments for the free variables in the abstraction represented by the apply function. If a constructor is applied to non-values, we insert let-bindings for the non-value arguments before performing this step.

2. **Remove the apply function:** Since the apply function is never called, its definition is no longer needed.

   Remove the data-type definition: Since the constructors of the data type are never used, and the only case dispatch on them (the apply function) has been removed, the data-type definition is no longer needed.

Refunctionalization is akin to the Scott-encoding [69] of a data type, i.e., the functional representation of a data type as its case-dispatch function.

2.4 Example: from first-order to higher-order environments

In this section, symmetrically to Section 2.2, we show how refunctionalizing a traditional first-order implementation of an environment as an association list yields a traditional higher-order implementation of an environment as a function.

Here is a traditional first-order implementation using an association list:

```
structure First_Order_Environment' :> ENVIRONMENT
= struct
    type ide = string
    type 'a env = (ide * 'a) list
    exception UNBOUND of ide

    val empty = []

    fun extend (x, v, e) = (x, v) :: e

    fun lookup (y, e)
        = let fun visit []
                  = raise (UNBOUND y)
                | visit ((x, v) :: e)
                  = if x = y
                    then v
                    else visit e
          in visit e
          end
  end
```

Here, an environment is an association list, i.e., a list of pairs of identifiers and values. An empty environment is represented as the empty list and extending a given environment with a given identifier and a given value is achieved by pairing them and consing the pair on the given environment. Looking up an identifier in an environment is achieved by traversing the association list in search for the first matching identifier and raising the exception `UNBOUND` if none does.

Let us refunctionalize this representation of environments, which we can do locally because of the opaque signature ascription (`:>`) that isolates the definition of this structure from its uses. We identify the association list as a first-order data structure with a sole point of consumption, namely the `visit` function. To make `visit` fit the pattern that the apply function is explicitly applied to two arguments, we first apply the first step of lambda-lifting [35, 49] to `visit`, i.e., we make it scope-insensitive by passing it explicitly its free variable `y`:

```
fun lookup (y, e)
    = let fun visit ([], y)
              = raise (UNBOUND y)
            | visit ((x, v) :: e, y)
              = if x = y
                then v
                else visit (e, y)
      in visit (e, y)
      end
```

(The second step of lambda-lifting would be to make the definition of `visit` float to the same lexical level as that of `lookup`, thereby obtaining recursive equations.)

With that, we follow the steps outlined in Section 2.3:

1. **Replace applications of the apply function:** We replace calls to the apply function, `visit (e, y)`, with calls to the applied function, `e y`.

   Replace data-type constructor applications: We replace data-type constructor applications with abstractions based on the apply function:

   - `[]` is replaced by `fn y => raise (UNBOUND y)`, and
   - `(x, v) :: e` is replaced by `fn y => if x = y then v else e y`.

2. We then remove the definition of `visit`.

The result of refunctionalization is the traditional higher-order representation of environments as a function mapping identifiers to values.

3 Towards putting programs in defunctionalized form

The hallmark of a defunctionalized data type and an apply function is that the apply function is the single point of consumption for the data type (this characterization is due to Danvy and Nielsen [33]). In this section, we present two simple transformations for programs that almost, but don't quite, have this form: disentangling (Section 3.1) abstracts data-type consumptions into functions, and merging (Section 3.2) combines multiple candidate apply functions for the same data type when possible. Both transformations are illustrated in Sections 4 and 5.

3.1 Disentangling

A program can fail to be in the image of defunctionalization if there are multiple points of consumption for elements of a data type or if the single point of consumption is not in a separate function. This situation can occur, for example, when a program is defunctionalized and then the apply function is inlined.

'Disentangling' consists of abstracting each case expression dispatching on a candidate data type into a separate function. The arguments of the function are the free variables of the case expression and their types are the types of the free variables. The return type is the type of the case expression.

Disentangling was first used for the transition function of the SECD machine [24]. Disentangling the transition function of an abstract machine consists of splitting single transitions that dispatch simultaneously on multiple components of the state into multiple serial transitions that dispatch separately on single components of the state. We illustrate this technique in Figure 3 of Section 4 for a push-down automaton recognizing Dyck words.

3.2 Merging apply functions

After disentangling, every case dispatch on a data type is abstracted into a function which is a candidate apply function. If there is a single apply function, then the program can be refunctionalized. If there are multiple apply functions, then they can be merged under some conditions.

A technique that frequently works for abstract machines is to merge the apply functions using the universal property of sum types. Specifically, two functions $apply_1 : \delta \times \tau_1 \to \tau'$ and $apply_2 : \delta \times \tau_2 \to \tau'$ can be merged into a single function $apply : \delta \times (\tau_1 + \tau_2) \to \tau'$ that performs a case dispatch on its second argument. To reflect this merging, calls to $apply_1$ and $apply_2$ are adjusted to call *apply* with their second argument injected into the appropriate summand. (Dually, we can also merge the co-domains of the apply functions.)

Merging apply functions was first used to refunctionalize Burge's version of the SECD machine with the J operator [32]. This technique works when the possible apply functions all have the same return types. In abstract machines this is usually the case because the transition functions are all tail-recursive and thus share the same answer type, as illustrated in Figure 4 of Section 4 for a push-down automaton recognizing Dyck words.

4 A worked-out example: recognizing Dyck words

Dyck words are well-balanced words of left and right parentheses, which we represent in ML as follows:

```
datatype parenthesis = L | R

type word = parenthesis list
```

For example, the list `[L, L, R, L, R, R]` represents a Dyck word whereas the list `[R, L]` does not.

Dyck words are classically recognized with an abstract machine implementing a push-down automaton. This state-transition system operates iteratively over a given list and a counter reflecting the number of open parentheses seen in the list so far:

```
datatype nat = ZERO | SUCC of nat
```

The machine starts with a given word and a zero counter. At each iteration, one of the following transitions takes place:

- if the list of parentheses is empty and the counter is zero, a final, accepting state is reached;
- if the list of parentheses is empty and the counter is positive, a final, non-accepting state is reached;
- if the first parenthesis is a left one, the tail of the list is taken and the counter is incremented;
- if the first parenthesis is a right one and if the counter is zero, a final, non-accepting state is reached;
- if the first parenthesis is a right one and if the counter is positive, the tail of the list is taken and the counter is decremented.

Figure 2 displays an ML program implementing this transition system.

```
(*  recognize : word -> bool  *)
fun recognize ps
    = let (*  run : word * nat -> bool  *)
          fun run (     [], ZERO  ) = true
            | run (     [], SUCC c) = false
            | run (L :: ps, c     ) = run (ps, SUCC c)
            | run (R :: ps, ZERO  ) = false
            | run (R :: ps, SUCC c) = run (ps, c     )
      in run (ps, ZERO)
      end
```

Figure 2: A Dyck-words recognizer in Standard ML

The goal of this section is to refunctionalize this Dyck-words recognizer with respect to the counter. Graphically, we proceed as follows:

Figure 2 —disentangle→ Figure 3 —merge / apply functions→ Figure 4 —refunctionalize→ Figure 5 —transform to direct style→ Figure 6

In Figure 2, both parameters of `run` serve as induction variables: `run` dispatches both over the list of parentheses and over the counter. We therefore disentangle `run` by introducing two specialized, mutually recursive versions: one with respect to an empty word (`run_nil`) and the other with respect to a right parenthesis (`run_par`). The resulting disentangled program is displayed in Figure 3: `run` solely dispatches over the list of parentheses, and `run_nil` and `run_par` solely dispatch over the counter.

The transition system displayed in Figure 3 operates in lockstep with the original transition system,[3] but it is not in defunctionalized form with respect to the counter: the counter

[3]The new machine takes one or two steps for each step in the original one.

```
fun recognize_disentangled ps
    = let (*  run : word * nat -> bool  *)
          fun run (     [], c) = run_nil c
            | run (L :: ps, c) = run (ps, SUCC c)
            | run (R :: ps, c) = run_par (c, ps)
          (*  run_nil : nat -> bool  *)
          and run_nil ZERO     = true
            | run_nil (SUCC n) = false
          (*  run_par : nat * word -> bool  *)
          and run_par (ZERO  , ps) = false
            | run_par (SUCC c, ps) = run (ps, c)
      in run (ps, ZERO)
      end
```

Figure 3: Dyck-words recognizer: disentangled version

```
fun recognize_merged ps
    = let (*  run : word * nat -> bool  *)
          fun run (     [], c) = run_aux (c, NONE)
            | run (L :: ps, c) = run (ps, SUCC c)
            | run (R :: ps, c) = run_aux (c, SOME ps)
          (*  run_aux : nat * word option -> bool  *)
          and run_aux (ZERO  , NONE   ) = true
            | run_aux (SUCC c, NONE   ) = false
            | run_aux (ZERO  , SOME ps) = false
            | run_aux (SUCC c, SOME ps) = run (ps, c)
      in run (ps, ZERO)
      end
```

Figure 4: Dyck-words recognizer: merged version

```
fun recognize_refunctionalized ps
    = let (*  run : word * (word option -> bool) -> bool  *)
          fun run (     [], c) = c NONE
            | run (L :: ps, c) = run (ps, fn NONE    => false
                                        | SOME ps => run (ps, c))
            | run (R :: ps, c) = c (SOME ps)
      in run (ps, fn NONE    => true
                  | SOME ps => false)
      end
```

Figure 5: Dyck-words recognizer: refunctionalized version

```
fun recognize_in_direct_style ps
    = callcc (fn exit => let (*  run : word -> word option  *)
                             fun run []        = NONE
                               | run (L :: ps) = (case run ps
                                                   of NONE    => throw exit false
                                                    | SOME ps => run ps)
                               | run (R :: ps) = SOME ps
                         in case run ps
                             of NONE    => true
                              | SOME ps => false
                         end)
```

Figure 6: Dyck-words recognizer: refunctionalized version expressed in direct style

is dispatched upon both in `run_nil : nat -> bool` and in `run_par : nat * word -> bool`. We therefore merge `run_nil` and `run_par` into a function `run_aux : nat * word option -> bool`. The resulting merged program is displayed in Figure 4. It is in defunctionalized form with respect to the data type `nat` and its apply function `run_aux`.

The refunctionalized counterpart of the merged program in Figure 4 is displayed in Figure 5. Its function `word option -> bool` is the refunctionalized counterpart of the data type `nat` and the apply function `run_aux`. This program is in CPS, and except for errors, it uses its continuations linearly and in order. Its direct-style counterpart is displayed in Figure 6. It is a recursive program where not all calls are in tail position. As an aid to the eye, we have shaded the non-tail calls in grey. Errors are handled with `callcc` and `throw` [31], though of course using an exception would do as well here.

The counter that was explicit in the original recognizer (Figure 2), the disentangled one (Figure 3), the merged one (Figure 4), and the refunctionalized one (Figure 5) is now implicit in the direct-style program (Figure 6). In other words, the original recognizer was implemented by a tail-recursive push-down automaton that managed an explicit data stack—namely the counter. This explicit data stack is now implicit in the control stack of the language processor for the recursive program in direct style.

5 A worked-out example: Dijkstra's shunting-yard algorithm

The shunting-yard algorithm is used to parse an arithmetic expression with operator precedence from infix form (as a stream of tokens: literals, operators, and parentheses) to postfix form (again, as a stream of tokens) or to an abstract-syntax tree. It is named for the railroad shunting yard because it uses a pair of tracks (stacks): one to store operand subtrees and one to store operators until all the operands are complete.

This bottom-up parser is attributed to Dijkstra,[4] and was developed for and used in one of the first Algol 60 compilers. It is still used today to process binary expressions in the GCC parser for C and for Objective-C, for example.

The goal of this section is to refunctionalize a shunting-yard parser with respect to its stack of operators. The version we consider here maps a stream of tokens (integers, addition operator, multiplication operator, left parenthesis, or right parenthesis) to an abstract-syntax tree:

```
datatype token = LIT of int | ADD | MUL | L_P | R_P

datatype expression = INT of int
                    | PLUS of expression * expression
                    | TIMES of expression * expression
```

For example, the ML program implementing the algorithm maps the list of tokens

```
[LIT 10, MUL, LIT 20, ADD, LIT 30, MUL, L_P, LIT 40, ADD, LIT 50, R_P]
```

into the abstract-syntax tree

```
PLUS (TIMES (INT 10, INT 20), TIMES (INT 30, PLUS (INT 40, INT 50)))
```

[4] "In the summer of 1959 I had discovered how to implement subroutines that could call themselves, by the end of the year I saw how to use a stack for the evaluation of expressions and how to translate expressions from the usual infix notation into 'reverse Polish' (only we did not call it that)." [38]

that represents $10 \times 20 + 30 \times (40 + 50)$. Using `++` and `**` as ML infix notation for `PLUS` and `TIMES`, the result reads `((INT 10) ** (INT 20)) ++ ((INT 30) ** ((INT 40) ++ (INT 50)))`. We make use of this infix notation in Figures 7, 8, and 9 to construct lists of expressions.

The algorithm is defined with an abstract machine implementing a pushdown automaton with two stacks. This state-transition system operates iteratively over a stack of subtrees and a stack of operators (`XADD` and `XMUL`) and parenthetical separators (`XPAR`):

```
datatype operator = XADD | XMUL | XPAR
```

The machine starts with a given list of tokens, an empty stack of subtrees, and an empty stack of operators. At each iteration, one of the following transitions takes place:

- if the first token of the input list is a literal (`LIT n`), the tail of the list is taken and the corresponding expression (`INT n`) is pushed on the stack of subtrees;
- if the first token of the input list is an operator `x` (`ADD` or `MUL`), then
  - if the stack of operators is empty, the operator on top of this stack is a parenthesis (`XPAR`), or its precedence is strictly lower than that of `x`, the tail of the input list is taken and the operator corresponding to `x` (`XADD` or `XMUL`) is pushed on the stack of operators;
  - otherwise, the precedence of the operator on top of the stack is higher than or the same as that of `x`; this operator is popped off the stack, two expressions are popped off the stack of subtrees, the corresponding expression combining the popped operator and the popped expressions is pushed back, and the operator corresponding to `x` is pushed on the stack;
- if the first token is a left parenthesis (`L_P`), the tail of the input list is taken and the corresponding operator (`XPAR`) is pushed on the stack;
- if the first token is a right parenthesis (`R_P`), operators are popped off the stack (and for each of them, two expressions are popped off the stack of subtrees and the corresponding expression combining the popped operator and the popped expressions is pushed back) until a parenthesis is met on the stack of operators; the tail of the input list is taken and the parenthesis is popped from the stack of operators;
- otherwise, a final, non-accepting state is reached.

When the list of tokens is exhausted, the stack of operators is iteratively emptied in concert with the stack of subtrees as described above. The final result is the expression on top of the stack of subtrees.

Figure 7 displays a program written in Standard ML that implements the shunting-yard algorithm. We proceed as follows to refunctionalize it with respect to the stack of operators:

Figure 7 —disentangle→ Figure 8 —refunctionalize and transform to direct style→ Figure 9

As an ML program, the transition system in Figure 7 is difficult to read because several of the parameters of `run` serve as induction variables: `run` dispatches both over the list of tokens

```
fun parse ts = let
      (*  run : token list * expression list * operator list -> expression option  *)
      fun run (       ts as [],        e :: [],              []) = SOME e
        | run (       ts as [], e2 :: e1 :: es,       XADD :: xs) = run (ts, e1 ++ e2 :: es,         xs)
        | run (       ts as [], e2 :: e1 :: es,       XMUL :: xs) = run (ts, e1 ** e2 :: es,         xs)
        | run (    LIT n :: ts,             es,              xs) = run (ts,    INT n :: es,         xs)
        | run (      ADD :: ts,             es,              []) = run (ts,             es, XADD :: [])
        | run (      ADD :: ts,             es, xs as XPAR ::  _) = run (ts,             es, XADD :: xs)
        | run (ts as ADD ::  _, e2 :: e1 :: es,       XADD :: xs) = run (ts, e1 ++ e2 :: es,         xs)
        | run (ts as ADD ::  _, e2 :: e1 :: es,       XMUL :: xs) = run (ts, e1 ** e2 :: es,         xs)
        | run (      MUL :: ts,             es,              []) = run (ts,             es, XMUL :: [])
        | run (      MUL :: ts,             es, xs as XPAR ::  _) = run (ts,             es, XMUL :: xs)
        | run (      MUL :: ts,             es, xs as XADD ::  _) = run (ts,             es, XMUL :: xs)
        | run (ts as MUL ::  _, e2 :: e1 :: es,       XMUL :: xs) = run (ts, e1 ** e2 :: es,         xs)
        | run (      L_P :: ts,             es,              xs) = run (ts,             es, XPAR :: xs)
        | run (      R_P :: ts,             es,       XPAR :: xs) = run (ts,             es,         xs)
        | run (ts as R_P ::  _, e2 :: e1 :: es,       XADD :: xs) = run (ts, e1 ++ e2 :: es,         xs)
        | run (ts as R_P ::  _, e2 :: e1 :: es,       XMUL :: xs) = run (ts, e1 ** e2 :: es,         xs)
        | run (             ts,             es,              xs) = NONE
    in run (ts, [], [])
    end
```

Figure 7: Dijkstra's shunting-yard algorithm in ML

```
fun parse ts = let
      fun run_nil (         [], e :: []) = SOME e
        | run_nil (LIT n :: ts,      es) = run_nil (ts, INT n :: es)
        | run_nil (  ADD :: ts,      es) = run_add (ts,          es, [])
        | run_nil (  MUL :: ts,      es) = run_mul (ts,          es, [])
        | run_nil (  L_P :: ts,      es) = run_par (ts,          es, [])
        | run_nil (         ts,      es) = NONE

      and run_add (       ts as [], e2 :: e1 :: es, xs) = run_aux (ts, e1 ++ e2 :: es,         xs)
        | run_add (    LIT n :: ts,             es, xs) = run_add (ts,    INT n :: es,         xs)
        | run_add (ts as ADD ::  _, e2 :: e1 :: es, xs) = run_aux (ts, e1 ++ e2 :: es,         xs)
        | run_add (      MUL :: ts,             es, xs) = run_mul (ts,             es, XADD :: xs)
        | run_add (      L_P :: ts,             es, xs) = run_par (ts,             es, XADD :: xs)
        | run_add (ts as R_P ::  _, e2 :: e1 :: es, xs) = run_aux (ts, e1 ++ e2 :: es,         xs)
        | run_add (             ts,             es, xs) = NONE

      and run_mul (       ts as [], e2 :: e1 :: es, xs) = run_aux (ts, e1 ** e2 :: es,         xs)
        | run_mul (    LIT n :: ts,             es, xs) = run_mul (ts,    INT n :: es,         xs)
        | run_mul (ts as ADD ::  _, e2 :: e1 :: es, xs) = run_aux (ts, e1 ** e2 :: es,         xs)
        | run_mul (ts as MUL ::  _, e2 :: e1 :: es, xs) = run_aux (ts, e1 ** e2 :: es,         xs)
        | run_mul (      L_P :: ts,             es, xs) = run_par (ts,             es, XMUL :: xs)
        | run_mul (ts as R_P ::  _, e2 :: e1 :: es, xs) = run_aux (ts, e1 ** e2 :: es,         xs)
        | run_mul (             ts,             es, xs) = NONE

      and run_par (    LIT n :: ts,             es, xs) = run_par (ts,    INT n :: es,         xs)
        | run_par (      ADD :: ts,             es, xs) = run_add (ts,             es, XPAR :: xs)
        | run_par (      MUL :: ts,             es, xs) = run_mul (ts,             es, XPAR :: xs)
        | run_par (      L_P :: ts,             es, xs) = run_par (ts,             es, XPAR :: xs)
        | run_par (      R_P :: ts,             es, xs) = run_aux (ts,             es,         xs)
        | run_par (             ts,             es, xs) = NONE

      and run_aux (ts, es,         []) = run_nil (ts, es)
        | run_aux (ts, es, XADD :: xs) = run_add (ts, es, xs)
        | run_aux (ts, es, XMUL :: xs) = run_mul (ts, es, xs)
        | run_aux (ts, es, XPAR :: xs) = run_par (ts, es, xs)
    in run_nil (ts, [])
    end
```

Figure 8: The shunting-yard algorithm: disentangled version

```
fun parse ts =
    callcc (fn exit => let fun run_nil (             [],        e :: []) = SOME e
                             | run_nil (    LIT n :: ts,             es) = run_nil (ts, INT n :: es)
                             | run_nil (      ADD :: ts,             es) = run_nil (run_add (ts, es))
                             | run_nil (      MUL :: ts,             es) = run_nil (run_mul (ts, es))
                             | run_nil (      L_P :: ts,             es) = run_nil (run_par (ts, es))
                             | run_nil (             ts,             es) = NONE

                           and run_add (       ts as [], e2 :: e1 :: es) = (ts, e1 ++ e2 :: es)
                             | run_add (    LIT n :: ts,             es) = run_add (ts, INT n :: es)
                             | run_add (ts as ADD ::  _, e2 :: e1 :: es) = (ts, e1 ++ e2 :: es)
                             | run_add (      MUL :: ts,             es) = run_add (run_mul (ts, es))
                             | run_add (      L_P :: ts,             es) = run_add (run_par (ts, es))
                             | run_add (ts as R_P ::  _, e2 :: e1 :: es) = (ts, e1 ++ e2 :: es)
                             | run_add (             ts,             es) = throw exit NONE

                           and run_mul (       ts as [], e2 :: e1 :: es) = (ts, e1 ** e2 :: es)
                             | run_mul (    LIT n :: ts,             es) = run_mul (ts, INT n :: es)
                             | run_mul (ts as ADD ::  _, e2 :: e1 :: es) = (ts, e1 ** e2 :: es)
                             | run_mul (ts as MUL ::  _, e2 :: e1 :: es) = (ts, e1 ** e2 :: es)
                             | run_mul (      L_P :: ts,             es) = run_mul (run_par (ts, es))
                             | run_mul (ts as R_P ::  _, e2 :: e1 :: es) = (ts, e1 ** e2 :: es)
                             | run_mul (             ts,             es) = throw exit NONE

                           and run_par (    LIT n :: ts,             es) = run_par (ts, INT n :: es)
                             | run_par (      ADD :: ts,             es) = run_par (run_add (ts, es))
                             | run_par (      MUL :: ts,             es) = run_par (run_mul (ts, es))
                             | run_par (      L_P :: ts,             es) = run_par (run_par (ts, es))
                             | run_par (      R_P :: ts,             es) = (ts, es)
                             | run_par (             ts,             es) = throw exit NONE
                       in run_nil (ts, [])
                       end)
```

Figure 9: The shunting-yard algorithm: refunctionalized version expressed in direct style

and over the stack of operators. (In contrast, the stack of subtrees is only threaded passively and, except for errors in the input list of tokens, it does not influence the control flow of the program.)

We therefore disentangle `run` into several specialized, mutually recursive versions: one with respect to an empty stack of operators (`run_nil`), three with respect to each of the possible operators on top of the stack (`run_add`, `run_mul`, and `run_par`), and one to dispatch on the stack of operators (`run_aux`). The resulting disentangled program is displayed in Figure 8: all of `run_nil`, `run_add`, `run_mul`, and `run_par` solely dispatch over the list of tokens, and `run_aux` solely dispatches over the stack of operators. This transition system operates in lockstep with the original one implemented in Figure 7: the new machine takes one or two steps for each step in the original one.

In addition to being more readable, the transition system implemented in Figure 8 is also in defunctionalized form: the stack of operators and `run_aux` respectively form a data type and an apply function that are in the image of defunctionalization.

As for all transition systems in defunctionalized form, the refunctionalized counterpart of the program in Figure 8 is in CPS. Furthermore—again save for errors in the input list of tokens—it uses its continuations linearly and in order, and can therefore be expressed in direct style. We spare the reader with this CPS program (those who don't like this kind of things etc.), and display the corresponding direct-style program in Figure 9: it is a recursive program where not all calls are in tail position. As an aid to the eye, we have shaded the

non-tail calls in grey. Errors are handled with `callcc` and `throw` [31], though again using an exception would do as well.

The stack of operators that was explicit in Figures 7 and 8 is now implicit in the direct-style program of Figure 9. In other words, the original algorithm was implemented by a tail-recursive automaton with two stacks. The stack of operators is now implicit in the control stack of the language processor for this recursive program in direct style.

The refunctionalized shunting-yard algorithm still exhibits the standard features of a bottom-up parser: a return corresponds to a reduce action, a simple tail call with trivial arguments corresponds to a shift action, and a non-tail call corresponds to a state transition where the control-flow at return time implements the goto transition associated with a reduce action.

6 Perspectives

Let us list other applications of refunctionalization (Section 6.1) and then mention its current limitations and alternatives (Section 6.2).

6.1 Other applications

We had to intervene in the following situations to put an abstract machine into defunctionalized form:

The SECD machine: The SECD machine only needed to be disentangled to be put in defunctionalized form: the C and the D components can be refunctionalized into a control continuation and a dump continuation, respectively [24].

The SECD machine with the J operator: Two versions of the SECD machine with the J operator exist—the original one by Landin and Burge [16] and a version due to Felleisen [40]. Disentangling Felleisen's version yields a machine that is in defunctionalized form (and uses a control delimiter). In contrast, disentangling Landin and Burge's version is not sufficient to put it in defunctionalized form; its apply functions also need to be merged [32].

Properly tail-recursive stack inspection: In Clements and Felleisen's abstract machine for properly tail-recursive stack inspection [18], stack inspection is achieved by traversing the current context and checking its permission tables. Unzipping this context into an ordinary CEK-machine context and a list of permission tables yields a CEK machine with a state and an error facility: this machine can be refunctionalized and mapped back to direct style all the way to a compositional monadic evaluation function with a state+error monad [4]. Conversely, via refocusing [34], it syntactically corresponds to a new version of Fournet and Gordon's λ_{sec}-calculus [9, Section 7].

Strong-reduction strategies: In our ongoing work on strongly reducing abstract machines [54, 55], we use both disentangling and merging to refunctionalize a number of abstract machines [21, 22, 45, 50, 51] into compositional normalization functions.

6.2 Limitations and alternatives

Dynamic CPS: The original abstract machine for dynamic delimited-continuation operators [41] is not in defunctionalized form. Disentangling it and then merging its apply functions was not enough for our purpose: we also needed to thread a tail of delimited continuations [12]. It was then simple to make it correspond, through refocusing, to a reduction semantics, and through refunctionalization, to a compositional evaluation function. Still, the trail of delimited continuations is an ad-hoc device for which we have found no other use so far.

Towards de-objectification and re-objectification: Object orientation offers an alternative to merging apply functions with different domain and codomain types: the data type could be represented by an abstract class with a method for each of the apply functions. Each variant of the data type would be represented by a subclass of this abstract class with a field for each of the variant's free variables. This 'reobjectified' program can be implemented in a language without objects using any of the standard encodings of objects in objectless languages. The original 'deobjectified' program would then be such an encoding, specialized to flat class hierarchies.

7 Conclusion

We have investigated the applicability of refunctionalization for transforming programs and implementing programming languages. Elsewhere, we have illustrated its relevance to programming [13, 30, 33] and to specifying programming languages [2–4, 11, 24, 32].

7.1 Wirth

In some sense, our work on defunctionalization and refunctionalization provides a concrete illustration for the paragraph that follows the definition of the quicksort algorithm, in Section 2.3.3 of Niklaus Wirth's textbook "Algorithms and Data Structures" from 1985:

> *Procedure* `sort` *activates itself recursively. Such use of recursion in algorithms is a very powerful tool and will be discussed further in Chapter 3. In some programming languages of older provenience, recursion is disallowed for certain technical reasons. We will now show how this same algorithm can be expressed as a non-recursive procedure. Obviously, the solution is to express recursion as an iteration, whereby a certain amount of additional bookkeeping operations become necessary.*

Dijkstra's shunting-yard algorithm and Landin's SECD machine were directly written for programming languages of 'old provenience.' In Section 5 and in our earlier work [24, 32], we have shown how each is the defunctionalized and CPS-transformed counterpart of a recursive program in direct style. Incidentally, the same can be said of the quicksort algorithm in Wirth's book: CPS-transforming and defunctionalizing the recursive definition that precedes the paragraph above in the book yields the iterative definition that immediately follows in the book. So the 'certain amount of additional bookkeeping operations' mentioned above can be mechanized using the CPS transformation and defunctionalization. The same could be said for their respective correctness proofs: instead of developing them separately [46, 71], these proofs could be considered in the light of defunctionalization and refunctionalization [33, Section 5]—a future work.

7.2 Dijkstra

Dijkstra's case against the GOTO statement [37] has been mostly interpreted in the negative, as the forceful denunciation of a programming sin. It seems to us, though, that his message, 40 years on, can also be understood as an invitation to mindfulness when programming. Indeed consider a compiler from a structured language (e.g., one with while loops and conditional commands) to an unstructured language (e.g., one with labels and with conditional and unconditional jumps): on the one hand, this compiler yields programs that use GOTO statements; on the other hand, as denotations of structured programs, these unstructured programs only use GOTO statements to implement the control structures of structured programs. In that light, Dijkstra's implicit message is not so much that GOTO statements should be considered harmful, no matter what, than one should be mindful about staying in or straying from the image of the compiler when programming in the unstructured language:

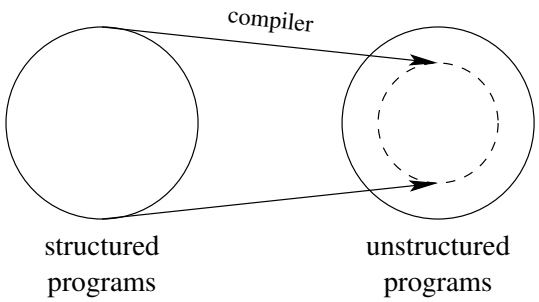

In fact, finding oneself straying with good reason is a clear indication that a useful control construct is missing in the source language. For example, C and Pascal programmers condone the use of GOTO for error cases because these languages lack an exception mechanism.

Dijkstra's implicit message applies to at least two situations involving a program transformation and its left inverse:

CPS transformation. When programming in continuation-passing style, one should be mindful of the continuation identifiers and of the parameters of continuations to stay in the image of the CPS transformation [23, 27]:

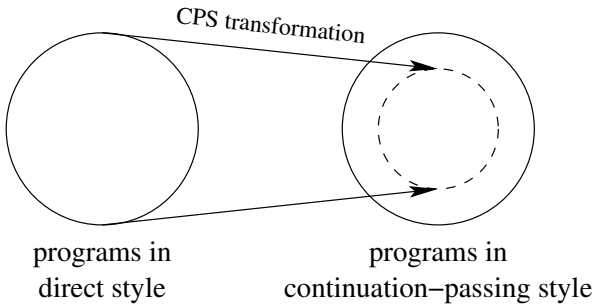

Yet programmers use "the extra expressive power of CPS" to stray with good reason:

- for a simple example, not using the current continuation identifier prevents the computation from continuing and therefore has the effect of aborting it; this effect can be obtained in direct style by adding an "abort" control operator;

- for a common example, using another continuation identifier than the current one makes the computation continue elsewhere; this effect can be obtained in direct style by adding, e.g., the control operator escape [59] or call/cc [19,31];
- for a more radical example, one may altogether leave the language of CPS, where all calls are tail calls [67], and mix CPS with non-tail calls, thereby introducing the notion of delimited control [25, Section 1.3]; this effect can be obtained in direct style by adding, e.g., the delimited-control operators shift and reset [28].

Defunctionalization. Since, as pointed out in Section 1.1, defunctionalizing a CPS program yields an abstract machine, one should be mindful about specifying abstract machines in defunctionalized form:

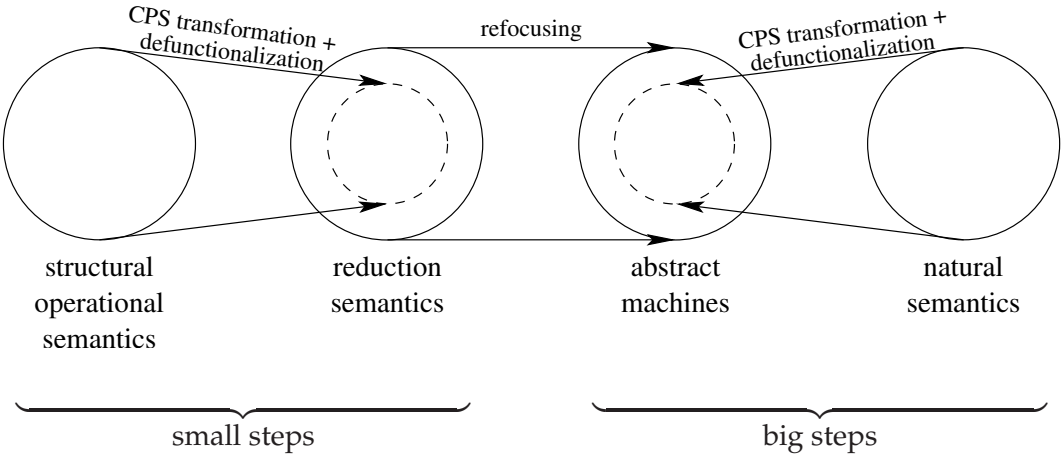

Indeed

- CPS-transforming and defunctionalizing a function implementing a small-step semantics yields a function implementing a reduction semantics [25];
- symmetrically, CPS-transforming and defunctionalizing a function implementing a big-step semantics yields a function implementing an abstract machine [2–4,7, 11,24,32]; and
- refocusing the evaluation function of a reduction semantics yields a function implementing an abstract machine [8,9,34].

Straying from the image of the CPS transformation makes it possible for the reduction semantics on the left and for the abstract machine on the right to specify, e.g., control effects.

In any case, most abstract machines have been designed independently of defunctionalization. It is our experience that a number of them are in defunctionalized form, and that a number of others can be transformed to be so, using the techniques described here (see Section 6) or using, e.g., GADTs to make them well typed [58]. We currently have no sense about why one would want an abstract machine to stray from being in defunctionalized form when using a functional meta-language for specifying computation.

So overall, we see the CPS transformation and defunctionalization as useful guidelines that are consistent with Dijkstra's implicit message.

Getting back to the shunting-yard algorithm, it seems very plausible that Dijkstra programmed it initially in the image of the above-mentioned compiler, while mentally picturing it as a structured program:

```
initialize the operator stack to be empty;
initialize the operand stack to be empty;
WHILE there are tokens in the input stream
DO peek the first such token;
   IF this first token is a number
   THEN read it and push it on the operand stack
   ELSIF this first token is an operator
   THEN IF the stack of operators is empty or has, at its top,
           a parenthetical separator or an operator with a lower precedence
        THEN read this first token and push it on the operator stack
        ELSE reduce
   ELSIF this first token is a left parenthesis
   THEN read it and push a parenthetical separator on the operator stack
   ELSIF this first token is a right parenthesis
   THEN IF the top operator is a parenthetical separator
        THEN read this first token and pop the top operator off the stack
        ELSE reduce
OD;
WHILE the operator stack is non-empty
DO reduce
OD
(* The result is on top of the operand stack. *)
```

In Section 5, we have (1) observed that the shunting-yard algorithm takes the form of an abstract machine, and (2) nudged it to be in defunctionalized form. Refunctionalizing it yields a program in CPS. Mapping this program back to direct style yields a functional parser, farther and farther away from the Turing tarpit:

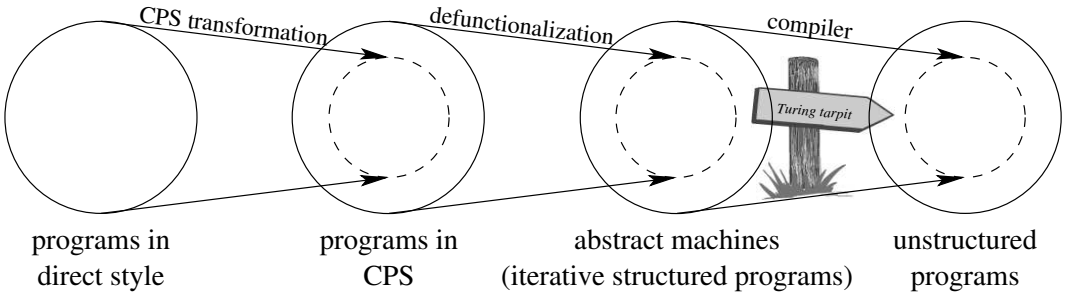

To close, when Dijkstra submitted "A Case against the GO TO Statement" [36] to CACM 40 years ago, Wirth used his editorial discretion to change the title to "Go To Statement Considered Harmful" [37]. All proportions kept, we would be grateful to our editor if he could leave the title of the present article as it is now: witness the drawing just above, abstract machines in defunctionalized form are harmless. As for those that can be disentangled etc., it is our point that they are mostly so.

Acknowledgments: The first author is grateful to Tarmo Uustalu for his invitation to MPC '06, and doubly so for not changing our title.[5] Thanks are also due to Dariusz Biernacki, Kristian Støvring, and the anonymous reviewers of Science of Computer Programming for their comments. This work is partly supported by the Danish Natural Science Research Council, Grant no. 272-06-0530.

References

[1] Mads Sig Ager. *Partial Evaluation of String Matchers & Constructions of Abstract Machines*. PhD thesis, BRICS PhD School, University of Aarhus, Aarhus, Denmark, January 2006.

[2] Mads Sig Ager, Dariusz Biernacki, Olivier Danvy, and Jan Midtgaard. A functional correspondence between evaluators and abstract machines. In Dale Miller, editor, *Proceedings of the Fifth ACM-SIGPLAN International Conference on Principles and Practice of Declarative Programming (PPDP'03)*, pages 8–19, Uppsala, Sweden, August 2003. ACM Press.

[3] Mads Sig Ager, Olivier Danvy, and Jan Midtgaard. A functional correspondence between call-by-need evaluators and lazy abstract machines. *Information Processing Letters*, 90(5):223–232, 2004. Extended version available as the research report BRICS RS-04-3.

[4] Mads Sig Ager, Olivier Danvy, and Jan Midtgaard. A functional correspondence between monadic evaluators and abstract machines for languages with computational effects. *Theoretical Computer Science*, 342(1):149–172, 2005. Extended version available as the research report BRICS RS-04-28.

[5] Anindya Banerjee, Nevin Heintze, and Jon G. Riecke. Design and correctness of program transformations based on control-flow analysis. In Naoki Kobayashi and Benjamin C. Pierce, editors, *Theoretical Aspects of Computer Software, 4th International Symposium, TACS 2001*, number 2215 in Lecture Notes in Computer Science, pages 420–447, Sendai, Japan, October 2001. Springer-Verlag.

[6] 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*, SIGPLAN Notices, Vol. 32, No. 8, pages 25–37, Amsterdam, The Netherlands, June 1997. ACM Press.

[7] Małgorzata Biernacka, Dariusz Biernacki, and Olivier Danvy. An operational foundation for delimited continuations in the CPS hierarchy. *Logical Methods in Computer Science*, 1(2:5):1–39, November 2005. A preliminary version was presented at the Fourth ACM SIGPLAN Workshop on Continuations (CW'04).

[8] Małgorzata Biernacka and Olivier Danvy. A concrete framework for environment machines. *ACM Transactions on Computational Logic*, 200? In press. Available as the research report BRICS RS-06-3.

[9] Małgorzata Biernacka and Olivier Danvy. A syntactic correspondence between context-sensitive calculi and abstract machines. *Theoretical Computer Science*, 375(1-3):76–108, 2007. Extended version available as the research report BRICS RS-06-18.

[5] But just in case, it should be "Refunctionalization at Work."

[10] Dariusz Biernacki. *The Theory and Practice of Programming Languages with Delimited Continuations*. PhD thesis, BRICS PhD School, University of Aarhus, Aarhus, Denmark, December 2005.

[11] Dariusz Biernacki and Olivier Danvy. From interpreter to logic engine by defunctionalization. In Maurice Bruynooghe, editor, *Logic Based Program Synthesis and Transformation, 13th International Symposium, LOPSTR 2003*, number 3018 in Lecture Notes in Computer Science, pages 143–159, Uppsala, Sweden, August 2003. Springer-Verlag.

[12] Dariusz Biernacki, Olivier Danvy, and Kevin Millikin. A dynamic continuation-passing style for dynamic delimited continuations. *ACM Transactions on Programming Languages and Systems*, 200? Accepted for publication (February 2007). A preliminary version is available as the research report BRICS RS-06-15.

[13] Dariusz Biernacki, Olivier Danvy, and Chung-chieh Shan. On the static and dynamic extents of delimited continuations. *Science of Computer Programming*, 60(3):274–297, 2006.

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

[15] 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.

[16] William H. Burge. *Recursive Programming Techniques*. Addison-Wesley, 1975.

[17] Henry Cejtin, Suresh Jagannathan, and Stephen Weeks. Flow-directed closure conversion for typed languages. In Gert Smolka, editor, *Proceedings of the Ninth European Symposium on Programming*, number 1782 in Lecture Notes in Computer Science, pages 56–71, Berlin, Germany, March 2000. Springer-Verlag.

[18] John Clements and Matthias Felleisen. A tail-recursive semantics for stack inspection. *ACM Transactions on Programming Languages and Systems*, 26(6):1029–1052, 2004. A preliminary version was presented at the 12th European Symposium on Programming (ESOP 2003).

[19] William Clinger, Daniel P. Friedman, and Mitchell Wand. A scheme for a higher-level semantic algebra. In *Algebraic Methods in Semantics*, pages 237–250. Cambridge University Press, 1985.

[20] Charles Consel. A tour of Schism: A partial evaluation system for higher-order applicative languages. In David A. Schmidt, editor, *Proceedings of the Second ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, pages 145–154, Copenhagen, Denmark, June 1993. ACM Press.

[21] Pierre Crégut. Strongly reducing variants of the Krivine abstract machine. *Higher-Order and Symbolic Computation*, 20(3):209–230, 2007. A preliminary version was presented at the 1990 ACM Conference on Lisp and Functional Programming.

[22] Pierre-Louis Curien. *Categorical Combinators, Sequential Algorithms and Functional Programming*. Progress in Theoretical Computer Science. Birkhaüser, 1993.

[23] Olivier Danvy. Back to direct style. *Science of Computer Programming*, 22(3):183–195, 1994. A preliminary version was presented at the Fourth European Symposium on Programming (ESOP 1992).

[24] Olivier Danvy. A rational deconstruction of Landin's SECD machine. In Clemens Grelck, Frank Huch, Greg J. Michaelson, and Phil Trinder, editors, *Implementation and Application of Functional Languages, 16th International Workshop, IFL'04*, number 3474 in Lecture Notes in Computer Science, pages 52–71, Lübeck, Germany, September 2004. Springer-Verlag. Recipient of the 2004 Peter Landin prize. Extended version available as the research report BRICS RS-03-33.

[25] Olivier Danvy. *An Analytical Approach to Program as Data Objects*. DSc thesis, Department of Computer Science, University of Aarhus, Aarhus, Denmark, October 2006.

[26] Olivier Danvy. Refunctionalization at work. In Tarmo Uustalu, editor, *Preliminary proceedings of the 8th International Conference on Mathematics of Program Construction (MPC '06)*, Kuressaare, Estonia, July 2006. Invited talk.

[27] Olivier Danvy, Belmina Dzafic, and Frank Pfenning. On proving syntactic properties of CPS programs. In *Third International Workshop on Higher-Order Operational Techniques in Semantics*, volume 26 of *Electronic Notes in Theoretical Computer Science*, pages 19–31, Paris, France, September 1999.

[28] Olivier Danvy and Andrzej Filinski. Abstracting control. In Mitchell Wand, editor, *Proceedings of the 1990 ACM Conference on Lisp and Functional Programming*, pages 151–160, Nice, France, June 1990. ACM Press.

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

[30] Olivier Danvy and Mayer Goldberg. There and back again. *Fundamenta Informaticae*, 66(4):397–413, 2005. A preliminary version was presented at the 2002 ACM SIGPLAN International Conference on Functional Programming (ICFP 2002).

[31] 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.

[32] Olivier Danvy and Kevin Millikin. A rational deconstruction of Landin's J operator. Research Report BRICS RS-06-17, DAIMI, Department of Computer Science, University of Aarhus, Aarhus, Denmark, December 2006. A preliminary version appears in the proceedings of 17th International Workshop on the Implementation and Application of Functional Languages. Accepted to Logical Methods in Computer Science.

[33] Olivier Danvy and Lasse R. Nielsen. Defunctionalization at work. In Harald Søndergaard, editor, *Proceedings of the Third International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP'01)*, pages 162–174, Firenze, Italy, September 2001. ACM Press. Extended version available as the research report BRICS RS-01-23.

[34] Olivier Danvy and Lasse R. Nielsen. Refocusing in reduction semantics. Research Report BRICS RS-04-26, DAIMI, Department of Computer Science, University of Aarhus, Aarhus, Denmark, November 2004. A preliminary version appeared in the informal proceedings of the Second International Workshop on Rule-Based Programming (RULE 2001), Electronic Notes in Theoretical Computer Science, Vol. 59.4.

[35] Olivier Danvy and Ulrik P. Schultz. Lambda-lifting in quadratic time. *Journal of Functional and Logic Programming*, 2004(1), July 2004. Available online at <http://danae.uni-muenster.de/lehre/kuchen/JFLP/>. A preliminary version was presented at the Sixth International Symposium on Functional and Logic Programming (FLOPS 2002).

[36] Edsger W. Dijkstra. A case against the GO TO statement. EWD 215, 1968. Available online at <http://www.cs.utexas.edu/users/EWD/>.

[37] Edsger W. Dijkstra. Go To statement considered harmful. *Communications of the ACM*, 11(3):147–148, March 1968. Letter to the editor.

[38] Edsger W. Dijkstra. From my life. EWD 1166, 1993. Available online at <http://www.cs.utexas.edu/users/EWD/>.

[39] Christopher J. Dutchyn. Specializing continuations: A model for dynamic join points. In Curtis Clifton, Gary T. Leavens, and Mira Mezini, editors, *Proceedings of the Sixth Workshop on Foundations of Aspect-Oriented Languages (FOAL '07)*, ACM International Conference Proceedings Series, pages 45–57, Vancouver, British Columbia, Canada, March 2007. ACM Press. Available online at <http://www.cs.iastate.edu/~leavens/FOAL/index-2007.shtml>.

[40] Matthias Felleisen. Reflections on Landin's J operator: a partly historical note. *Computer Languages*, 12(3/4):197–207, 1987.

[41] Matthias Felleisen. The theory and practice of first-class prompts. In Jeanne Ferrante and Peter Mager, editors, *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Programming Languages*, pages 180–190, San Diego, California, January 1988. ACM Press.

[42] Matthias Felleisen and Daniel P. Friedman. Control operators, the SECD machine, and the λ-calculus. In Martin Wirsing, editor, *Formal Description of Programming Concepts III*, pages 193–217. Elsevier Science Publishers B.V. (North-Holland), Amsterdam, 1986.

[43] Jean-Christophe Filliâtre and François Pottier. Producing all ideals of a forest, functionally. *Journal of Functional Programming*, 13(5):945–956, 2003.

[44] Paul T. Graunke, Robert Bruce Findler, Shriram Krishnamurthi, and Matthias Felleisen. Automatically restructuring programs for the web. In Martin S. Feather and Michael Goedicke, editors, *16th IEEE International Conference on Automated Software Engineering (ASE 2001)*, pages 211–222, Coronado Island, San Diego, California, USA, November 2001. IEEE Computer Society.

[45] Benjamin Grégoire and Xavier Leroy. A compiled implementation of strong reduction. In Simon Peyton Jones, editor, *Proceedings of the 2002 ACM SIGPLAN International Conference on Functional Programming (ICFP'02)*, SIGPLAN Notices, Vol. 37, No. 9, pages 235–246, Pittsburgh, Pennsylvania, September 2002. ACM Press.

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

[47] Peter Henderson and James H. Morris Jr. A lazy evaluator. In Susan L. Graham, editor, *Proceedings of the Third Annual ACM Symposium on Principles of Programming Languages*, pages 95–103. ACM Press, January 1976.

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

[49] 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.

[50] Werner E. Kluge. *Abstract Computing Machines: A Lambda Calculus Perspective*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2005.

[51] Pierre Lescanne. From λσ to λν a journey through calculi of explicit substitutions. In Hans-J. Boehm, editor, *Proceedings of the Twenty-First Annual ACM Symposium on Principles of Programming Languages*, pages 60–69, Portland, Oregon, January 1994. ACM Press.

[52] John McCarthy. Another samefringe. *SIGART Newsletter*, 61, February 1977.

[53] Jan Midtgaard. *Transformation, Analysis, and Interpretation of Higher-Order Procedural Programs*. PhD thesis, BRICS PhD School, University of Aarhus, Aarhus, Denmark, June 2007.

[54] Kevin Millikin. *A Structured Approach to the Transformation, Normalization and Execution of Computer Programs*. PhD thesis, BRICS PhD School, University of Aarhus, Aarhus, Denmark, May 2007.

[55] Johan Munk. A study of syntactic and semantic artifacts and its application to lambda definability, strong normalization, and weak normalization in the presence of state. Master's thesis, DAIMI, Department of Computer Science, University of Aarhus, Aarhus, Denmark, May 2007.

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

[57] Gordon D. Plotkin. Call-by-name, call-by-value and the λ-calculus. *Theoretical Computer Science*, 1:125–159, 1975.

[58] François Pottier and Nadji Gauthier. Polymorphic typed defunctionalization and concretization. *Higher-Order and Symbolic Computation*, 19(1):125–162, 2006. A preliminary version was presented at the Thirty-First Annual ACM Symposium on Principles of Programming Languages (POPL 2004).

[59] John C. Reynolds. Definitional interpreters for higher-order programming languages. In *Proceedings of 25th ACM National Conference*, pages 717–740, Boston, Massachusetts, 1972. Reprinted in Higher-Order and Symbolic Computation 11(4):363–397, 1998, with a foreword [61].

[60] John C. Reynolds. The discoveries of continuations. *Lisp and Symbolic Computation*, 6(3/4):233–247, 1993.

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

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

[63] David A. Schmidt. State transition machines for lambda calculus expressions. In Neil D. Jones, editor, *Semantics-Directed Compiler Generation*, number 94 in Lecture Notes in Computer Science, pages 415–440, Aarhus, Denmark, 1980. Springer-Verlag.

[64] David A. Schmidt. State-transition machines for lambda-calculus expressions. *Higher-Order and Symbolic Computation*, 20(3):319–332, 2007. Journal version of [63], with a foreword [65].

[65] David A. Schmidt. State-transition machines, revisited. *Higher-Order and Symbolic Computation*, 20(3):319–332, 2007.

[66] 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.

[67] Guy L. Steele Jr. Rabbit: A compiler for Scheme. Master's thesis, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, May 1978. Technical report AI-TR-474.

[68] Andrew P. 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.

[69] Christopher P. Wadsworth. Some unusual λ-calculus numeral schemes. In *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, pages 215–230. Academic Press, 1980.

[70] David H. D. Warren. Higher-order extensions to PROLOG: are they needed? In J. E. Hayes, Donald Michie, and Y.-H. Pao, editors, *Machine Intelligence*, volume 10, pages 441–454. Ellis Horwood, 1982.

[71] Kwangkeun Yi. "Proof-directed debugging" revisited for a first-order version. *Journal of Functional Programming*, 16(6):663–670, 2006.

Recent BRICS Report Series Publications

RS-08-4 Olivier Danvy and Kevin Millikin. *Refunctionalization at Work*. June 2008. ii+25 pp. To appear in *Science of Computer Programming*. A preliminary version is available as the research report BRICS RS-07-7.

RS-08-3 Johan Munk. *A Study of Syntactic and Semantic Artifacts and its Application to Lambda Definability, Strong Normalization, and Weak Normalization in the Presence of State*. April 2008. xi+144 pp.

RS-08-2 Gudmund Skovbjerg Frandsen and Piotr Sankowski. *Dynamic Normal Forms and Dynamic Characteristic Polynomial*. April 2008. 21 pp. To appear in ICALP '08.

RS-08-1 Anders Møller. *Static Analysis for Event-Based XML Processing*. jan 2008. 23 pp. Appears in PLAN-X '08.

RS-07-18 Jan Midtgaard. *Control-Flow Analysis of Functional Programs*. December 2007. iii+38 pp.

RS-07-17 Luca Aceto, Willem Jan Fokkink, and Anna Ingólfsdóttir. *A Cancellation Theorem for 7BCCSP*. December 2007. 30 pp.

RS-07-16 Olivier Danvy and Kevin Millikin. *On the Equivalence between Small-Step and Big-Step Abstract Machines: A Simple Application of Lightweight Fusion*. November 2007. ii+11 pp. To appear in Information Processing Letters (extended version). Supersedes BRICS RS-07-8.

RS-07-15 Jooyong Lee. *A Case for Dynamic Reverse-code Generation*. August 2007. ii+10 pp.

RS-07-14 Olivier Danvy and Michael Spivey. *On Barron and Strachey's Cartesian Product Function*. July 2007. ii+14 pp.

RS-07-13 Martin Lange. *Temporal Logics Beyond Regularity*. July 2007. 82 pp.

RS-07-12 Gerth Stølting Brodal, Rolf Fagerberg, Allan Grønlund Jørgensen, Gabriel Moruz, and Thomas Mølhave. *Optimal Resilient Dynamic Dictionaries*. July 2007.