

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity (draft)

Jeremy Gibbons^a

^a Department of Computer Science, University of Oxford

Abstract *Context:* Reynolds showed us how to use continuation-passing style and defunctionalization to transform a recursive interpreter for a language into an abstract machine for programs in that language. The same techniques explain other programming tricks, including zippers and accumulating parameters. *Inquiry:* Buried within all those applications there is usually a hidden appeal to the algebraic property of associativity. *Approach:* Our purpose in this paper is to entice associativity out of the shadows and into the limelight. *Knowledge:* We revisit some well-known applications (factorial, fast reverse, tree flattening, and a compiler for a simple expression language) to spotlight their dependence on associativity. *Grounding:* We replay developments of these programs through a series of program transformations and data refinements, justified by equational reasoning. *Importance:* Understanding the crucial role played by associativity clarifies when continuation-passing style and defunctionalization can help and when they cannot, and may prompt other applications of these techniques.

ACM CCS 2012

- **Software and its engineering** → **Functional languages; Data types and structures; Patterns; Interpreters; Formal software verification;**
- **Theory of computation** → **Program reasoning;**

Keywords program transformation, accumulating parameters, continuation-passing style, defunctionalization, abstract machines

The Art, Science, and Engineering of Programming

Perspective The Art of Programming

Area of Submission General-purpose programming; program derivation



© Jeremy Gibbons
This work is licensed under a “CC BY 4.0” license.
Submitted to *The Art, Science, and Engineering of Programming*.

1 Introduction

In a seminal paper [27], Reynolds showed how to use *continuation-passing style* (CPS) followed by *defunctionalization* to take a direct-style recursive interpreter written in some (presumably well understood) *defining language* implementing some (typically less well understood) *defined language*, and transform it into an abstract machine for executing programs in the defined language. CPS ensures that the order of execution of the resulting machine is independent of that of the defining language, for example whether the defining language is lazy or eager; and defunctionalization ensures that the abstract machine is first-order even if the defining language was higher-order.

Danvy and colleagues have subsequently elaborated on Reynolds' approach in a long sequence of papers [13, 2, 1, 3, 4, 12], using it to reconstruct many published abstract machines and to construct some new ones. One can also see CPS and defunctionalization applied in contexts other than language interpreters and abstract machines—for example, Huet's Zipper [15] and McBride's Clowns and Jokers [21]. Moreover, they shed light on some other standard techniques, such as accumulating parameters [9], translation between right-to-left and left-to-right folds [8, §3.5], and fast reverse [16].

This much is fairly well known. However, behind all these applications of the two techniques is an appeal to associativity, which is not so well known. Perhaps it should not come as a surprise that CPS transformation—which after all is a matter of sequentializing code that would otherwise be tree-structured—is connected to associativity. Nevertheless, we feel that the associativity aspect deserves to be highlighted. The purpose of this paper is to explain and explore the connection between CPS, defunctionalization, and associativity.

The paper uses the dependently typed language Idris as a programming notation, for which a brief primer is given in Appendix A. The paper is mostly independent of any specifics of the language; in particular, the only real reliance on Idris's dependent types is for implementing a generalized form of composition in Section 5, and is not essential to the central argument. The source of the paper is a literate script, and the extracted code is provided as supplementary material.

2 Warm-up: Factorial

We start with a very simple example, in order to set out the main ideas in a familiar context—a straightforward direct-style recursive implementation of factorial:

$$\begin{aligned} \text{fact} &: \text{Integer} \rightarrow \text{Integer} \\ \text{fact } 0 &= 1 \\ \text{fact } n &= n \times \text{fact } (n - 1) \end{aligned}$$

We deliberately work through the details here, so that we can present the later examples more briskly.

2.1 Continuation-passing style

The first step is to convert the direct-style definition to *continuation-passing style*. This is achieved by introducing an additional argument k , an *accumulating parameter*, thereby *generalizing* the function $fact$ to another function $fact'_2$ specified by

$$\begin{aligned} fact'_2 &: Integer \rightarrow (Integer \rightarrow r) \rightarrow r \\ fact'_2 n k &= k (fact n) \quad \text{-- specification of } fact'_2 \end{aligned}$$

This really is a generalization, because $fact$ may be retrieved by setting $k = id$. Indeed, we define

$$\begin{aligned} fact_2 &: Integer \rightarrow Integer \\ fact_2 n &= fact'_2 n id \end{aligned}$$

and so $fact_2 = fact$. (As a convention throughout the paper, we use the same name but different indices for different versions of morally the same function. Sometimes the different versions really are extensionally equal; more generally, they may be related by data refinement.)

From the original definition of $fact$, and the specification of $fact'_2$ in terms of it, using simple unfold–fold transformations [10], it is now straightforward to synthesize a definition of $fact'_2$ that no longer depends on $fact$. For the base case, we have

$$\begin{aligned} &fact'_2 0 k \\ &= [[\text{specification of } fact'_2 \]] \\ &\quad k (fact 0) \\ &= [[\text{definition of } fact \]] \\ &\quad k 1 \end{aligned}$$

and for the inductive step, that is, when $n > 0$:

$$\begin{aligned} &fact'_2 n k \\ &= [[\text{specification of } fact'_2 \]] \\ &\quad k (fact n) \\ &= [[\text{definition of } fact, \text{ given } n > 0 \]] \\ &\quad k (n \times fact (n - 1)) \\ &= [[\text{composition} \]] \\ &\quad (k \cdot (n \times)) (fact (n - 1)) \\ &= [[\text{specification of } fact'_2 \]] \\ &\quad fact'_2 (n - 1) (k \cdot (n \times)) \end{aligned}$$

We therefore define:

$$\begin{aligned} fact'_2 &: Integer \rightarrow (Integer \rightarrow r) \rightarrow r \\ fact'_2 0 k &= k 1 \\ fact'_2 n k &= fact'_2 (n - 1) (k \cdot (n \times)) \end{aligned}$$

This is the continuation-passing style implementation. Clearly the continuation argument k is being used as an accumulating parameter: it starts off as id , accumulates

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity

more information (here, by function composition) as the computation proceeds, and yields the final result by some projection (here, by application to 1).

An equivalent way of writing the second clause of $fact'_2$ is

$$fact'_2 n k = fact'_2 (n - 1) (\lambda m \Rightarrow k (n \times m))$$

in which the right-hand side might be read “compute the factorial of $n - 1$; get a result, and call it m ; then return $k (n \times m)$ ”. This version is more idiomatically in continuation-passing style, and we use this style without comment for later examples.

2.2 Defunctionalization

Now $fact'_2$ is *tail-recursive*: when it recurses, the result of the recursive call is also the final result. But it is still higher-order: it manipulates continuations. However, in the context of its use as an auxiliary function to $fact_2$, we don't need the full generality of $fact'_2$: the continuations are always of type $Integer \rightarrow Integer$, so we could have given $fact'_2$ the more specific type

$$fact'_2 : Integer \rightarrow (Integer \rightarrow Integer) \rightarrow Integer$$

Moreover, the continuations aren't even arbitrary functions of type $Integer \rightarrow Integer$: they're specifically compositions of the form $(a \times) \cdot (b \times) \cdots (c \times)$. Having made this observation, we can choose a different representation of functions in this special form, and transform $fact'_2$ by *data refinement* to use this different representation instead. If we pick a first-order representation, we get a first-order implementation of $fact'_2$. In particular, we can represent a function of the form $(a \times) \cdot (b \times) \cdots (c \times)$ —that is, the composition of a possibly empty sequence of multiplications by various factors a, b, \dots, c —by the list $[a, b, \dots, c]$ of those factors. The *abstraction function* $factabs_3$ converts from the new ‘concrete’ representation back to the old ‘abstract’ one:

$$\begin{aligned} factabs_3 &: List Integer \rightarrow (Integer \rightarrow Integer) \\ factabs_3 [] &= id \\ factabs_3 (n :: k) &= (n \times) \cdot factabs_3 k \end{aligned}$$

—that is, $factabs_3 k m = foldr (\times) m k$. Data refinement then starts from the specification

$$fact'_3 n k = fact'_2 n (factabs_3 k) \quad \text{-- specification of } fact'_3$$

and leads to the new implementation

$$\begin{aligned} fact_3 &: Integer \rightarrow Integer \\ fact_3 n &= fact'_3 n [] \quad \textbf{where} \\ fact'_3 &: Integer \rightarrow List Integer \rightarrow Integer \\ fact'_3 0 k &= foldr (\times) 1 k \quad \text{-- = product } k \\ fact'_3 n k &= fact'_3 (n - 1) (k \# [n]) \end{aligned}$$

This is still tail-recursive, but is now also first-order.

2.3 Associativity

Of course, $fact_3$ is not the tail-recursive program for factorial that you would sit down and write from first principles. That program arises from a further data refinement, representing the composition $(a \times) \cdot (b \times) \cdot \dots \cdot (c \times)$ by the product $a \times b \times \dots \times c$ of factors—a single integer, rather than a list of integers. Now the abstraction function back to the original continuations used by $fact'_2$ is

$$\begin{aligned} factabs_4 &: Integer \rightarrow (Integer \rightarrow Integer) \\ factabs_4 k &= (k \times) \end{aligned}$$

Applying the data refinement starts from the specification

$$fact'_4 n k = fact'_2 n (factabs_4 k) \quad \text{-- specification of } fact'_4$$

and leads to the new implementation

$$\begin{aligned} fact_4 &: Integer \rightarrow Integer \\ fact_4 n &= fact'_4 n 1 \quad \textbf{where} \\ fact'_4 &: Integer \rightarrow Integer \rightarrow Integer \\ fact'_4 0 k &= k \\ fact'_4 n k &= fact'_4 (n - 1) (k \times n) \end{aligned}$$

This is still tail-recursive and first-order, but now also manipulates only scalar data. But note that this final data refinement is *valid only because of associativity of multiplication*. The crucial part of the development is the inductive step in the auxilliary function, in the case $n > 0$:

$$\begin{aligned} & fact'_4 n k \\ &= \llbracket \text{specification of } fact'_4 \rrbracket \\ & \quad fact'_2 n (factabs_4 k) \\ &= \llbracket \text{definition of } fact'_2, \text{ given that } n > 0 \rrbracket \\ & \quad fact'_2 (n - 1) (\lambda m \Rightarrow factabs_4 k (n \times m)) \\ &= \llbracket \text{definition of } factabs_4 \rrbracket \\ & \quad fact'_2 (n - 1) (\lambda m \Rightarrow k \times (n \times m)) \\ &= \llbracket \text{associativity of multiplication} \rrbracket \\ & \quad fact'_2 (n - 1) (\lambda m \Rightarrow (k \times n) \times m) \\ &= \llbracket \text{definition of } factabs_4 \rrbracket \\ & \quad fact'_2 (n - 1) (\lambda m \Rightarrow factabs_4 (k \times n) m) \\ &= \llbracket \text{specification of } fact'_4 \rrbracket \\ & \quad fact'_4 (n - 1) (k \times n) \end{aligned}$$

Suppose we had started instead with a non-associative operator, as in the ‘subtractorial’ function:

$$\begin{aligned} subt &: Integer \rightarrow Integer \\ subt 0 &= 1 \\ subt n &= n - subt (n - 1) \end{aligned}$$

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity

We could have followed the same steps to get a tail-recursive $subt_2$ and a first-order $subt_3$, but we would need to do something different in order to get a scalar $subt_4$. It's a nice exercise to work out just what that difference would be—but because subtraction is not associative, a different insight is needed than for factorial.

3 Reverse

Let's look at another familiar linear example: list reversal. This goes through similar steps, but the dependence on associativity arises in a different way. We start with the naive, quadratic-time, direct-style reverse function:

$$\begin{aligned} reverse &: List\ a \rightarrow List\ a \\ reverse\ [] &= [] \\ reverse\ (x :: xs) &= reverse\ xs ++ [x] \end{aligned}$$

3.1 Continuation-passing style

CPS conversion (with the more specific type for the auxiliary function, but this time making it a local definition) gives:

$$\begin{aligned} reverse_2 &: List\ a \rightarrow List\ a \\ reverse_2\ xs &= reverse'_2\ xs\ id\ \mathbf{where} \\ reverse'_2 &: List\ a \rightarrow (List\ a \rightarrow List\ a) \rightarrow List\ a \\ reverse'_2\ []\ k &= k\ [] \\ reverse'_2\ (x :: xs)\ k &= reverse'_2\ xs\ (\lambda zs \Rightarrow k\ (zs ++ [x])) \end{aligned}$$

3.2 Defunctionalization

The continuations are all functions of the form $(++[a]) \cdot (++[b]) \cdot \dots \cdot (++[c])$, compositions of functions that each append a single element. Defunctionalization picks a first-order encoding of such higher-order values; which encoding shall we pick? The most natural thing to try is to represent the composition as the list $[a, b, \dots, c]$ of appendees. But then the abstraction function, which we need in the base case of the auxiliary function, is list reversal again! That's no help.

3.3 Associativity

The next most natural thing is to represent the composition $(++[a]) \cdot (++[b]) \cdot \dots \cdot (++[c])$ as the list $[c, \dots, b, a]$ of appendees in reverse order—after all, the abstraction function then no longer needs to reverse anything:

$$\begin{aligned} revabs_3 &: List\ a \rightarrow (List\ a \rightarrow List\ a) \\ revabs_3\ ys &= \lambda zs \Rightarrow zs ++ ys \end{aligned}$$

Is this more helpful? Let's work through the development. We start with the specification

$$\text{reverse}'_3 \text{ xs } k = \text{reverse}'_2 \text{ xs } (\text{revabs}_3 \text{ } k) \quad \text{-- specification of } \text{reverse}'_3$$

and calculate for the base case of the auxiliary function:

$$\begin{aligned} & \text{reverse}'_3 \text{ [] } \text{ys} \\ = & \text{[[specification of } \text{reverse}'_3 \text{]]} \\ & \text{reverse}'_2 \text{ [] } (\text{revabs}_3 \text{ } \text{ys}) \\ = & \text{[[definition of } \text{reverse}'_2 \text{]]} \\ & \text{revabs}_3 \text{ } \text{ys []} \\ = & \text{[[definition of } \text{revabs}_3 \text{]]} \\ & \text{[] ++ ys} \\ = & \text{[[definition of ++]]} \\ & \text{ys} \end{aligned}$$

and for the inductive step:

$$\begin{aligned} & \text{reverse}'_3 (x :: \text{xs}) \text{ys} \\ = & \text{[[specification of } \text{reverse}'_3 \text{]]} \\ & \text{reverse}'_2 (x :: \text{xs}) (\text{revabs}_3 \text{ } \text{ys}) \\ = & \text{[[definition of } \text{reverse}'_2 \text{]]} \\ & \text{reverse}'_2 \text{ xs } (\lambda \text{zs} \Rightarrow \text{revabs}_3 \text{ } \text{ys } (\text{zs ++ [x]})) \\ = & \text{[[definition of } \text{revabs}_3 \text{]]} \\ & \text{reverse}'_2 \text{ xs } (\lambda \text{zs} \Rightarrow (\text{zs ++ [x]} ++ \text{ys})) \\ = & \text{[[associativity of ++, and definition]]} \\ & \text{reverse}'_2 \text{ xs } (\lambda \text{zs} \Rightarrow \text{zs ++ (x :: \text{ys}))} \\ = & \text{[[definition of } \text{revabs}_3 \text{]]} \\ & \text{reverse}'_2 \text{ xs } (\text{revabs}_3 (x :: \text{ys})) \\ = & \text{[[specification of } \text{reverse}'_3 \text{]]} \\ & \text{reverse}'_3 \text{ xs } (x :: \text{ys}) \end{aligned}$$

We end up with the well-known linear-time fast reverse function using an accumulator:

$$\begin{aligned} \text{reverse}_3 & : \text{List } a \rightarrow \text{List } a \\ \text{reverse}_3 \text{ xs} & = \text{reverse}'_3 \text{ xs []} \textbf{ where} \\ & \text{reverse}'_3 : \text{List } a \rightarrow \text{List } a \rightarrow \text{List } a \\ & \text{reverse}'_3 \text{ [] } k &= k \\ & \text{reverse}'_3 (x :: \text{xs}) k &= \text{reverse}'_3 \text{ xs } (x :: k) \end{aligned}$$

But we only get there by exploiting associativity of ++; without that, there is no productive defunctionalization of the CPS version of the program in the first place.

4 Flattening trees

Let's now turn our attention to non-linear data; in particular, binary trees:

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity

```
data Tree a = Tip a | Bin (Tree a) (Tree a)
```

Here is the obvious, recursive, direct-style function for flattening a tree to a list:

```
flatten : Tree a → List a
flatten (Tip x) = [x]
flatten (Bin t u) = flatten t ++ flatten u
```

Note that, like *reverse*, this is again quadratic-time, because of the nested applications of ++; and again, like *reverse*, there is a linear-time implementation obtained via an accumulating parameter. We might suspect that CPS and defunctionalization will again take us there. Let's see.

4.1 Continuation-passing style

Following our nose with CPS conversion gives:

```
flatten2 : Tree a → List a
flatten2 t = flatten'2 t id where
  flatten'2 : Tree a → (List a → List a) → List a
  flatten'2 (Tip x) k = k [x]
  flatten'2 (Bin t u) k = flatten'2 t (λxs ⇒ flatten'2 u (λys ⇒ k (xs ++ ys)))
```

Note that because binary trees are a non-linear datatype, there a choice to be made when using CPS to linearize the traversal: left-to-right, or right-to-left? This definition is left-to-right, visiting left child *t* before right child *u* (we will revisit this choice later).

4.2 Defunctionalization

It is perhaps no longer immediately obvious how to defunctionalize these continuations; what is the general form that they take? We need not wait for inspiration to strike: we can proceed methodically, following the path laid by Reynolds and explored by Danvy. There are three places where continuations are constructed, for the three call sites to *flatten'₂*:

- *id* -- no free variables
- $\lambda ys \Rightarrow k (xs ++ ys)$ -- free variables *xs, k*
- $\lambda xs \Rightarrow \text{flatten}'_2 u (\lambda ys \Rightarrow k (xs ++ ys))$ -- free variables *u, k*

So we will need a datatype with three different constructors. (Compare this with *fact₃*, which constructed continuations at the two call sites to the auxiliary function *fact'₃*, so used a datatype with two constructors, namely lists; and similarly for *reverse₂*.)

The last two of the three call sites to *flatten'₂* construct continuations with free variables, representing the context of the recursive call; the corresponding constructors will be *closures*, taking arguments to capture values for those free variables. We therefore introduce the following representation:


```

data FlatCont3 a = FlatRoot3
    | FlatLeftTree3 (Tree a) (FlatCont3 a)
    | FlatRightList3 (List a) (FlatCont3 a)

```

The idea is that $FlatRoot_3$ is the continuation used at the root of the tree, $FlatLeftTree_3$ for the recursive call on a left child (whose context includes the right sibling, another tree), and $FlatRightList_3$ for the recursive call on a right child (whose context includes the flattening of the left sibling, a list). The abstraction function is typed

$$flatabs_3 : FlatCont_3 a \rightarrow (List a \rightarrow List a)$$

and the corresponding auxilliary function is specified by

$$flatten'_3 t k = flatten'_2 t (flatabs_3 k) \quad \text{-- specification of } flatten'_3$$

Working through the data refinement yields the following implementation:

```

flatten3 : Tree a → List a
flatten3 t = flatten'3 t FlatRoot3 where
  mutual
    flatten'3 : Tree a → FlatCont3 a → List a
    flatten'3 (Tip x) k = flatabs3 k [x]
    flatten'3 (Bin t u) k = flatten'3 t (FlatLeftTree3 u k)
    flatabs3 : FlatCont3 a → (List a → List a)
    flatabs3 FlatRoot3 = id
    flatabs3 (FlatLeftTree3 u k) = λxs ⇒ flatten'3 u (FlatRightList3 xs k)
    flatabs3 (FlatRightList3 xs k) = λys ⇒ flatabs3 k (xs ++ ys)

```

Intuitively, the defunctionalized continuation is a stack of ‘context frames’, each of which is either a tree (a right sibling, the additional context when descending into a left child) or a list (the flattening of a left sibling, the additional context when descending into a right child).

4.3 Defunctionalization again

By happy accident, we need not have defined a new algebraic datatype for the defunctionalized continuations, since $FlatCont_3 a$ happens to be isomorphic to a sequence of choices:

```

FlatCont4 : Type → Type
FlatCont4 a = List (Either (Tree a) (List a))

```

(here, the sum type $Either X Y$ has values of the form $Left x$ for $x : X$ and $Right y$ for $y : Y$). Using this representation instead gives the equivalent program:

```

flatten4 : Tree a → List a
flatten4 t = flatten'4 t [ ] where
  mutual

```

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity

$$\begin{aligned}
 \text{flatten}'_4 &: \text{Tree } a \rightarrow \text{FlatCont}_4 a \rightarrow \text{List } a \\
 \text{flatten}'_4 (\text{Tip } x) k &= \text{flatabs}_4 k [x] \\
 \text{flatten}'_4 (\text{Bin } t u) k &= \text{flatten}'_4 t (\text{Left } u :: k) \\
 \text{flatabs}_4 &: \text{FlatCont}_4 a \rightarrow (\text{List } a \rightarrow \text{List } a) \\
 \text{flatabs}_4 [] &= \text{id} \\
 \text{flatabs}_4 (\text{Left } u :: k) &= \lambda xs \Rightarrow \text{flatten}'_4 u (\text{Right } xs :: k) \\
 \text{flatabs}_4 (\text{Right } xs :: k) &= \lambda ys \Rightarrow \text{flatabs}_4 k (xs ++ ys)
 \end{aligned}$$

With built-in lists, it is easier to see that the defunctionalized continuation behaves as a stack. However, the naming convention is a bit awkward with the built-in sum type: a right child u is in the *Left* side of the sum, because it is the context for descending into its left sibling.

Nevertheless, this is not the linear-time flattening function we were hoping for: it's still quadratic, because every flattening xs of a left sibling is retraversed to append the flattening ys of its right sibling (in the *Right* case of flatabs_4). Where did we go wrong?

4.4 Continuation-passing style again

As noted above, there was actually an arbitrary decision made earlier, not present in the direct-style program flatten and forced upon us by the CPS conversion to flatten_2 . The direct-style program is agnostic about whether left children are visited before or after right children; that is determined by the evaluation semantics of the defining language, and the definition of $++$. The CPS program makes explicit that the left child is visited before the right. Making this order of evaluation explicit is precisely why Reynolds used CPS for his definitional interpreters; and 'left before right' turns out to be precisely the wrong thing to do for tree flattening. Here is the other CPS version, visiting the right child before the left (but still delivering the tips in left to right order):


$$\begin{aligned}
 \text{flatten}_5 &: \text{Tree } a \rightarrow \text{List } a \\
 \text{flatten}_5 t &= \text{flatten}'_5 t \text{ id } \mathbf{where} \\
 \text{flatten}'_5 &: \text{Tree } a \rightarrow (\text{List } a \rightarrow \text{List } a) \rightarrow \text{List } a \\
 \text{flatten}'_5 (\text{Tip } x) k &= k [x] \\
 \text{flatten}'_5 (\text{Bin } t u) k &= \text{flatten}'_5 u (\lambda ys \Rightarrow \text{flatten}'_5 t (\lambda xs \Rightarrow k (xs ++ ys)))
 \end{aligned}$$

These continuations defunctionalize in the mirror image of what happened earlier: the additional context for a right child u is its left sibling t , and the context for a left child t is the flattening ys of its right sibling. So we introduce a new representation of defunctionalized continuations:

$$\begin{aligned}
 \text{FlatCont}_6 &: \text{Type} \rightarrow \text{Type} \\
 \text{FlatCont}_6 a &= \text{List } (\text{Either } (\text{List } a) (\text{Tree } a))
 \end{aligned}$$

Using this representation instead gives the following defunctionalized right-to-left traversal:

$$\begin{aligned}
 \text{flatten}_6 &: \text{Tree } a \rightarrow \text{List } a \\
 \text{flatten}_6 t &= \text{flatten}'_6 t [] \mathbf{where}
 \end{aligned}$$



■ **Figure 1** A tree to be flattened, with subtree t_2 highlighted

mutual

$$\begin{aligned}
 \text{flatten}'_6 &: \text{Tree } a \rightarrow \text{FlatCont}_6 a \rightarrow \text{List } a \\
 \text{flatten}'_6 (\text{Tip } x) k &= \text{flatabs}_6 k [x] \\
 \text{flatten}'_6 (\text{Bin } t \ u) k &= \text{flatten}'_6 u (\text{Right } t :: k) \\
 \text{flatabs}_6 &: \text{FlatCont}_6 a \rightarrow (\text{List } a \rightarrow \text{List } a) \\
 \text{flatabs}_6 [] &= \text{id} \\
 \text{flatabs}_6 (\text{Left } ys :: k) &= \lambda xs \Rightarrow \text{flatabs}_6 k (xs ++ ys) \\
 \text{flatabs}_6 (\text{Right } t :: k) &= \lambda ys \Rightarrow \text{flatten}'_6 t (\text{Left } ys :: k)
 \end{aligned}$$

Sadly, this is still quadratic-time: in the *Left* case of flatabs_6 , the constructed flattening xs is retraversed to append ys . More work is required to get linear time. And not surprisingly, that additional work involves associativity.

4.5 Associativity

Consider applying flatten_6 to the tree shown in Figure 1, via

$$xs = \text{flatten}_6 (\text{Bin } (\text{Bin } t_1 (\text{Bin } t_2 t_3)) t_4)$$

where t_1, t_2, t_3, t_4 are the four subtrees; and let $xs_i = \text{flatten } t_i$ for $i = 1, 2, 3, 4$. The figure highlights the left–right–left path leading from the root to subtree t_2 . The context when visiting subtree t_2 is the continuation $[\text{Left } xs_3, \text{Right } t_1, \text{Left } xs_4]$; that is,

$$xs = \text{flatten}'_6 t_2 [\text{Left } xs_3, \text{Right } t_1, \text{Left } xs_4]$$

Having obtained $xs_2 = \text{flatten } t_2$, the continuation is applied, leading eventually to

$$xs = (\text{flatten } t_1 ++ (xs_2 ++ xs_3)) ++ xs_4$$

This reveals that the particular interleaving of *Left* and *Right* frames in the continuation is irrelevant: *Left* frames xs_3, xs_4 get appended, *Right* frames t_1 get prepended, and—by associativity of $++$!—these two operations commute. So we could just as well separate the *Left* frames from the *Right* frames:

$$\text{flatabs}_6 k = \text{flatabs}_6 (\text{map Left } (\text{lefts } k) ++ \text{map Right } (\text{rights } k))$$

where $\text{lefts} : \text{List } (\text{Either } a \ b) \rightarrow \text{List } a$ and $\text{rights} : \text{List } (\text{Either } a \ b) \rightarrow \text{List } b$ extract the *Left*- and *Right*-tagged elements respectively of a list of *Eithers*. This in turn suggests another data refinement: we can discard the interleaving information, and

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity

represent the defunctionalized continuation instead as an uninterleaved pair of lists, ($List (List a), List (Tree a)$), consisting respectively of the data to be appended and that to be prepended.

But we can (and should) go further: all the *Left*-tagged lists are eventually going to be concatenated, so we might as well concatenate them into a single list as we go. That is,

$$flatabs_6 k = flatabs_6 ([Left (concat (lefts k))] ++ map Right (rights k))$$

This suggests yet another data refinement, to represent the defunctionalized continuation as a pair consisting of a single list of elements (to be appended) and a list of trees (whose flattenings will be prepended). We therefore introduce the new representation of defunctionalized continuations

$$\begin{aligned} FlatCont_7 &: Type \rightarrow Type \\ FlatCont_7 a &= (List a, List (Tree a)) \end{aligned}$$

with abstraction function specified by

$$flatabs_7 (ys, ts) xs = concat (map flatten (reverse ts)) ++ xs ++ ys$$

-- specification of $flatabs_7$

Applying this data refinement leads to the following program:

$$\begin{aligned} flatten_7 &: Tree a \rightarrow List a \\ flatten_7 t &= flatten'_7 t ([], []) \textbf{ where} \\ &\textbf{mutual} \\ &flatten'_7 : Tree a \rightarrow FlatCont_7 a \rightarrow List a \\ &flatten'_7 (Tip x) (ys, ts) = flatabs_7 (ys, ts) [x] \\ &flatten'_7 (Bin t u) (ys, ts) = flatten'_7 u (ys, t :: ts) \\ &flatabs_7 : FlatCont_7 a \rightarrow (List a \rightarrow List a) \\ &flatabs_7 (ys, []) = \lambda xs \Rightarrow xs ++ ys \\ &flatabs_7 (ys, t :: ts) = \lambda xs \Rightarrow flatten'_7 t (xs ++ ys, ts) \end{aligned}$$

This still contains two occurrences of $++$, in the two clauses for $flatabs_7$. But note that $flatabs_7$ is only applied in a single place, and then only to a singleton list; so those $++$ s take constant time. This program therefore takes linear time overall; and the transformations that brought us there depended crucially on associativity of $++$.

This is still not quite the program you might have written from first principles, but it is only one final data refinement away. Instead of the call $flatten'_7 t (ys, ts)$ with an isolated ‘tree in focus’ t and a possibly empty stack ts of postponed trees, we push t onto ts to make a non-empty stack $t :: ts$ of trees:

$$\begin{aligned} flatten_8 &: Tree a \rightarrow List a \\ flatten_8 t &= flatten'_8 ([], [t]) \textbf{ where} \\ &\textbf{mutual} \\ &flatten'_8 : (List a, List (Tree a)) \rightarrow List a \end{aligned}$$

$$\begin{aligned}
\text{flatten}'_8 (ys, \text{Tip } x :: ts) &= \text{flatabs}_8 (ys, ts) [x] \\
\text{flatten}'_8 (ys, \text{Bin } t \ u :: ts) &= \text{flatten}'_8 (ys, u :: t :: ts) \\
\text{flatabs}_8 : (\text{List } a, \text{List } (\text{Tree } a)) &\rightarrow (\text{List } a \rightarrow \text{List } a) \\
\text{flatabs}_8 (ys, []) &= \lambda xs \Rightarrow xs ++ ys \\
\text{flatabs}_8 (ys, t :: ts) &= \lambda xs \Rightarrow \text{flatten}'_8 (xs ++ ys, t :: ts)
\end{aligned}$$

There is only one call site for flatabs_8 ; if we make a case analysis there on whether ts is empty, we can inline flatabs_8 and eliminate the additional definition. We also take the opportunity to curry the auxiliary function, and to flip its arguments:

$$\begin{aligned}
\text{flatten}_9 : \text{Tree } a &\rightarrow \text{List } a \\
\text{flatten}_9 t &= \text{flatten}'_9 [t] [] \textbf{ where} \\
\text{flatten}'_9 : \text{List } (\text{Tree } a) &\rightarrow \text{List } a \rightarrow \text{List } a \\
\text{flatten}'_9 [\text{Tip } x] \quad ys &= x :: ys \\
\text{flatten}'_9 (\text{Tip } x :: ts) \quad ys &= \text{flatten}'_9 ts (x :: ys) \\
\text{flatten}'_9 (\text{Bin } t \ u :: ts) \quad ys &= \text{flatten}'_9 (u :: t :: ts) ys
\end{aligned}$$

This program operates on a non-empty stack of trees, maintaining an accumulating parameter for the flattening so far, which is constructed from right to left. It is tail-recursive, first-order, and linear-time. This is the program you would write from first principles: a model application of an accumulating parameter to tree flattening.

5 Interpreters and compilers


Reynolds' original motivation [27] for investigating CPS and defunctionalization was to elucidate the derivation of abstract machines from language interpreters. Associativity also shows up there—in a way not exploited by Reynolds himself, nor by the extensive subsequent elaboration by Danvy and colleagues [13, 2, 1, 3, 4, 12], but it is central to the approach taken by Wand [29, 30]. However, in contrast to our earlier examples using factorial, reverse, and flattening, associativity does not arise from the problem domain—that is, the operators of the language being implemented—but from (a generalized form of) sequential composition of subprograms. Our point here is to remind ourselves of Wand's approach, which seems to have become lost over the intervening four decades.

For these purposes, it suffices to take an extremely simple language—one of arithmetic expressions constructed from numeric literals and addition. Hutton [17] shows that this same language serves as a basis for discussing many aspects of language semantics and implementation. But because we are interested specifically in associativity, we diverge from Hutton's language to one of literals and (non-associative) *subtraction*; we do not want to confuse associativity of sequential composition in the machine implementation with associativity of addition in the source language.

data $\text{Expr} = \text{Lit Integer} \mid \text{Diff Expr Expr}$

Here is the straightforward, direct-style, recursive evaluator for this language:

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity



■ **Figure 2** The expression $\text{Diff} (\text{Diff} (\text{Lit } 3) (\text{Lit } 4)) (\text{Lit } 5)$

$$\begin{aligned}
 \text{eval} &: \text{Expr} \rightarrow \text{Integer} \\
 \text{eval} (\text{Lit } n) &= n \\
 \text{eval} (\text{Diff } e \ e') &= \text{eval } e - \text{eval } e'
 \end{aligned}$$

so with

$$\text{expr} = \text{Diff} (\text{Diff} (\text{Lit } 3) (\text{Lit } 4)) (\text{Lit } 5)$$

as shown in Figure 2 we have $\text{eval } \text{expr} = (3 - 4) - 5 = -6$.

5.1 CPS and defunctionalization

If we follow the same process as for tree flattening, we obtain a tail-recursive evaluator via CPS:

$$\begin{aligned}
 \text{eval}_2 &: \text{Expr} \rightarrow \text{Integer} \\
 \text{eval}_2 \ e &= \text{eval}'_2 \ e \ \text{id} \ \mathbf{where} \\
 \text{eval}'_2 &: \text{Expr} \rightarrow (\text{Integer} \rightarrow \text{Integer}) \rightarrow \text{Integer} \\
 \text{eval}'_2 (\text{Lit } n) \ k &= k \ n \\
 \text{eval}'_2 (\text{Diff } e \ e') \ k &= \text{eval}'_2 \ e \ (\lambda m \Rightarrow \text{eval}'_2 \ e' \ (\lambda n \Rightarrow k \ (m - n)))
 \end{aligned}$$

and a tail-recursive, first-order evaluator via defunctionalization:

$$\begin{aligned}
 \mathbf{data} \ \text{EvalFrame}_3 &= \text{EvalLeftExpr}_3 \ \text{Expr} \mid \text{EvalRightValue}_3 \ \text{Integer} \\
 \text{EvalCont}_3 &: \text{Type} \\
 \text{EvalCont}_3 &= \text{List } \text{EvalFrame}_3 \\
 \text{eval}_3 &: \text{Expr} \rightarrow \text{Integer} \\
 \text{eval}_3 \ e &= \text{eval}'_3 \ e \ [] \ \mathbf{where} \\
 \mathbf{mutual} & \\
 \text{eval}'_3 &: \text{Expr} \rightarrow \text{EvalCont}_3 \rightarrow \text{Integer} \\
 \text{eval}'_3 (\text{Lit } n) \ k &= \text{evalabs}_3 \ k \ n \\
 \text{eval}'_3 (\text{Diff } e \ e') \ k &= \text{eval}'_3 \ e \ (\text{EvalLeftExpr}_3 \ e' :: k) \\
 \text{evalabs}_3 &: \text{EvalCont}_3 \rightarrow (\text{Integer} \rightarrow \text{Integer}) \\
 \text{evalabs}_3 [] & \quad \quad \quad n = n \\
 \text{evalabs}_3 (\text{EvalLeftExpr}_3 \ e' :: k) & \quad m = \text{eval}'_3 \ e' \ (\text{EvalRightValue}_3 \ m :: k) \\
 \text{evalabs}_3 (\text{EvalRightValue}_3 \ m :: k) & \quad n = \text{evalabs}_3 \ k \ (m - n)
 \end{aligned}$$

The defunctionalized continuations $EvalCont_3$ are stacks of frames, with either unevaluated expressions (the right siblings of left visitees) or integers (the evaluations of left siblings of right visitees); we have defined a new datatype $EvalFrame_3$ to allow more specific names than with plain $Either\ Expr\ Integer$.

Now $eval_3$ is readily seen as an abstract machine for evaluating expressions: evaluating a $Diff$ pushes the right child on the stack then moves down to the left; when eventually a result is obtained, the next task is popped off the stack; if this result was from a left child, its right sibling is popped off the stack to visit next and the result pushed in its place; and if the result was from a right child, the corresponding left result is popped from the stack, the difference computed, and the remainder of the stack used; when the stack is empty, the result obtained is the final result.

5.2 Compilation

However, $eval_3$ cannot be seen as a *compiler* for expressions: the abstract machine manipulates unevaluated expressions ‘at run-time’ on the stack, whereas a compiler should have eliminated these. Nevertheless, there is a well-known stack-based compilation scheme for such expression languages. There are instructions to push a value onto the stack, and to replace the top two elements of the stack with their difference:

data $Instr = PushI\ Integer\ |\ SubI$

$Prog_4 : Type$

$Prog_4 = List\ Instr$

Program execution is a left fold over the instruction sequence, inducing a stack transformation (a function from stacks to stacks):

$exec_4 : Prog_4 \rightarrow List\ Integer \rightarrow List\ Integer$

$exec_4\ p\ s = foldl\ step\ s\ p$ **where**

$step\ ns \quad (PushI\ n) = n :: ns$

$step\ (n :: m :: ns)\ SubI \quad = (m - n) :: ns \quad \text{-- note reversal of arguments}$

Compilation is a traversal of the expression tree:

$compile_4 : Expr \rightarrow Prog_4$

$compile_4\ (Lit\ n) \quad = [PushI\ n]$

$compile_4\ (Diff\ e\ e') = compile_4\ e\ ++\ compile_4\ e'\ ++\ [SubI]$


For example,

$compile_4\ expr = [PushI\ 3, PushI\ 4, SubI, PushI\ 5, SubI]$

In the special case of a program of the form $compile_4\ e$, execution consists of consuming an initial stack (which is not inspected) and producing a final stack with the evaluation of e pushed on top:

$exec_4\ (compile_4\ e)\ s = eval\ e :: s$

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity



■ **Figure 3** Generalized composition $b^r gf$

In particular, an expression may be evaluated by compiling and executing a program on the initially empty stack, then extracting the sole element of the final stack:

$$\begin{aligned} eval_4 &: Expr \rightarrow Integer \\ eval_4 e &= \mathbf{case} \mathit{exec}_4 (\mathit{compile}_4 e) [] \mathbf{of} [n] \Rightarrow n \end{aligned}$$

This compilation scheme entails an abstract machine, and it evidently involves stacks; so it seems like it ought to be obtainable using CPS and defunctionalization. Can we get there following the same process as with the other examples? We can—but only if we exploit associativity.

5.3 Generalized composition

Wand [30] observes that the essence of the key transformation is associativity of a generalization of function composition. Consider the recursive case of the CPS interpreter:

$$eval'_2 (\mathit{Diff} e e') k = eval'_2 e (\lambda m \Rightarrow eval'_2 e' (\lambda n \Rightarrow k (m - n)))$$

The right-hand side can be seen as routing one argument k from the surrounding context into the continuation for the left child e , and routing two arguments k, m into the continuation for the right child e' . To arrange this routing, Wand [30] therefore introduces the generalized composition operation b^r , such that

$$b^r gf x_1 \dots x_r = g (f x_1 \dots x_r)$$

as illustrated in Figure 3 (although Wand writes B_r where we write b^r). That is,

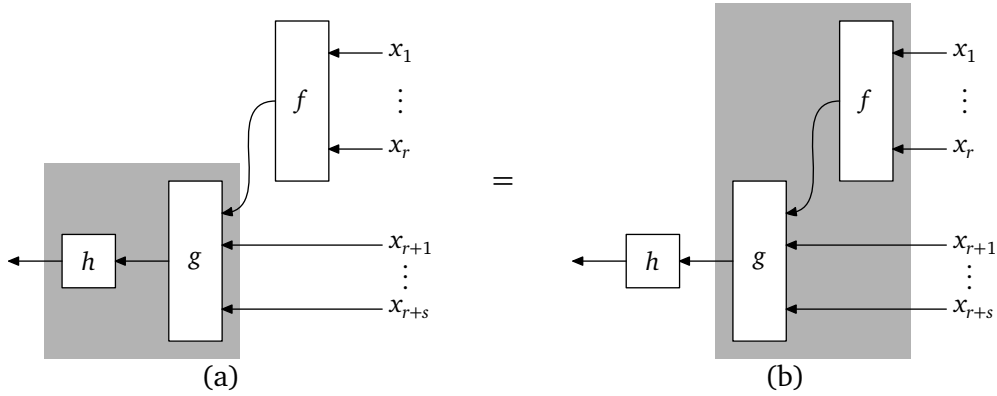
$$\begin{aligned} b^0 gf &= gf \\ b^{r+1} gf &= (b^r g) \cdot f \end{aligned}$$

Equivalently,

$$\begin{aligned} b^0 &= id \\ b^{r+1} &= (b^r g) \cdot (\cdot) \end{aligned}$$

so that $b^1 = (\cdot)$, $b^2 = (\cdot) \cdot (\cdot)$, $b^3 = (\cdot) \cdot (\cdot) \cdot (\cdot)$ (respectively Smullyan's 'bluebird', 'blackbird', and 'bunting' combinators [28]), and in general $b^r = (\cdot) \cdot \dots \cdot (\cdot)$, the composition of r instances of (\cdot) .

Using generalized composition, we have:



■ **Figure 4** Associativity of generalized composition: (a) $b^r (b^{s+1} h g) f$, (b) $b^{r+s} h (b^r g f)$

$$\begin{aligned}
 & \text{eval}'_2 (\text{Diff } e \ e') \\
 = & \text{[[definition of } \text{eval}'_2 \text{]]} \\
 & \lambda k \Rightarrow \text{eval}'_2 e (\lambda m \Rightarrow \text{eval}'_2 e' (\lambda n \Rightarrow k (m - n))) \\
 = & \text{[[since } \lambda k \Rightarrow g (f k) \text{ is } b^1 g (\lambda k \Rightarrow f k) \text{]]} \\
 & b^1 (\text{eval}'_2 e) (\lambda k m \Rightarrow \text{eval}'_2 e' (\lambda n \Rightarrow k (m - n))) \\
 = & \text{[[since } \lambda k m \Rightarrow g (f k m) \text{ is } b^2 g (\lambda k m \Rightarrow f k m) \text{]]} \\
 & b^1 (\text{eval}'_2 e) (b^2 (\text{eval}'_2 e') (\lambda k m n \Rightarrow k (m - n)))
 \end{aligned}$$

Moreover, generalized composition is—of course!—(pseudo-)associative:

$$\begin{aligned}
 & (b^r (b^{s+1} h g) f) x_1 \dots x_{r+s} \\
 = & (b^{s+1} h g) (f x_1 \dots x_r) x_{r+1} \dots x_{r+s} \\
 = & h (g (f x_1 \dots x_r) x_{r+1} \dots x_{r+s}) \\
 = & h ((b^r g f) x_1 \dots x_{r+s}) \\
 = & (b^{r+s} h (b^r g f)) x_1 \dots x_{r+s}
 \end{aligned}$$

That is,

$$b^r (b^{s+1} h g) f = b^{r+s} h (b^r g f)$$

as illustrated in Figure 4. We call this ‘pseudo-associativity’ rather than plain ‘associativity’ because the indices change under rebracketing. It is this associativity that is the crucial transformation allowing us to rotate tree-shaped code to obtain linear code, thereby eliminating expressions from the run-time artifacts, as we shall see.

5.4 Implementing generalized composition

The first few arities of generalized composition can be defined as follows:

$$\begin{aligned}
 b^0 & : (b \rightarrow c) \rightarrow b \rightarrow c \\
 b^0 g f & = g f \\
 b^1 & : (b \rightarrow c) \rightarrow (a \rightarrow b) \rightarrow (a \rightarrow c) \\
 b^1 g f & = \lambda x \Rightarrow g (f x)
 \end{aligned}$$

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity

$$b^2 : (b \rightarrow c) \rightarrow (a \rightarrow a' \rightarrow b) \rightarrow (a \rightarrow a' \rightarrow c)$$

$$b^2 gf = \lambda xy \Rightarrow g (f xy)$$

However, the general case has a dependent type, depending on the arity of f . In order to accommodate different types for the different argument positions, we index the definition by a *list of types*. Thus, we introduce the type function *Arrow as b* for list of types *as* and type *b*:

$$\text{Arrow} : \text{List Type} \rightarrow \text{Type} \rightarrow \text{Type}$$

$$\text{Arrow} [] \quad b = b$$

$$\text{Arrow} (a :: as) b = a \rightarrow \text{Arrow as } b$$

For example

$$\text{Arrow} [\text{Char}, \text{Bool}] \text{String} = \text{Char} \rightarrow \text{Bool} \rightarrow \text{String}$$

at arity 2. Then we can define generalized composition, by induction over the arity:

$$b : \{as : \text{List Type}\} \rightarrow (b \rightarrow c) \rightarrow \text{Arrow as } b \rightarrow \text{Arrow as } c$$

$$b \{as = []\} \quad gf = gf$$

$$b \{as = _ :: _\} gf = b g \cdot f$$

Rewriting $eval'_2$ to use generalized composition yields a different implementation of the evaluator:

$$eval_5 : \text{Expr} \rightarrow \text{Integer}$$

$$eval_5 e = eval'_5 e \text{ halt } \mathbf{where}$$

$$eval'_5 : \text{Expr} \rightarrow (\text{Integer} \rightarrow \text{Integer}) \rightarrow \text{Integer}$$

$$eval'_5 (\text{Lit } n) \quad = \text{ret } n$$

$$eval'_5 (\text{Diff } e e') = b^1 (eval'_5 e) (b^2 (eval'_5 e') \text{ sub})$$

where for later convenience we have introduced three abbreviations:

$$\text{halt} = \text{id}$$

$$\text{ret } n = \lambda k \Rightarrow k n$$

$$\text{sub} = \lambda k \Rightarrow \lambda m \Rightarrow \lambda n \Rightarrow k (m - n)$$

5.5 Tree-shaped code

Note that $eval'_5$ is not tail-recursive any more, because of the imposition of the generalized compositions. Nevertheless, it suggests a data representation:

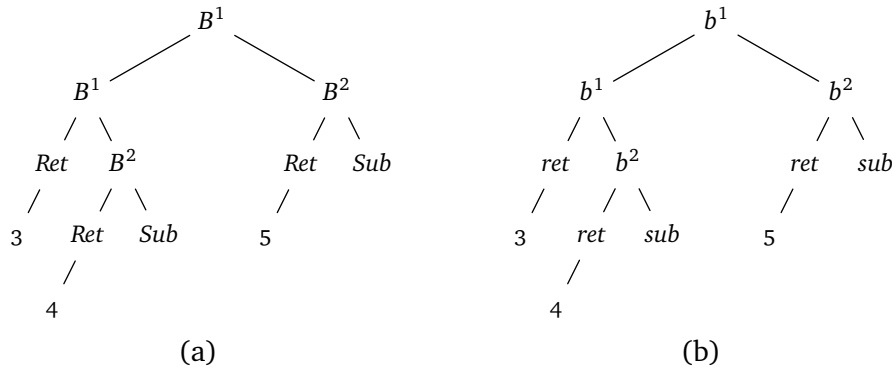
$$\mathbf{data} \text{ExprRep}_6 : \text{List Type} \rightarrow \text{Type} \mathbf{where}$$

$$\text{Ret}_6 : \text{Integer} \rightarrow \text{ExprRep}_6 []$$

$$\text{Sub}_6 : \text{ExprRep}_6 [\text{Integer}, \text{Integer}]$$

$$B_6^1 : \text{ExprRep}_6 [] \rightarrow \text{ExprRep}_6 [\text{Integer}] \rightarrow \text{ExprRep}_6 []$$

$$B_6^2 : \text{ExprRep}_6 [] \rightarrow \text{ExprRep}_6 [\text{Integer}, \text{Integer}] \rightarrow \text{ExprRep}_6 [\text{Integer}]$$



■ **Figure 5** (a) Tree-shaped code and (b) its interpretation

A data structure of type $\text{ExprRep}_6 r$ represents an evaluation function of type $(\text{Integer} \rightarrow \text{Integer}) \rightarrow \text{Arrow } r \text{ Integer}$. The first argument is the finalizing continuation. The index r specifies how many additional integer values are required to complete the evaluation to an integer; for example, $\text{Ret}_6 n :: \text{ExprRep}_6 []$ will yield a result with no additional values needed, whereas $\text{Sub}_6 :: \text{ExprRep}_6 [\text{Integer}, \text{Integer}]$ requires two additional integer values.

We can translate expressions into this representation:

$$\begin{aligned} \text{rep}_6 &: \text{Expr} \rightarrow \text{ExprRep}_6 [] \\ \text{rep}_6 (\text{Lit } n) &= \text{Ret}_6 n \\ \text{rep}_6 (\text{Diff } e \ e') &= B_6^1 (\text{rep}_6 e) (B_6^2 (\text{rep}_6 e') \text{Sub}_6) \end{aligned}$$

which gives a tree as shown in Figure 5(a). We can then evaluate the representative:

$$\begin{aligned} \text{abs}_6 &: \text{ExprRep}_6 r \rightarrow (\text{Integer} \rightarrow \text{Integer}) \rightarrow \text{Arrow } r \text{ Integer} \\ \text{abs}_6 (\text{Ret}_6 n) &= \text{ret } n \\ \text{abs}_6 \text{Sub}_6 &= \text{sub} \\ \text{abs}_6 (B_6^1 x y) &= b^1 (\text{abs}_6 x) (\text{abs}_6 y) \\ \text{abs}_6 (B_6^2 x y) &= b^2 (\text{abs}_6 x) (\text{abs}_6 y) \end{aligned}$$

yielding the interpretation shown in Figure 5(b). The two steps constitute another evaluator:

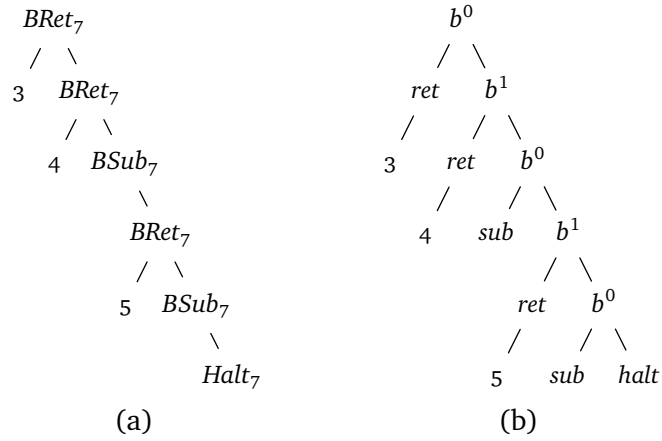
$$\begin{aligned} \text{eval}_6 &: \text{Expr} \rightarrow \text{Integer} \\ \text{eval}_6 e &= \text{abs}_6 (\text{rep}_6 e) \text{halt} \end{aligned}$$

which evidently deforests to eval_5 . But this is still not promising as a compiler, because the intermediate representation ExprRep_6 is still tree-shaped, as Figure 5 illustrates, whereas we were hoping for a linear sequence of instructions.

5.6 Linear code

Fortunately, associativity of generalized composition justifies rotation of the tree-shaped code in Figure 5 into linear form, as shown in Figure 6. Observe that

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity



■ **Figure 6** (a) Linear code and (b) its interpretation

$$\begin{aligned}
 & eval_5 e \\
 = & \llbracket \text{definition} \rrbracket \\
 & eval'_5 e \text{ halt} \\
 = & \llbracket b^0 \text{ is simply application} \rrbracket \\
 & b^0 (eval'_5 e) \text{ halt}
 \end{aligned}$$

Then we can unroll the definition of $eval'_5$ on $expr$, as in Figure 5(b), and exploit pseudo-associativity of the generalized compositions:

$$\begin{aligned}
 & eval_5 expr \\
 = & \llbracket \text{above} \rrbracket \\
 & b^0 (eval'_5 expr) \text{ halt} \\
 = & \llbracket \text{definition of } eval'_5 \rrbracket \\
 & b^0 (b^1 (b^1 (ret\ 3) (b^2 (ret\ 4) sub)) (b^2 (ret\ 5) sub)) \text{ halt} \\
 = & \llbracket \text{pseudo-associativity: } b^0 (b^1 h g) f = b^0 h (b^0 g f) \rrbracket \\
 & b^0 (b^1 (ret\ 3) (b^2 (ret\ 4) sub)) (b^0 (b^2 (ret\ 5) sub) \text{ halt}) \\
 = & \llbracket \text{pseudo-associativity: } b^0 (b^1 h g) f = b^0 h (b^0 g f) \rrbracket \\
 & b^0 (ret\ 3) (b^0 (b^2 (ret\ 4) sub) (b^0 (b^2 (ret\ 5) sub) \text{ halt})) \\
 = & \llbracket \text{pseudo-associativity: } b^0 (b^2 h g) f = b^1 h (b^0 g f) \rrbracket \\
 & b^0 (ret\ 3) (b^1 (ret\ 4) (b^0 sub (b^0 (b^2 (ret\ 5) sub) \text{ halt}))) \\
 = & \llbracket \text{pseudo-associativity: } b^0 (b^2 h g) f = b^1 h (b^0 g f) \rrbracket \\
 & b^0 (ret\ 3) (b^1 (ret\ 4) (b^0 sub (b^1 (ret\ 5) (b^0 sub \text{ halt}))))
 \end{aligned}$$

The end result is shown in Figure 6(b); note in particular that the generalized compositions now all nest to the right. In the same way as this specific example, pseudo-associativity allows us to rotate any tree-shaped evaluation into a linear form, in which no composition has another composition in its left argument. We can capture and enforce this property in another data representation:

```

data ExprRep7 : List Type → Type where
  Halt7 : ExprRep7 [Integer]
  BRet7 : Integer → ExprRep7 (Integer :: r) → ExprRep7 r
  BSub7 : ExprRep7 (Integer :: r) → ExprRep7 (Integer :: Integer :: r)
  
```

Now a data structure of type $ExprRep_7$ r represents an evaluation function of type $Arrow$ r $Integer$, and as before the index r denotes the number of additional integer values required to complete the evaluation. But the datatype allows only a single ('right') child per non-leaf node. Note that in the process of rotating, the arities can get arbitrarily high—in particular, for right-nested expressions of the form $4 - (3 - (2 - (1 - 1)))$ as in the 'subtractorial' function *subt* from Section 2. The representation must therefore allow arbitrary indices, in contrast to $ExprRep_6$ which used indices only up to arity 2.

Given a function to append two representations:

$$\begin{aligned} append_7 &: ExprRep_7\ r \rightarrow ExprRep_7\ (Integer :: s) \rightarrow ExprRep_7\ (r \# s) \\ append_7\ Halt_7\ y &= y \\ append_7\ (BRet_7\ n\ k)\ y &= BRet_7\ n\ (append_7\ k\ y) \\ append_7\ (BSub_7\ k)\ y &= BSub_7\ (append_7\ k\ y) \end{aligned}$$

we can transform from $ExprRep_6$ to $ExprRep_7$ by rotating the branching constructors B_6^1, B_6^2 :

$$\begin{aligned} rotate_7 &: ExprRep_6\ r \rightarrow ExprRep_7\ r \\ rotate_7\ (Ret_6\ n) &= BRet_7\ n\ Halt_7 \\ rotate_7\ Sub_6 &= BSub_7\ Halt_7 \\ rotate_7\ (B_6^1\ x\ y) &= append_7\ (rotate_7\ x)\ (rotate_7\ y) \\ rotate_7\ (B_6^2\ x\ y) &= append_7\ (rotate_7\ x)\ (rotate_7\ y) \end{aligned}$$

But instead of going via the tree-shaped representation, we can convert expressions directly into the linear representation:

$$\begin{aligned} rep_7 &: Expr \rightarrow ExprRep_7\ [] \\ rep_7\ (Lit\ n) &= BRet_7\ n\ Halt_7 \\ rep_7\ (Diff\ e\ e') &= append_7\ (rep_7\ e)\ (append_7\ (rep_7\ e')\ (BSub_7\ Halt_7)) \end{aligned}$$

Either way, we can then interpret $ExprRep_7$ representations:

$$\begin{aligned} abs_7 &: ExprRep_7\ r \rightarrow Arrow\ r\ Integer \\ abs_7\ Halt_7 &= halt \\ abs_7\ (BRet_7\ n\ k) &= ret\ n\ (abs_7\ k) \\ abs_7\ (BSub_7\ k) &= flip\ (sub\ (abs_7\ k)) \quad \text{-- note reversal of arguments again} \end{aligned}$$

Appending representations corresponds to composing their interpretations:

$$b\ (abs_7\ y)\ (abs_7\ x) = abs_7\ (append_7\ x\ y)$$

(note that backwards composition on the left-hand side becomes forwards sequencing on the right-hand side). This gives yet another implementation of the evaluator:

$$\begin{aligned} eval_7 &: Expr \rightarrow Integer \\ eval_7\ e &= abs_7\ (rep_7\ e) \end{aligned}$$

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity

Since non-leaf nodes in the datatype $ExprRep_7$ have by construction a single child, this is a linear representation rather than a tree-shaped one, and there is an obvious compilation into the linear type of programs from before:

$$\begin{aligned} compile_7 &: Expr \rightarrow Prog_4 \\ compile_7 &= compileRep_7 \cdot rep_7 \text{ where} \\ compileRep_7 &: ExprRep_7 r \rightarrow Prog_4 \\ compileRep_7 Halt_7 &= [] \\ compileRep_7 (BRet_7 n k) &= PushI n :: compileRep_7 k \\ compileRep_7 (BSub_7 k) &= SubI :: compileRep_7 k \end{aligned}$$

Thus, $Halt_7$ corresponds to the empty program, and $BRet_7$ and $BSub_7$ each prefix one instruction onto a program.

The type $ExprRep_7 r$ corresponds to programs which yield a single integer when run on an initial stack of shape r ; there are no representations corresponding to incomplete programs such as $[PushI\ 3, PushI\ 4]$ which result in more than one value being pushed onto the stack.

6 Discussion

Reynolds [27] introduced the two-step process of conversion to continuation-passing style followed by defunctionalization of the continuations so introduced, in order to describe and classify interpreters used as a mechanism of language definition. Danvy and colleagues [13, 2] have extensively explored Reynolds's process, using it to explain the correspondence between various published interpreters and abstract machines, and to derive some new abstract machines from existing interpreters and vice versa. Our contribution has been to highlight the appeals to associativity in applications of this process, prominent in Wand's early application [29, 30] of Reynolds's approach but mostly unspoken since.

Reynolds and Danvy et al. focus on the lambda calculus as the defined language, using the process to explain the lambda calculus by implementing it in a simpler defining language. Subsequent work has focussed on the process, and has therefore chosen simpler defined languages than the lambda calculus as illustrations. In particular, a lot can be said in terms of a mere expression language consisting of numeric constants and addition, without touching the full Turing completeness of the lambda calculus. Hutton and colleagues have used this device in a series of papers [19, 5, 18, 6, 17]—so much so that McBride [23, 22] called it 'Hutton's Razor', although the device is actually much older than Hutton's uses [24, 30].

We've used Hutton's Razor ourselves in Section 5, although we switched from addition to subtraction in order to be clear about where associativity is used. Following Reynolds's process leads to an abstract machine for evaluating arithmetic expressions, but not a compiler, because it cannot make a clear phase distinction between processing the input expression and performing arithmetic operations. Nevertheless, there is a well-known compilation scheme from expressions to a stack machine [24]. Bahr and

Hutton [5] do manage to derive this compiler; but they do so by introducing stacks and stack transformers, motivated by operational reasoning:

The next step is to transform the evaluation function into a version that utilises a stack, in order to make the manipulation of argument values explicit. [5, §2.2]

The next step is to transform [the evaluator] into CPS, in order to make the flow of control explicit. [5, §2.3]

The operational justification (“in order to make . . . explicit”) seems a pity, because CPS and defunctionalization already introduce stacks and make control flow explicit—these artifacts appear just by turning the handle, with no operational insight required.

Our development in Section 5 shows that the operational reasoning is not needed in order to arrive at a compiler. Instead, we got there by exploiting associativity of a generalized composition operator, an observation due to Wand [30], and explored in detail for more sophisticated languages by Henson [14, Chapter 8]. Wand and Henson used an untyped setting; precise typing of the generalized composition requires (at least lightweight) dependent types. In fact, those dependent types express stack safety, making stack underflow a type error [25]. If one were to fully bite the dependently typed bullet, one could go further: to express not only the number and types of stack elements, but even their actual values [22, 26, 20], completing the proof of correctness of the compiler.

We have also shown that the same process applies to simpler problems than compilers, simpler even than compilers for basic expression languages. We showed in Section 2 that it leads from the direct recursive implementation of factorial to the simple imperative loop for computing it—once associativity of multiplication has been exploited. Similarly, with associativity, it leads from the naive quadratic-time implementation of reverse to the linear-time accumulating-parameter version, essentially an application of Hughes’s ‘novel representation’ [16] or difference lists in logic programming [11]. And again in conjunction with associativity, it leads to a linear-time tail-recursive tree traversal. Admittedly, there is a simpler linear-time tree traversal, if one dispenses with tail recursion [7, Exercise 71]:

$$\begin{aligned} \text{flatten } f &= \text{flatten}' t \text{ [] } \mathbf{where} \\ \text{flatten}' (\text{Tip } x) \quad xs &= x :: xs \\ \text{flatten}' (\text{Bin } t \ u) \quad xs &= \text{flatten}' t (\text{flatten}' u \ xs) \end{aligned}$$

And if one follows the tree traversal development steps on the identity function on trees

$$\text{id} : \text{Tree } a \rightarrow \text{Tree } a$$

then the defunctionalized continuations $\text{List } (\text{Either } (\text{Tree } a) (\text{Tree } a))$ correspond to the zipper [15]. Do the same for the map function on trees

$$\text{treeMap} : (a \rightarrow b) \rightarrow \text{Tree } a \rightarrow \text{Tree } b$$

and the defunctionalized continuations $\text{List } (\text{Either } (\text{Tree } b) (\text{Tree } a))$ are tree dissections [21]—but neither of these exploits associativity.

Acknowledgements

I'm very grateful for constructive feedback from members of the Algebra of Programming research group at Oxford, especially Guillaume Boisseau, Josh Ko, Shin Cheng Mu, Richard Bird, Geraint Jones, Nick Wu, and Cezar Ionescu, and to Mitch Wand for comments on his paper. I also owe thanks to members of IFIP Working Group 2.1 and participants at BOBKonf 2021 for helpful questions and comments, to reviewers of an earlier version for constructive suggestions, and to Keith Clark and Tom Schrijvers for bibliographic assistance.

References

- [1] Mads Sig Ager, Dariusz Biernacki, Olivier Danvy, and Jan Midtgaard. From interpreter to compiler and virtual machine: A functional derivation. Technical Report BRICS RS-03-14, DAIMI, Department of Computer Science, University of Aarhus, 2003. <https://www.brics.dk/RS/03/14/>.
- [2] Mads Sig Ager, Dariusz Biernacki, Olivier Danvy, and Jan Midtgaard. A functional correspondence between evaluators and abstract machines. In *Principles and Practice of Declarative Programming*, pages 8–19. ACM, 2003. doi:10.1145/888251.888254.
- [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. doi:10.1016/j.ipl.2004.02.012.
- [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. doi:10.1016/j.tcs.2005.06.008.
- [5] Patrick Bahr and Graham Hutton. Calculating correct compilers. *Journal of Functional Programming*, 25:e14, 2015. doi:10.1017/S0956796815000180.
- [6] Patrick Bahr and Graham Hutton. Calculating correct compilers II: Return of the register machines. *Journal of Functional Programming*, 30(e25), 2020. doi:10.1017/S0956796820000209.
- [7] Richard Bird. *Thinking Functionally with Haskell*. Cambridge University Press, 2015.
- [8] Richard Bird and Philip Wadler. *An Introduction to Functional Programming*. Prentice-Hall, 1988.
- [9] Richard S. Bird. The promotion and accumulation strategies in transformational programming. *ACM Transactions on Programming Languages and Systems*, 6(4):487–504, October 1984. doi:10.1145/1780.1781.
- [10] Rod M. Burstall and John Darlington. A transformational system for developing recursive programs. *Journal of the ACM*, 24(1):44–67, January 1977. doi:10.1145/321992.321996.

- [11] Keith L. Clark and Sten-Åke Tärnlund. A first-order theory of data and programs. In Bruce Gilchrist, editor, *Proceedings of IFIP Congress*, pages 939–944, Amsterdam, 1977. North-Holland.
- [12] Olivier Danvy. Defunctionalized interpreters for programming languages. In *International Conference on Functional Programming*, pages 131–142. ACM, 2008. doi:10.1145/1411204.1411206.
- [13] Olivier Danvy and Lasse R. Nielsen. Defunctionalization at work. In *Principles and Practice of Declarative Programming*, pages 162–174. ACM, 2001. doi:10.1145/773184.773202.
- [14] Martin C. Henson. *Elements of Functional Languages*. Blackwell Scientific Publishing, 1987.
- [15] Gérard Huet. The zipper. *Journal of Functional Programming*, 7(5):549–554, September 1997. doi:10.1017/S0956796897002864.
- [16] John Hughes. A novel representation of lists and its application to the function ‘reverse’. *Information Processing Letters*, 22:141–144, 1986. doi:10.1016/0020-0190(86)90059-1.
- [17] Graham Hutton. It’s as easy as 1, 2, 3. Unpublished draft, <https://www.cs.nott.ac.uk/~pszgmh/123.pdf>, January 2021.
- [18] Graham Hutton and Patrick Bahr. Compiling a 50-year journey. *Journal of Functional Programming*, 27(e20), 2017. doi:10.1017/S0956796817000120.
- [19] Graham Hutton and Joel J. Wright. Compiling exceptions correctly. In Dexter Kozen and Carron Shankland, editors, *Mathematics of Program Construction*, volume 3125 of *Lecture Notes in Computer Science*, pages 211–227. Springer, 2004. doi:10.1007/978-3-540-27764-4_12.
- [20] Hsiang-Shang Ko. McBride’s Razor. <https://josh-hs-ko.github.io/blog/0010/>, 2020.
- [21] Conor McBride. Clowns to the left of me, jokers to the right: Dissecting data structures. In *Principles of Programming Languages*, pages 287–295. ACM, 2008. doi:10.1145/1328438.1328474.
- [22] Conor McBride. Ornamental algebras, algebraic ornaments. Unpublished draft, <https://personal.cis.strath.ac.uk/conor.mcbride/pub/OAAO/LitOrn.pdf>, 2011.
- [23] Conor McBride. She’s cutting and pasting it: Higgledy piggledy programming. <https://personal.cis.strath.ac.uk/conor.mcbride/pub/she/higpig.html>, 2012.
- [24] John McCarthy and James Painter. Correctness of a compiler for arithmetic expressions. In Jacob T. Schwartz, editor, *Mathematical Aspects of Computer Science*, volume 19 of *Proceedings of Symposia in Applied Mathematics*, pages 33–41. American Mathematical Society, 1967. doi:10.1090/psapm/019.
- [25] James McKinna and Joel Wright. A type-correct, stack-safe, provably correct, expression compiler for Epigram. Unpublished draft, <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.105.4086>, 2006.
- [26] Alberto Pardo, Emmanuel Gunther, Miguel Pagano, and Marcos Viera. An internalist approach to correct-by-construction compilers. In David Sabel and

Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity

- Peter Thiemann, editors, *Principles and Practice of Declarative Programming*. ACM, 2018. doi:10.1145/3236950.3236965.
- [27] John C. Reynolds. Definitional interpreters for higher-order programming languages. In John J. Donovan and Rosemary Shields, editors, *Proceedings of the 25th ACM National Conference*, pages 717–740. ACM, 1972. Reprinted in *Higher-Order and Symbolic Computation* 11(4):363–397, 1998. doi:10.1145/800194.805852.
- [28] Raymond Smullyan. *To Mock a Mockingbird, and Other Logic Puzzles*. Oxford University Press, 1985.
- [29] Mitchell Wand. Continuation-based program transformation strategies. *Journal of the ACM*, 27(1):164–180, January 1980.
- [30] Mitchell Wand. Deriving target code as a representation of continuation semantics. *ACM Transactions on Programming Languages and Systems*, 4(3):496–517, 1982. doi:10.1145/357172.357179.

A Appendix: A primer on Idris syntax

For the purposes of this paper, Idris is mostly used simply as a strongly typed pure functional programming language, similar in spirit to Haskell and Standard ML.

Syntactically, like both those languages, Idris uses square brackets `[1, 2, 3]` for list values and `[]` for the empty list. It uses Haskell’s `++` rather than ML’s `@` for list append, Haskell’s `·` rather than ML’s `◦` for function composition, and Haskell’s `--` for line comments (which ML lacks). Lambda abstractions $\lambda x \Rightarrow e$ use Haskell’s λ (a backslash in ASCII) instead of ML’s `fn`, but ML’s \Rightarrow instead of Haskell’s \rightarrow . Case expressions `case x of p \Rightarrow e | p' \Rightarrow e'` also use ML’s \Rightarrow instead of Haskell’s \rightarrow . Idris uses the ML convention of `:` for type declarations and `::` for list cons, the opposite way round from the Haskell convention. Datatype declarations:

```
data Tree a = Tip a | Bin (Tree a) (Tree a)
```

use the same syntax as Haskell; but the list datatype is named (as in *List Integer*) rather than using brackets.

Each definition must have a type declaration, rather than allowing the type to be inferred as in ML and (most of the time) in Haskell. Functions must be defined before use; mutually recursive functions must be given in a **mutual** block:

```
parity : Nat  $\rightarrow$  Bool
parity n = odd n where
  mutual
    odd : Nat  $\rightarrow$  Bool
    odd Z    = False
    odd (S n) = even n
    even : Nat  $\rightarrow$  Bool
    even Z   = True
    even (S n) = odd n
```

Idris is dependently typed: types are first-class citizens, and the type of types is *Type*. So what would be a type synonym in Haskell:

```
type FlatCont4 a = [Either [a] (Tree a)]
```

is just an ordinary definition (of a function from types to types) in Idris:

```
FlatCont4 : Type → Type
FlatCont4 a = List (Either (List a) (Tree a))
```

We make only light use of the power of dependent types in this paper. The *Arrow* type from Section 5:

```
Arrow : List Type → Type → Type
Arrow [ ]      b = b
Arrow (a :: as) b = a → Arrow as b
```

is just another ordinary definition, constructing a type *Arrow as b* indexed by a list of argument types *as* and a result type *b*. In Haskell it would require the lightweight dependently typed features of the GHC compiler (specifically, type level analogues of value-level lists, and we would also need the corresponding type-level function to append type-level lists); in ML it could not be expressed at all. Generalized composition:

```
b : {as : List Type} → (b → c) → Arrow as b → Arrow as c
b {as = [ ]}      gf = gf
b {as = _ :: _}   gf = b g · f
```

is defined by induction over the argument types index *as*; the curly brackets denote that *as* is an implicit argument—it could be inferred from the overall type (as *b, c* are), but it made explicit to enable pattern matching on whether or not *as* is empty in the two clauses of the definition. In Haskell, it would require a family of singleton types in order to connect the value level with the type level; again, it is inexpressible in ML.

About the author

Jeremy Gibbons is Professor of Computing at the University of Oxford, where he leads the *Algebra of Programming* research group.

email: jeremy.gibbons@cs.ox.ac.uk

www: <http://www.cs.ox.ac.uk/jeremy.gibbons/>

orcid: 0000-0002-8426-9917