

15–150: Principles of Functional Programming

Some Notes on Continuations

Michael Erdmann*

Spring 2024

These notes provide a brief introduction to continuations as a programming technique. Continuations have a rich history and have made important contributions to the understanding of functional programming and compilation. One can distinguish *first-class continuations* (which are not available in Standard ML, although the SML of New Jersey implementation provides them) and *functions as continuations*. In these notes, we talk only about the latter; thus we do not refer to an extension of the language, but to a particular higher-order programming technique.

1 A Tail-Recursive Append Function

Recall the definition of function @ that appends two lists

```
(* @ : 'a list * 'a list -> 'a list *)
fun @ (nil, L) = L
  | @ (x::xs, L) = x :: @(xs, L)
infixr @
```

Note that this function is not tail-recursive, since it applies the list constructor `::` to the result of the recursive call. Nonetheless, this function is quite efficient and the definition above is fully satisfactory from a pragmatic point of view.

But one may still ask if there is any way to write an append function in tail-recursive form. The answer is “yes”, although it will be less efficient than the direct version above. In fact, it is a deep property of SML that *every* function can be rewritten in tail-recursive form!

Suppose that we have a function $f : t \rightarrow s$ and would like to rewrite it as a function f' in tail-recursive form. The basic idea is to give f' an additional argument, called the *continuation*, which represents the computation that should be done on the result of f . In the base case, instead of returning a result, we call the continuation. This means that f' should have type

$$f' : t \rightarrow (s \rightarrow 'b) \rightarrow 'b.$$

In the recursive case we add whatever computation should be done on the result to the continuation. So a continuation is like a functional accumulator argument.

When we use this function to compute f we give it the *initial continuation* which is often the identity function, indicating no further computation is done on the result.

*Modified from a draft by Frank Pfenning.

Applying this basic idea to `append` yields the following definition:

```
local
  (* @' : 'a list * 'a list * ('a list -> 'b) -> 'b *)
  fun @' (nil, L, cont) = cont L
    | @' (x::xs, L, cont) = @' (xs, L, fn r => cont (x::r))
in
  fun @ (X, L) = @' (X, L, fn r => r)
end
```

This first line of `@'` implements the idea that instead of returning the result `L`, we apply the continuation to `L`.

The second line of `@'` implements the idea that instead of constructing

```
x :: @ (xs, L)
```

we call `@'` recursively on `xs` and `L`, adding the task of prepending `x` to the argument `r` of the continuation `cont`.

Consider the following sample computation (here, we have renamed some instances of the variables bound in the continuations and directly substituted variable bindings for instances of `x` and `cont` in the bodies of the lambda expressions, in order to make the reductions easier to read):

```
@ ([1,2], [3,4])
=> @' ([1,2], [3,4], fn r => r)
=> @' ([2], [3,4], fn r1 => (fn r => r) (1::r1))
=> @' ([], [3,4], fn r2 => (fn r1 => (fn r => r) (1::r1)) (2::r2))
=> (fn r2 => (fn r1 => (fn r => r) (1::r1)) (2::r2)) [3,4]
=> (fn r1 => (fn r => r) (1::r1)) (2::[3,4])
= (fn r1 => (fn r => r) (1::r1)) [2,3,4]
=> (fn r => r) (1::[2,3,4])
= (fn r => r) [1,2,3,4]
=> [1,2,3,4]
```

Even though the function is now tail-recursive, in this example we haven't saved anything since we have to build up and then call the continuation. Compare the above with the direct computation:

```
@ ([1,2], [3,4])
=> 1 :: @ ([2], [3,4])
=> 1 :: (2 :: @ ([], [3,4]))
=> 1 :: (2 :: [3,4])
=> 1 :: [2,3,4]
=> [1,2,3,4]
```

Continuations are useful not necessarily because of efficiency gains but because they provide a direct handle on future computation. We will soon see the benefit of having such a handle.

2 Proving Correctness

The correctness statement for the continuation-passing tail-recursive version of append expresses formally the English explanation given informally earlier:

Theorem 1 *For any list values X and L and continuation values c , all of appropriate type:*

$$\text{@}' (X, L, c) \cong c (\text{@} (X, L)).$$

Proof: By structural induction on X .

Base Case: $X = \text{nil}$.

We need to show that $\text{@}' (\text{nil}, L, c) \cong c (\text{@} (\text{nil}, L))$ for all values L and c .

Observe that

$$\text{@}' (\text{nil}, L, c) \implies c (L) \quad [\text{first clause of @}']$$

and

$$c (\text{@} (\text{nil}, L)) \implies c (L) \quad [\text{first clause of @}]$$

so both sides reduce to the same expression, and thus are equivalent.

Inductive Case: $X = x::xs$. Assume inductively that, for all values L and c' ,

$$\text{@}' (xs, L, c') \cong c' (\text{@} (xs, L)).$$

We have to show that, for all values L and c ,

$$\text{@}' (x::xs, L, c) \cong c (\text{@} (x::xs, L)).$$

Showing:

$$\begin{aligned} & \text{@}' (x::xs, L, c) \\ \cong & \text{@}' (xs, L, \text{fn } r \Rightarrow c (x::r)) && [\text{by second clause of @}'] \\ \cong & (\text{fn } r \Rightarrow c (x::r)) (\text{@} (xs, L)) && [\text{by IH with } c' \text{ the continuation } (\text{fn } r \Rightarrow c (x::r))] \\ \cong & c (x::(\text{@} (xs, L))) && [\text{since @ is total (by an earlier handout)}] \\ \cong & c (\text{@} (x::xs, L)) && [\text{by second clause of @}] \end{aligned}$$

That establishes the induction step and completes the proof.

□

3 Control with Continuations

In the function that appends two lists, continuations can be applied, but they do not help in writing simpler or more efficient code. But in functions with more complex control flow, they can often be used to advantage. In a future lecture, we will see a major application, namely the implementation of backtracking in an acceptor for regular expressions.

Here, we consider a simpler example, namely a function that traverses a list and returns `SOME` of the shortest prefix of the list such that a predicate p is false on the next element. If p holds on all elements, the function returns `NONE` to indicate no such prefix exists.

First, a direct implementation, using a helper function that conses an element onto an optional list.

```
(* consOpt : 'a * ('a list) option -> ('a list) option *)
fun consOpt (x, NONE) = NONE
  | consOpt (x, SOME (L)) = SOME (x::L)

(* findprefix : ('a -> bool) -> 'a list -> ('a list) option *)
fun findprefix p nil = NONE
  | findprefix p (x::xs) =
    if p(x) then consOpt (x, findprefix p xs)
    else SOME(nil)
```

Note an inefficiency in this function: if p is true for every element in the list, `consOpt` is called on every element after the end of the list has been reached, passing along `NONE` all the way up to be returned as the overall result.

Instead, we would like to return `NONE` immediately once we have reached the end of the list and found no element on which p is false. Using continuations we can achieve this easily, simply by ignoring the continuation in this case.

```
local
  (* findpref : ('a -> bool)
     -> 'a list * ('a list -> 'b option)
     -> 'b option *)
  (* Used as : ('a -> bool)
     -> ('a list * ('a list -> ('a list) option))
     -> ('a list) option *)
  fun findpref p (nil, k) = NONE
    | findpref p (x::xs, k) =
      if p(x) then findpref p (xs, fn r => k (x::r))
      else k(nil)
in
  (* findpref' : ('a -> bool) -> 'a list -> ('a list) option *)
  fun findpref' p L = findpref p (L, fn r => SOME r)
end
```

Note that in the “success” case (we find an element where p is false) we call the continuation k on the expected initial answer (`nil` in this program). In the case of a “failure” (we do not find a place where p is false) we discard the continuation and return directly (with `NONE` in this program). In this pattern of use we refer to the continuation k as a *success continuation*.

4 Further Examples

The pattern of control required for the `findprefix` function can also be achieved by using *exceptions*, which we will discuss in more detail in a future lecture. However, for more complicated control patterns, continuations are often a natural way to program a solution. Many backtracking algorithms (such as the regular expression matcher in an upcoming lecture) are of this nature.

Another class of examples includes *co-routines*, which are functions that call each other in an egalitarian way (as distinct from *subroutines*, where one is subordinate to the other). Applications of continuations also arise when the call pattern is at cross-purposes with the natural modularization of the code. In this case we pass a continuation to the function in another module that can then “return” control by invoking the continuation. In the language of operating systems this is often referred to as an *up-call*.

To summarize the important principles that guide the implementation of a function passing a continuation:

- In base cases, we apply the continuation instead of directly returning a value.
- In recursive cases, the continuation acts as a functional accumulator.
- In exceptional cases, we may discard (or duplicate) the continuation to circumvent the normal control flow.