

References; Continuation passing style

Lecture 10

Thursday, February 27, 2025

1 References

We introduce constructs for creating, reading, and updating memory locations, also called *references*. The resulting language is still a functional language (since functions are first-class values), but expressions can have side-effects, that is, they can modify state.

The syntax is defined as follows.

$$e ::= x \mid \lambda x. e \mid e_0 e_1 \mid \mathbf{ref} e \mid !e \mid e_1 := e_2 \mid \ell$$

$$v ::= \lambda x. e \mid \ell$$

Expression $\mathbf{ref} e$ creates a new memory location (like a malloc), and sets the initial contents of the location to (the result of) e . The expression $\mathbf{ref} e$ itself evaluates to a memory location ℓ . Think of a location as being like a pointer to a memory address. The expression $!e$ assumes that e evaluates to a memory location, and $!e$ evaluates to the current contents of the memory location. Expression $e_1 := e_2$ assumes that e_1 evaluates to a memory location ℓ , and updates the contents of ℓ with (the result of) e_2 . Locations ℓ are not intended to be used directly by a programmer: they are not part of the *surface syntax* of the language, the syntax that a programmer would write. They are introduced only by the operational semantics.

We define a small-step CBV operational semantics. We use configurations $\langle e, \sigma \rangle$, where e is an expression, and σ is a map from locations to values.

$$E ::= [\cdot] \mid E e \mid v E \mid \mathbf{ref} E \mid !E \mid E := e \mid v := E$$

$$\frac{\langle e, \sigma \rangle \longrightarrow \langle e', \sigma' \rangle}{\langle E[e], \sigma \rangle \longrightarrow \langle E[e'], \sigma' \rangle}$$

$$\beta\text{-REDUCTION} \frac{}{\langle (\lambda x. e) v, \sigma \rangle \longrightarrow \langle e\{v/x\}, \sigma \rangle} \quad \text{ALLOC} \frac{}{\langle \mathbf{ref} v, \sigma \rangle \longrightarrow \langle \ell, \sigma[\ell \mapsto v] \rangle} \ell \notin \text{dom}(\sigma)$$

$$\text{DEREF} \frac{}{\langle !\ell, \sigma \rangle \longrightarrow \langle v, \sigma \rangle} \sigma(\ell) = v \quad \text{ASSIGN} \frac{}{\langle \ell := v, \sigma \rangle \longrightarrow \langle v, \sigma[\ell \mapsto v] \rangle}$$

References do not add any expressive power to the lambda calculus: it is possible to translate lambda calculus with references to the pure lambda calculus. Intuitively, this is achieved by explicitly representing the store, and threading the store through the evaluation of the program. The details are left as an exercise.

2 Continuations

So far we have seen a number of language features that extend lambda calculus, and have translated many of these into the pure lambda calculus, using a direct semantic translation. That is, the control structure of the source language translated into the corresponding control structure in the target language:

$$\mathcal{T}[\lambda x. e] = \lambda x. \mathcal{T}[e]$$

$$\mathcal{T}[e_1 e_2] = \mathcal{T}[e_1] \mathcal{T}[e_2]$$

This style of translation works well when the source language is similar to the target language. However, when the control structures of the source and target languages differ considerable, it doesn't work as well.

In this lecture, we'll learn about *continuations* and *continuation-passing style*. These can be used to define the semantics of control-flow constructs such as exceptions.

Continuations are a programming technique that may be used directly by a programmer, or used in program transformations by a compiler.

Intuitively, a continuation represents “the rest of the program.” Consider the program

```
if foo < 10 then 32 + 6 else 7 + bar
```

and consider the evaluation of the expression `foo < 10`. When we finish evaluating this subexpression, we will evaluate the if statement, and then evaluate the appropriate branch. The *continuation* of the subexpression `foo < 10` is the rest of the computation that will occur after we evaluate the subexpression. We can write this continuation as a function that takes the result of the subexpression:

$$(\lambda y. \text{if } y \text{ then } 32 + 6 \text{ else } 7 + \text{bar}) (\text{foo} < 10)$$

The evaluation order and result of this program will be the same as the original expression; the difference is that we extracted the continuation of the subexpression in to a function.

The nice thing about continuations is that it makes the control explicit, and this is especially useful in the case of functional programs, where control is not explicit otherwise. In fact, we can rewrite a program to make continuations more explicit.

Let's consider another program, and convert it so that continuations are explicit. Let's think about the following applied lambda calculus program.

$$(\lambda x. x) ((1 + 2) + 3) + 4$$

We'll start by defining a continuation for the outermost evaluation context, which takes a value, and applies the identity function to it.

$$k_0 = \lambda v. (\lambda x. x) v$$

The evaluation context that is evaluated next-to-last takes a value, adds 4 to it, and then passes the result to k_0 .

$$k_1 = \lambda a. k_0 (a + 4)$$

Likewise, for the next evaluation contexts.

$$k_2 = \lambda b. k_1 (b + 3)$$

$$k_3 = \lambda c. k_2 (c + 2)$$

The program itself is now equivalent to $k_3 1$. Since `let $x = e$ in e'` is just syntactic sugar for $(\lambda x. e') e$, we can actually rewrite the above as

```
let c = 1 in
let b = c + 2 in
let a = b + 3 in
let v = a + 4 in
( $\lambda x. x$ ) v
```

This is fairly close to some machine instructions of the form:

```
set c, 1
add b, c, 2
add a, b, 3
add v, a, 4
call id, v
```

Using continuations, functions can be transformed into “functions that don’t return”—functions that take, besides the usual arguments, an additional argument representing a continuation. When the function finishes, it invokes the continuation on its result, instead of returning the result to its caller. Writing functions in this way is usually referred to as Continuation-Passing Style, or CPS for short. For instance, the CPS version of factorial looks like the following:

$$FACT_{cps} = Y \lambda f. \lambda n, k. \text{if } n = 0 \text{ then } k \ 1 \ \text{else } f \ (n - 1) \ (\lambda v. k \ (n * v))$$

Note that the last thing that code in $FACT_{cps}$ does is call a function (either k or f), and does not do anything with the result.

Continuation-Passing Style is an important concept in the compilation of functional languages and is used as an intermediate compiler representation (it has been used in compilers for Scheme, ML, etc). The main advantage is that CPS makes the control flow explicit and makes it easier to translate functional code to machine code where control is explicit (in the form of sequences of machine instructions and jumps). For instance, a CPS call can be easily translated into a jump to the invoked method, since the invoked function does not return the control.

2.1 CPS translation

We can translate lambda calculus programs into continuation-passing style. We define a translation function $CPS[\cdot]$, which takes a CBV lambda calculus expression, and translates the expression to a CBV lambda calculus expression in continuation-passing style.

Let’s consider a translation from lambda calculus with pairs and integers. The syntax of the source language is as follows.

$$e ::= x \mid \lambda x. e \mid e_1 e_2 \mid n \mid e_1 + e_2 \mid (e_1, e_2) \mid \#1 e \mid \#2 e$$

The translation $CPS[e]$ will produce a function whose argument is the continuation to which to pass the result. That is, for all expressions e , the translation is of the form $CPS[e] = \lambda k. \dots$, where k is a continuation. We will both assume and guarantee that for any expression e , the translation $CPS[e] = \lambda k. \dots$ will apply k to the result of evaluating e .

For convenience, instead of writing

$$CPS[e] = \lambda k. \dots$$

we write

$$CPS[e]k = \dots$$

$$\begin{aligned} CPS[n]k &= k \ n \\ CPS[e_1 + e_2]k &= CPS[e_1] (\lambda n. CPS[e_2] (\lambda m. k \ (n + m))) && n \text{ is not a free variable of } e_2 \\ CPS[(e_1, e_2)]k &= CPS[e_1] (\lambda v. CPS[e_2] (\lambda w. k \ (v, w))) && v \text{ is not a free variable of } e_2 \\ CPS[\#1 e]k &= CPS[e] (\lambda v. k \ (\#1 v)) \\ CPS[\#2 e]k &= CPS[e] (\lambda v. k \ (\#2 v)) \\ \\ CPS[x]k &= k \ x \\ CPS[\lambda x. e]k &= k \ (\lambda x, k'. CPS[e]k') && k' \text{ is not a free variable of } e \\ CPS[e_1 e_2]k &= CPS[e_1] (\lambda f. CPS[e_2] (\lambda v. f \ v \ k)) && f \text{ is not a free variable of } e_2 \end{aligned}$$

We translate a function $\lambda x. e$ to a function that takes an additional argument k' , which is the continuation after the function application. That is, k' is the continuation to which we hand the result of evaluating the

function body. In function application, we see that in addition to the actual argument, we also give the continuation as the additional argument.

Let's see an example translation and execution...

$$\begin{aligned}
\mathcal{CPS}[(\lambda a. a + 6) 7] ID &= \mathcal{CPS}[(\lambda a. a + 6)] (\lambda f. \mathcal{CPS}[7] (\lambda v. f v ID)) \\
&= (\lambda f. \mathcal{CPS}[7] (\lambda v. f v ID)) (\lambda a, k'. \mathcal{CPS}[a + 6] k') \\
&= (\lambda f. (\lambda v. f v ID) 7) (\lambda a, k'. \mathcal{CPS}[a + 6] k') \\
&= (\lambda f. (\lambda v. f v ID) 7) (\lambda a, k'. \mathcal{CPS}[a] (\lambda n. \mathcal{CPS}[6] (\lambda m. k' (m + n)))) \\
&= (\lambda f. (\lambda v. f v ID) 7) (\lambda a, k'. \mathcal{CPS}[a] (\lambda n. (\lambda m. k' (m + n)) 6)) \\
&= (\lambda f. (\lambda v. f v ID) 7) (\lambda a, k'. (\lambda n. (\lambda m. k' (m + n)) 6) a) \\
&\rightarrow (\lambda v. (\lambda a, k'. (\lambda n. (\lambda m. k' (m + n)) 6) a) v ID) 7 \\
&\rightarrow (\lambda a, k'. (\lambda n. (\lambda m. k' (m + n)) 6) a) 7 ID \\
&\rightarrow (\lambda n. (\lambda m. ID (m + n)) 6) 7 \\
&\rightarrow (\lambda m. ID (m + 7)) 6 \\
&\rightarrow ID (6 + 7) \\
&\rightarrow ID 13 \\
&\rightarrow 13
\end{aligned}$$