

1 Continuation-Passing Style

Consider the statement `if $x \leq 0$ then x else $x + 1$` . We can think of this as $(\lambda y. \text{if } y \text{ then } x \text{ else } x + 1) (x \leq 0)$. To evaluate this, we would first evaluate the argument $x \leq 0$ to obtain a Boolean value, then apply the function $\lambda y. \text{if } y \text{ then } x \text{ else } x + 1$ to this value. The function $\lambda y. \text{if } y \text{ then } x \text{ else } x + 1$ is called a *continuation*, because it specifies what is to be done with the result of the current computation in order to continue the computation.

Given an expression e , it is possible to transform the expression into a function that takes a continuation k and applies it to the value of e . The transformation is applied recursively. This is called *continuation-passing style* (CPS). There are a number of advantages to this style:

- The resulting expressions have a much simpler evaluation semantics, since the sequence of reductions to be performed is specified by a series of continuations. The next reduction to be performed is always uniquely determined, and the remainder of the computation is handled by a continuation. Thus evaluation contexts are not necessary to specify the evaluation order.
- In practice, function calls and function returns can be handled in a uniform way. Instead of returning, the called function simply calls the continuation.
- In a recursive function, any computation to be performed on the value returned by a recursive call can be bundled into the continuation. Thus every recursive call becomes tail-recursive. For example, the factorial function

$$\text{fact } n = \text{if } n = 0 \text{ then } 1 \text{ else } n * \text{fact } (n - 1)$$

becomes

$$\text{fact}' n k = \text{if } n = 0 \text{ then } k 1 \text{ else } \text{fact}' (n - 1) (\lambda v. k (n * v)).$$

One can show inductively that $\text{fact}' n k = k (\text{fact } n)$, therefore $\text{fact}' n \lambda x. x = \text{fact } n$. This transformation effectively trades stack space for heap space in the implementation.

- Continuation-passing gives a convenient mechanism for non-local flow of control, such as `goto` statements and exception handling.

2 CPS Semantics

Our grammar for the λ -calculus was:

$$e ::= x \mid \lambda x. e \mid e_0 e_1$$

Our grammar for the CPS λ -calculus will be:

$$v ::= x \mid \lambda x_0, x_1, \dots, x_n. e \quad e ::= v_0 v_1 \cdots v_n$$

This is a highly constrained syntax. Barring reductions inside the scope of a λ -abstraction operator, the expressions v are all irreducible. The only reducible expression is $v_0 v_1 \cdots v_n$. If $n \geq 1$, there exactly one redex $v_0 v_1$, and both the function and the argument are already fully reduced. The small step semantics has a single rule

$$(\lambda x_0, x_1, \dots, x_n. e) v_0 v_1 \cdots v_n \rightarrow e \{v_o/x_i\}^{(i \in 1..n)},$$

and we do not need any evaluation contexts.

The big step semantics is also quite simple, with only a single rule:

$$\frac{e\{v_i/x_i\}^{(i \in 1..n)} \Downarrow v'}{(\lambda x_0, x_1, \dots, x_n. e) v_0 v_1 \dots v_n \Downarrow v'}$$

The resulting proof tree will not be very tree-like. The rule has one premise, so a proof will be a stack of inferences, each one corresponding to a step in the small-step semantics. This allows for a much simpler interpreter that can work in a straight line rather than having to make multiple recursive calls.

The fact that we can build a simpler interpreter for the language is a strong hint that this language is lower-level than the λ -calculus. Because it is lower-level (and actually closer to assembly code), CPS is typically used in functional language compilers as an intermediate representation. It also is a good code representation if one is building an interpreter.

3 CPS Conversion

Despite the restricted syntax of CPS, we have not lost any expressive power. Given a λ -calculus expression e , it is possible to define a translation $\llbracket e \rrbracket$ that translates it into CPS. This translation is known as *CPS conversion*. It was first described by John Reynolds. The translation takes an arbitrary λ -term e and produces a CPS term $\llbracket e \rrbracket$, which is a function that takes a continuation k as an argument. Intuitively, $\llbracket e \rrbracket k$ applies k to the value of e .

We want our translation to satisfy $e \xrightarrow{*}_{\text{CBV}} v$ iff $\llbracket e \rrbracket k \xrightarrow{*}_{\text{CPS}} \llbracket v \rrbracket k$ for primitive values v and any variable $k \notin \text{FV}(e)$, and $e \uparrow_{\text{CBV}}$ iff $\llbracket e \rrbracket k \uparrow_{\text{CPS}}$.

The translation is (adding numbers as primitive values):

$$\begin{aligned} \llbracket n \rrbracket k &\triangleq kn \\ \llbracket x \rrbracket k &\triangleq kx \\ \llbracket \lambda x. e \rrbracket k &\triangleq k(\lambda x. \llbracket e \rrbracket) = k(\lambda x k'. \llbracket e \rrbracket k') \\ \llbracket e_0 e_1 \rrbracket k &\triangleq \llbracket e_0 \rrbracket (\lambda f. \llbracket e_1 \rrbracket (\lambda v. fvk)). \end{aligned}$$

(Recall $\llbracket e \rrbracket k \triangleq e'$ really means $\llbracket e \rrbracket \triangleq \lambda k. e'$.)

3.1 An Example

In the CBV λ -calculus, we have

$$(\lambda xy. x) 1 \rightarrow \lambda y. 1$$

Let's evaluate the CPS-translation of the left-hand side using the CPS evaluation rules.

$$\begin{aligned} \llbracket (\lambda xy. x) 1 \rrbracket k &= \llbracket \lambda x. \lambda y. x \rrbracket (\lambda f. \llbracket 1 \rrbracket (\lambda v. fvk)) \\ &= (\lambda f. \llbracket 1 \rrbracket (\lambda v. fvk)) (\lambda x. \llbracket \lambda y. x \rrbracket) \\ &\rightarrow \llbracket 1 \rrbracket (\lambda v. (\lambda x. \llbracket \lambda y. x \rrbracket) vk) \\ &= (\lambda v. (\lambda x. \llbracket \lambda y. x \rrbracket) vk) 1 \\ &\rightarrow (\lambda x. \llbracket \lambda y. x \rrbracket) 1 k \\ &\rightarrow \llbracket \lambda y. 1 \rrbracket k. \end{aligned}$$

4 CPS Semantics for FL!

4.1 Syntax

Since FL! has references, we need to add a store σ to our notation. Thus we now have translations with the form $\mathcal{E}\llbracket e \rrbracket \rho k \sigma$, which means, “Evaluate e in the environment ρ with store σ and send the resulting value and the new store to the continuation k .” A continuation is now a function of a value and a store; that is, a continuation k should have the form $\lambda v \sigma. \dots$.

The translation for variables is as follows:

$$\mathcal{E}\llbracket x \rrbracket \rho k \sigma \triangleq k(\text{lookup } \rho x) \sigma$$

Note that if we think about this translation as a function and η -reduce away the σ , we obtain:

$$\mathcal{E}\llbracket x \rrbracket \rho k = \lambda \sigma. k(\text{lookup } \rho x) \sigma = k(\text{lookup } \rho x).$$

In particular, in the η -reduced version, we have the same translation that we had for FL. In general, any expression in FL! that is not state-aware can be η -reduced to the same translation as FL. Thus in order to translate to FL!, we need to add translation rules only for the functionality that is state-aware.

Assume that we have extended our `lookup` and `update` functions to apply to stores. The translations of the state-aware constructs in FL! are defined as follows:

$$\begin{aligned} \mathcal{E}\llbracket \text{ref } e \rrbracket \rho k &\triangleq \mathcal{E}\llbracket e \rrbracket \rho (\lambda v \sigma'. \text{let } (\ell, \sigma'') = (\text{malloc } \sigma' v) \text{ in } k(\text{Loc } \ell) \sigma'') \\ \mathcal{E}\llbracket !e \rrbracket \rho k &\triangleq \mathcal{E}\llbracket e \rrbracket \rho (\lambda \ell \sigma'. k(\text{lookup } \sigma' \ell) \sigma') \\ \mathcal{E}\llbracket e_1 := e_2 \rrbracket \rho k &\triangleq \mathcal{E}\llbracket e_1 \rrbracket \rho (\lambda \ell. \mathcal{E}\llbracket e_2 \rrbracket \rho (\lambda v \sigma'. k \text{ null } (\text{update } \sigma' \ell v))) \end{aligned}$$