

- Pass by reference
- Continuation-passing style (CPS)
- CPS semantics
- CPS conversion

1 Functions with mutable variables

$$\begin{aligned} \mathcal{R}[\lambda x.e]\rho\sigma &= (\lambda y\sigma_{\text{dyn}}.\mathbf{let} \ell = \mathit{MALLOC}\sigma_{\text{dyn}} \mathbf{in} \mathbf{let} \sigma'' = \mathit{UPDATE} \sigma_{\text{dyn}} \ell y \mathbf{in} \mathcal{R}[e](\mathit{EXTEND} \rho \text{ "x" } \ell)\sigma', \sigma) \\ \mathcal{R}[e_1 e_2]\rho\sigma &= \mathbf{let} (f, \sigma') = \mathcal{R}[e_1]\rho\sigma \mathbf{in} \\ &\quad \mathbf{let} (v, \sigma'') = \mathcal{R}[e_2]\rho\sigma' \mathbf{in} fv\sigma'' \end{aligned}$$

Discussion:

- Choices we made: function applications don't affect state, in function application use left to right evaluation.
- In defining function application, the allocation is required because variables must be bound to locations, whereas the argument of the function is a value (C is call-by-value).
- The value that a function returns is an rvalue rather than an lvalue. Why? Returning an lvalue would mean we could assign new values to locations allocated during computation, e.g. `fib 5 = 1`; We choose a semantics where we get back the value of a computation rather than its location.
- Note that we do not need to explicitly define the store in the translation of function application because $f v \sigma''$ returns a pair of a location and a store.
- The store returned by a function application includes the newly allocated memory location ℓ . This is important for strong typing because otherwise the location could be returned to the caller. However, the semantics is not faithful to C, where returning an address to an argument creates a *dangling pointer*:

```
int *f(int x) –
    return &x;
„
...
int *y = f(0);
*y++; // behavior is “implementation-dependent”.
```

C does this because it allocates local variables on the *stack* rather than on the *heap*. When a function returns, stack-allocated locations are released.

2 Pass (call) by Reference

It is conceivable to define a language where arguments to functions are lvalues. In this case, the body of functions can update the argument and therefore change the value stored in the location passed in. Examples of such languages are PASCAL, MODULA, Ada, Matlab (sort of), FORTRAN.

Consider a simple syntax for pass-by-reference:

$$e ::= \dots \mid \lambda_r x.e \mid \mathbf{rcall} e_1 e_2$$

Note that we are explicitly saying when we pass by reference because we do not have a type system that distinguishes between the two cases.

Example program:

```

let  x = 1 in
let  f =  $\lambda_r y. y = y + 1$  in
      rcall f x

```

In this program, y is an alias for x , so the value at the location of x is updated to 2.

To incorporate pass-by-reference into the semantics, we just need to change the rules for functions:

$$\begin{aligned}
 \mathcal{R}[\lambda_r x. e] \rho \sigma &= (\lambda y \sigma_{\text{dyn}}. \mathcal{R}[e](\text{EXTEND } \rho \text{ "x" } y) \sigma_{\text{dyn}}, \sigma) \\
 \mathcal{R}[\text{rcall } e_1 \ e_2] \rho \sigma &= \text{let } (f, \sigma') = \mathcal{R}[e_1] \rho \sigma \text{ in} \\
 &\quad \text{let } (\ell, \sigma'') = \mathcal{L}[e_2] \rho \sigma' \text{ in } f \ell \sigma''
 \end{aligned}$$

Since we are given a location, we don't need to allocate new memory, and we don't have to worry about dangling pointers.

3 The limitations of direct translation

So far we have been building translations that preserve control flow. That is, the control flow in the target language corresponds directly to control flow in the source. This makes it difficult to describe language features that involve *non-local control flow*, such as errors, exceptions, goto, break, continue, and tail recursion. All of these features violate the simple stack-like control flow of the lambda calculus, which comes from the fact that functions impose a call/return discipline. To do a better job of capturing non-local control flow, we go beyond functions to *continuations*.

4 Continuations

Consider the statement **if** $x \leq 0$ **then** x **else** $x + 1$. Using evaluation contexts, we can separate the redex from the rest of the expression:

$$(\text{if } [\cdot] \text{ then } x \text{ else } x + 1)[x \leq 0]$$

Of course, we have another way of performing substitution: β reduction. We can also write 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. This function captures an evaluation context as an explicit, separate term. 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.

We can perform this kind of transformation on *every* evaluation context in the program. The result will be a term in *continuation-passing style* (CPS). Because it makes transfers of control explicit, CPS is a better target language when we want to capture more complex ways to transfer control.

5 Continuation-passing style

We can define a version of the lambda calculus that syntactically enforces continuation passing style.

Recall that 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:

$$v ::= x \mid \lambda x_1 \dots x_n. e \qquad e ::= v_0 v_1 \dots v_n$$

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 $v_0 v_1 \dots v_n$. In fact, we will only use $n = 1$ and $n = 2$ to get the full expressive power of the lambda calculus. When $n = 1$, we have a continuation. We will use $n = 2$ to model source-language functions.

The small-step semantics is particularly simple. It has a single rule for β reduction:

$$(\lambda x_1 \dots x_n. e) v_1 \dots v_n \longrightarrow e\{v_i/x_i\}^{(i \in 1..n)}$$

Notice that we do not need any evaluation contexts! The syntax ensures that if there is a β redex, it is at the top level.

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

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

The resulting proof tree will not be tree-like. The rule has one premise, so a proof will be a stack of rule instances, each one corresponding exactly to a step in the small-step semantics. In other words, the big- and small-step semantics correspond exactly.

Both semantics lead us to build an interpreter that runs a loop performing reductions, and does not need to make 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 lambda 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,

6 CPS conversion

Despite restricting the syntax in CPS, we haven't lost 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 developed as part of Steele's Rabbit Scheme compiler, but was described earlier by Reynolds.

The translation takes an arbitrary λ -term e and produces a CPS term $\llbracket e \rrbracket$ which is a function taking a continuation as argument. Intuitively, $\llbracket e \rrbracket k$ applies k to the result of e . We can think of a continuation as a drop box that expects a result and sends it on its way through the rest of the computation—but never returns.

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

One challenge is how to translate functions into CPS. Since all control transfers must be explicit applications of continuations, we must make the return from a function call explicit as a distinct continuation. Therefore, we represent a function $\lambda x. \dots$ in CPS form as a function that takes an explicit continuation argument: $\lambda k x. \dots$. The argument k corresponds to the return address to which the function will return.

Using this insight, the translation is as follows:

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

(Here we have defined $\llbracket e \rrbracket$ in the style of ML; thus $\llbracket e \rrbracket k = e'$ really means $\llbracket e \rrbracket = \lambda k. e'$.)

6.1 An example

In λ_{CBV} , 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. Note that we translate integers and other simple values just like we do variables:

$$\llbracket n \rrbracket k = k \mathcal{V} \llbracket n \rrbracket = k n$$

Converting to CPS, we get:

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