Lecture: Environments and Stores

In the last lecture, we used *substitution* to evaluate let-expressions and function application. Substitution is a familiar idea from algebra, which is one of the reasons why we used it. However, it is not a very efficient implementation strategy. Consider the following expression:

$$\text{let } x_1 = 1 \text{ in let } x_2 = 2 \text{ in } \cdots \text{let } x_n = n \text{ in } x_1 + \cdots + x_n$$

To evaluate this expression, we would first scan the entire program to substitute x_1 with 1, then scan again to substitute x_2 with 2, and so on, so evaluation by substitution would take $O(n^2)$ steps.

We now introduce an auxiliary data structure called an *environment* that stores the values bound to identifiers. Intuitively, instead of substituting, we lookup the value of identifiers in the environment.

1   Environment-based Semantics

Figure 5.1 presents an environment-based semantics for the language. Before we discuss the new semantics, we'll explain the more compact notation that the figure uses:

- Instead of spelling out every type of constant and binary operator in the definition of expressions, e, the figure uses the metavariable c to denote an arbitrary constant and op_2 to denote an arbitrary binary operator. In addition, instead of spelling out the semantics of every binary operator, we assume the presence of a partial function, δ, that consumes an operator and two constants and produces a result, if the result is defined. **Note**: we write $X \rightharpoonup Y$ to denote a partial function, whereas $X \to Y$ is a total function.

  The main reason for this change is that we can easily introduce new operators and new kinds of constants (e.g., strings) without changing the definition of e and $\Downarrow$.

- The syntax no longer has let-expressions. However, note that $\text{let } x = e_1 \text{ in } e_2$ is equivalent to $(\lambda x.e_2)\ e_1$. Therefore, we will continue using let-expressions in examples, but you should think of them as functions that have been immediately applied.

  As a brief exercise, you should try to prove that $\text{let } x = e_1 \text{ in } e_2 \Downarrow v$ if and only if $\lambda x.e_2\ e_1 \Downarrow v$.

The primary change in fig. 5.1 is that the semantics is defined by a ternary relation, where ρ (greek letter "rho") is the environment. An environment is a partial function that maps identifiers to values. Therefore, instead of using substitution, we have a new rule for identifiers (ID) that "looks up" identifiers in the environment. (Recall that we didn't have a rule for identifiers earlier, since they were eliminated by substitution.) The rule for constants (CONST) is updated to ignore ρ and the rules for operators and conditionals simply propagate ρ to their subexpressions.

The rules for functions (FUN) and applications (APP) are significantly different. To understand why, consider the following program:

$$\text{let } f = (\text{let } x = 20 \text{ in } \lambda y.x + y) \text{ in let } x = 10 \text{ in } f\ x$$

Using our old evaluation relation, we would substitute x with 20 and transform the named expression to $\lambda y.20 + y$, but leave the body unchanged, since it has another binding of x to 10. Therefore, we would effectively evaluate this program, where there is no scope for getting two x's mixed up:

$$\text{let } f = \lambda y.20 + y \text{ in let } x = 10 \text{ in } f\ x$$

If we use environments naively and simply evaluate the function to $\lambda y.x + y$ without substitution, we'll be in trouble. In the context where this function is applied x is bound to 10 instead of 20.

To solve this problem, functions are no longer values. Instead, functions evaluate to *closures*, which are triples $\langle \rho', x, e' \rangle$, where (1) ρ' is the environment in which the function was defined (i.e., in our example, the environment

Figure 5.1: Environment-based semantics. Contrast to fig. 4.2.

where x is correctly bound to 20), (2) x is the name of the function's argument, and (3) e' is the function body. In function application (the APP rule), we are careful to evaluate e' in the environment in which it was defined, and *not* the environment in which the application occurs.

Correctness   At this point, we have two ways of defining the semantics of an expression: either $e \Downarrow v$ or $\rho, e \Downarrow v$. The environment-based semantics is more efficient than the substitution-based semantics, but it *must* produce results that are equivalent to the original semantics. So, we may try to prove that both semantics map e to the same value v. However, the definition of v has changed: our old values had functions, but our new values have closures.

However, notice that:

$$\langle y \mapsto 10, x, x + y \rangle \approx \lambda x.x + 10$$

i.e., we can eliminate the environment in the closure by substituting its variables into the body. We could formalize this intuition to relate our old and new notions of values and prove that the two semantics are equivalent to each other.

Performance   Our motivation for introducing environments was that substitution is an expensive operation. But, are environments cheap? An obvious implementation strategy is to use a hash table to represent an environment, but this would not be effective e.g., in the APP rule, ρ is augmented with x to evaluate e_1 but e_2 is evaluated with just ρ. In addition, suppose a closure holds an environment ρ that is later extended. When the closure is eventually evaluated, it must be evaluated with the original environment (without any future extensions). Therefore, it seems as though an implementation would require a functional map data structure, which is not terrible efficient.

An efficient implementation would use a technique known as *de Brujin Indices*. de Brujin indices leverage the insight that we can always calculate the static distance between an identifier and the let-expression that introduced it (similarly for functions). Therefore, in this program:

$$\text{let } x = 10 \text{ in let } y = 10 \text{ in } x + y$$

We can replace y with v_0, since y is the innermost enclosing let-expression and x with v_1 since x is the next let-expression:

$$\text{let } \cdot = 10 \text{ in let } \cdot = 10 \text{ in } v_1 + v_0$$

The following example is a little trickier:

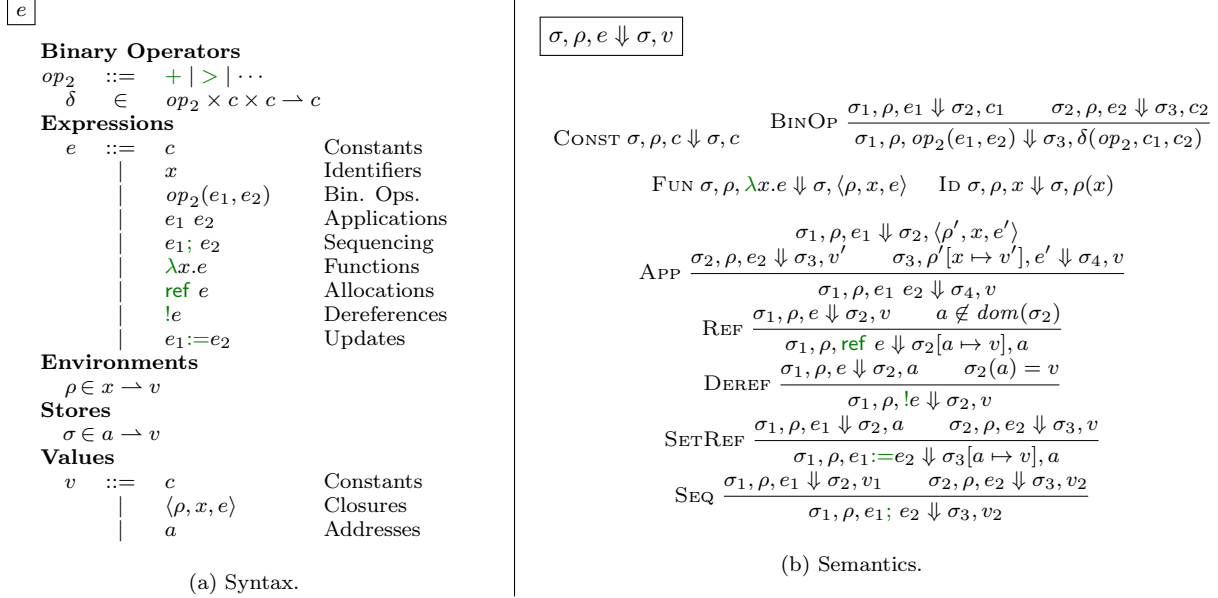$$\text{let } x = 10 \text{ in let } y = 5 + x \text{ in } x + y$$

$\boxed{e}$

Binary Operators

$op_2 \quad ::= \quad + \mid > \mid \cdots$

$\delta \quad \in \quad op_2 \times c \times c \rightharpoonup c$

Expressions

e	$::=$	c	Constants
	$\mid$	x	Identifiers
	$\mid$	$op_2(e_1, e_2)$	Bin. Ops.
	$\mid$	$e_1\ e_2$	Applications
	$\mid$	$e_1\, ; e_2$	Sequencing
	$\mid$	$\lambda x.e$	Functions
	$\mid$	$\mathsf{ref}\ e$	Allocations
	$\mid$	$!e$	Dereferences
	$\mid$	$e_1 := e_2$	Updates

Environments

$\rho \in x \rightharpoonup v$

Stores

$\sigma \in a \rightharpoonup v$

Values

v	$::=$	c	Constants
	$\mid$	$\langle \rho, x, e \rangle$	Closures
	$\mid$	a	Addresses

(a) Syntax.

$\boxed{\sigma, \rho, e \Downarrow \sigma, v}$

$$\text{CONST}\ \sigma, \rho, c \Downarrow \sigma, c \qquad \text{BINOP}\ \frac{\sigma_1, \rho, e_1 \Downarrow \sigma_2, c_1 \qquad \sigma_2, \rho, e_2 \Downarrow \sigma_3, c_2}{\sigma_1, \rho, op_2(e_1, e_2) \Downarrow \sigma_3, \delta(op_2, c_1, c_2)}$$

$$\text{FUN}\ \sigma, \rho, \lambda x.e \Downarrow \sigma, \langle \rho, x, e \rangle \qquad \text{ID}\ \sigma, \rho, x \Downarrow \sigma, \rho(x)$$

$$\text{APP}\ \frac{\begin{array}{c}\sigma_1, \rho, e_1 \Downarrow \sigma_2, \langle \rho', x, e' \rangle \\ \sigma_2, \rho, e_2 \Downarrow \sigma_3, v' \qquad \sigma_3, \rho'[x \mapsto v'], e' \Downarrow \sigma_4, v\end{array}}{\sigma_1, \rho, e_1\ e_2 \Downarrow \sigma_4, v}$$

$$\text{REF}\ \frac{\sigma_1, \rho, e \Downarrow \sigma_2, v \qquad a \notin dom(\sigma_2)}{\sigma_1, \rho, \mathsf{ref}\ e \Downarrow \sigma_2[a \mapsto v], a}$$

$$\text{DEREF}\ \frac{\sigma_1, \rho, e \Downarrow \sigma_2, a \qquad \sigma_2(a) = v}{\sigma_1, \rho, !e \Downarrow \sigma_2, v}$$

$$\text{SETREF}\ \frac{\sigma_1, \rho, e_1 \Downarrow \sigma_2, a \qquad \sigma_2, \rho, e_2 \Downarrow \sigma_3, v}{\sigma_1, \rho, e_1 := e_2 \Downarrow \sigma_3[a \mapsto v], a}$$

$$\text{SEQ}\ \frac{\sigma_1, \rho, e_1 \Downarrow \sigma_2, v_1 \qquad \sigma_2, \rho, e_2 \Downarrow \sigma_3, v_2}{\sigma_1, \rho, e_1\, ; e_2 \Downarrow \sigma_3, v_2}$$

(b) Semantics.

Figure 5.2: Mutable references.

This example has two references to x. The first x refers to the innermost let-expression v_0. However, the second x refers to the next let-expression, since y is v_0 in its context:

$$\mathsf{let}\ \cdot\ = 10\ \mathsf{in}\ \mathsf{let}\ \cdot\ = 5 + v_0\ \mathsf{in}\ v_1 + v_0$$

Since this technique gives us fixed indices for variables, we can use array-like data structures to lookup variables in constant time.

There are other issues to consider in a real implementation. For example, how long does it take to allocate a closure in memory? How much memory do closures consume? When can they be allocated on the stack? Are there cases where they don't need to be allocated at all?

2   Store-Passing

We've seen that some kinds of expressions can be rewritten to use other expressions in the language. For example, we've seen that let-expressions can be rewritten to function applications. Therefore, we can omit rules for let-expressions from the semantics. However, this form of rewriting is much harder to do for certain features, such as mutable. It is easier to change the semantics.

Figure 5.2 defines a language that supports mutable references in a style similar to OCaml. There are three main ideas here:

- An expression that creates a reference, such as $\mathsf{ref}\ 1+2$ is *not* a value. Instead, this expression stores the value of $1+2$ at a new address (a) in the store (σ). These addresses (or pointers, if you prefer) are a new kind of value.

- Certain expressions, such as allocation and dereferencing can change the contents of the store and these changes should then be visible for the rest of the execution of the program. Therefore, the store is threaded through the semantics: notice that the $\Downarrow$ relation produces a value and a new store.

- Although addresses (or pointers) are values in this language, they are opaque to the program: the only way to create a new pointer is to use ref and there is no support for pointer arithmetic as in C/C++. Therefore, the language is *memory safe*.