

## 1 Some thoughts about state

State refers to the ability to have mutation: a change in value over time. The functional languages we have studied so far, such as the  $\lambda$ -calculus and uML, have not had state, in the sense that once a value is created, it is impossible to change that value, and once a variable is bound it cannot be rebound to a new value (except by creating a new variable). Although state is not a necessary feature of a programming language (e.g.  $\lambda$ -calculus is Turing complete but does not have a notion of state), it is a feature of most languages in common use today. We distinguish between functional languages and imperative languages primarily on the basis of whether they are organized primarily around a stateless paradigm (with higher-order functions), or a stateful paradigm of mutable variables and data structures. It might be more useful, in fact, to distinguish between *stateless* and *stateful* languages, since stateful languages usually have functions too.

Why do many programmers like stateful programming languages? Perhaps because it matches how we think of the world. We view the world as being composed of entities (such as people or objects) with intrinsic unchanging identities. These entities change their state over time while maintaining the same underlying identity. If a programming language matches our mental model of the world, it's more comfortable for programmers to think about.

Thus, stateful languages always come equipped with an underlying notion of *identity*, such as variable identity or object identity. In these languages, the binding between identities and their current states can change over time. Stateful languages intrinsically have a notion of time, unlike pure functional languages such as uML or lambda calculus, which take an essentially timeless view of computation. In stateless languages, it doesn't matter exactly when something happens, because the outcome will be the same regardless. In stateful languages with multiple concurrent threads, it matters a great deal when something happens.

It can be argued that our perception of underlying identity is an illusion. In physics, identity does not exist: particles are identified by their state, not by any intrinsic identity. If two electrons in different states are swapped, no change has happened to the universe, because an electron is really the same thing as its state. Time is arguably an illusion as well: the universe is a multidimensional object in which physics says that time is much more symmetric to the other dimensions than we perceive. Being made of particles ourselves, we are only able to sense a single moment in time. Arguably, the universe is purely functional, and our minds deceive us. In fact, the observation that we are deceived into thinking that objects have an underlying identity dates back to Buddhist and even earlier Vedantic philosophy (the Upanishads).

The mismatch between reality and our perceptions becomes pragmatically relevant when we look at the difficulties in programming multiprocessor/multicore systems. Our consciousness is single-threaded, so it makes perfect sense to us to think about state evolving along a single timeline. But when with multiple simultaneous threads of computation, and communication delays, it is infeasible to give all threads a single, consistent view of the state evolution of the system. We find it very difficult to reason about how such multithreaded systems work and to write correct programs for them in a stateful language.

Designing workable programming languages for multicore systems is currently an active research topic. In sidestepping the intrinsic problems with state, a stateless language may be a better match for such systems than the stateful languages we commonly use!

## 2 First-class references

We extend uML with state, and call the new language uML<sub>1</sub>. We use a construct called a *reference cell* that allows a value to be rebound to different values over time. We can then use reference cells to model language features such as mutable variables. Reference cells are first-class values in this language, like in ML. The syntax of uML<sub>1</sub> is as follows:

|         |                    |  |
|---------|--------------------|--|
| $e ::=$ | $\dots$            |  |
|         | $\mathbf{ref} \ e$ | (create a reference cell with initial value from $e$ ) |
|         | $!e$               | (evaluation or dereference)                            |
|         | $e_1 := e_2$       | (assignment or mutation)                               |
|         | $e_1; e_2$         | (sequential composition)                               |

Informally, the meaning of these new expressions is as follows. Evaluating the expression  $\mathbf{ref} \ e$  creates a new reference cell having the initial value  $e$ . The expression  $!e$  evaluates to the current value of the reference cell  $e$ . The assignment  $e_1 := e_2$  assigns the value of  $e_2$  to the reference cell  $e_1$ , and in our semantics returns the value  $\mathbf{null}$ .

Other design choices could have been made. For example, the assignment operator could have returned the value of  $e_2$  so that expressions like  $e_0 := (e_1 := e_2)$  would make sense and could be used to assign the same value to multiple variables.

In this treatment, reference cells are first-class objects. This is more powerful than simple mutable variables as in C. It is also more dangerous, because it introduces *aliasing*, as the following example demonstrates:

```

let  $x = \mathbf{ref} \ 1$  in
  let  $y = x$  in
    let  $z = (x := 2)$  in
       $!y$ 

```

Here  $y$  is an alias for  $x$ . Dereferencing  $y$  gives us the value of  $x$ , which in this case has changed from 1 to 2. The expression therefore evaluates to 2. In other words, if you kick  $x$ ,  $y$  jumps!

Aliasing is a constant source of errors in stateful programming. A classic example is a procedure that is supposed to multiply two matrices  $a$  and  $b$  and write the result imperatively into a third matrix  $c$ . The programmer writes the obvious code, then later wants to set  $a$  to the product of  $a$  and  $b$ . This makes the matrix multiply compute something completely different from what was intended, and is very hard to debug!

## 2.1 Operational semantics

We now give an operational semantics for reference cell expressions in  $\mathbf{uML}_1$ . To do this, we define a *configuration* as a pair  $\langle e, \sigma \rangle$ , where  $e$  is an expression and  $\sigma$  is a function mapping reference cells to values. Reference cells are denoted generically by  $\ell$ . We extend the grammar of expressions and values to include reference cells.

|         |         |          |
|---------|---------|----------|
| $e ::=$ | $\dots$ | $  \ell$ |
| $v ::=$ | $\dots$ | $  \ell$ |

We inherit the existing reductions from  $\mathbf{uML}$ , lifting them to apply to  $\mathbf{uML}_1$  expressions. All these reductions have no side effects, so the state  $\sigma$  is unchanged by them:

$$\frac{e \rightarrow e'}{\langle e, \sigma \rangle \xrightarrow{\mathbf{uML}_1} \langle e', \sigma \rangle}$$

$$\begin{aligned}
\langle \mathbf{ref} \ v, \sigma \rangle &\longrightarrow \langle \ell, \sigma[\ell \mapsto v] \rangle && \text{(where } \ell \notin \text{dom}(\sigma) \text{)} \\
\langle \ell := v, \sigma \rangle &\longrightarrow \langle \mathbf{null}, \sigma[\ell \mapsto v] \rangle && \text{(where } \ell \in \text{dom}(\sigma) \text{)} \\
\langle !\ell, \sigma \rangle &\longrightarrow \langle \sigma(\ell), \sigma \rangle && \text{(where } \ell \in \text{dom}(\sigma) \text{)} \\
\langle \mathbf{null}; e, \sigma \rangle &\longrightarrow \langle e, \sigma \rangle
\end{aligned}$$

where  $\sigma[v \mapsto \ell]$  is the same function as  $\sigma$ , except that the value on input  $\ell$  is changed to  $v$ . The language has eager, left-to-right evaluation, so the evaluation contexts are given by this grammar:

$$E ::= \dots \mid \mathbf{ref} \ E \mid !E \mid E := e \mid \ell := E \mid E; e$$

The structural congruence rule is adapted to take into account the current state:

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

## 2.2 Translation semantics

$\llbracket e \rrbracket \rho \sigma$  is an expression that computes a pair  $(e', \sigma')$  in a naming environment  $\rho$  and state  $\sigma$ , where  $e'$  is the resulting expression and  $\sigma'$  is the store.

$$\begin{aligned} \llbracket n \rrbracket \rho \sigma &= (n, \sigma) \\ \llbracket x \rrbracket \rho \sigma &= (\rho \text{ ``}x\text{''}, \sigma) \\ \llbracket \text{if } e_b \text{ then } e_1 \text{ else } e_2 \rrbracket \rho \sigma &= \mathbf{let} (b, \sigma')^1 = \llbracket e_b \rrbracket \rho \sigma \mathbf{ in} \\ &\quad \mathbf{if } b \text{ then } \llbracket e_1 \rrbracket \rho \sigma' \mathbf{ else } \llbracket e_2 \rrbracket \rho \sigma' \\ \llbracket \text{ref } e \rrbracket \rho \sigma &= \mathbf{let} (x, \sigma') = \llbracket e \rrbracket \rho \sigma \mathbf{ in let } l = \text{MALLOC } \sigma' x \mathbf{ in } (l, \text{UPDATE } \sigma' l x) \\ \llbracket !e \rrbracket \rho \sigma &= \mathbf{let} (x, \sigma') = \llbracket e \rrbracket \rho \sigma \mathbf{ in } (\text{LOOKUP } \sigma' x, \sigma') \\ \llbracket e_1 := e_2 \rrbracket \rho \sigma &= \mathbf{let} (x_1, \sigma_1) = \llbracket e_1 \rrbracket \rho \sigma \mathbf{ in} \\ &\quad \mathbf{let} (x_2, \sigma_2) = \llbracket e_2 \rrbracket \rho \sigma_1 \mathbf{ in} \\ &\quad (\mathbf{null}, \text{UPDATE } \sigma_2 x_1 x_2) \\ \llbracket e_1; e_2 \rrbracket \rho \sigma &= \mathbf{let} (x, \sigma_1) = \llbracket e_1 \rrbracket \rho \sigma \mathbf{ in } \llbracket e_2 \rrbracket \rho \sigma_1 \\ \llbracket \mathbf{null} \rrbracket \rho \sigma &= (\mathbf{null}, \sigma) \\ \llbracket \mathbf{let } x = e_1 \mathbf{ in } e_2 \rrbracket \rho \sigma &= \mathbf{let} (y, \sigma') = \llbracket e_1 \rrbracket \rho \sigma \mathbf{ in } \llbracket e_2 \rrbracket (\text{EXTEND } \rho \text{ ``}x\text{'' } y) \sigma' \\ \llbracket \lambda x. e \rrbracket \rho_{lex} \sigma_{lex} &= \lambda y \sigma_{dyn}. \llbracket e \rrbracket (\text{EXTEND } \rho_{lex} \text{ ``}x\text{'' } y) \sigma_{dyn} \\ \llbracket e_1 e_2 \rrbracket \rho_{dyn} \sigma_{dyn} &= \mathbf{let} (f, \sigma_1) = \llbracket e_1 \rrbracket \rho_{dyn} \sigma_{dyn} \mathbf{ in} \\ &\quad \mathbf{let} (v, \sigma_2) = \llbracket e_2 \rrbracket \rho_{dyn} \sigma_1 \mathbf{ in} \\ &\quad f v \sigma_2 \end{aligned}$$

Here, we rely on certain abbreviations (*MALLOC*, *UPDATE*, and *LOOKUP*), for manipulating stores. The function *MALLOC* allocates a new memory location; *LOOKUP* reads the contents of a memory location; *UPDATE* produces a new store with one location changed. Further, we will need an initial store *EMPTY-STORE*. We don't care exactly how these terms are implemented, as long as they obey certain equations (the notion of equality here is equivalence after evaluation):

$$\begin{aligned} \text{LOOKUP } \text{EMPTY-STORE } \ell &= \mathbf{error} \\ \text{LOOKUP } \sigma (\text{MALLOC } \sigma) &= \mathbf{error} \\ \text{LOOKUP } (\text{UPDATE } \sigma \ell v) \ell &= v \\ \text{LOOKUP } (\text{UPDATE } \sigma \ell v) \ell' &= \text{LOOKUP } \sigma \ell' && (\text{if } \ell \neq \ell') \end{aligned}$$

Implementing terms with these properties is left as an exercise to the reader.

## 3 Mutable variables

Let's take a closer look at what happens in a language like C, where variables themselves are mutable.

### 3.1 Syntax

$$e ::= \dots \text{uML} \dots \mid e_1 = e_2 \mid *e \mid \&e \mid e_1; e_2$$

where we think of the constructs respectively as:

- $e_1 = e_2$  is assignment where  $e_1$  is an assignable variable name, for instance, an array reference.
- $*e$  is used to dereference a pointer  $e$  (acts sort of like ! in UML<sub>l</sub>).
- $\&e$  gives a pointer to the location of  $e$ .

### 3.2 Translation semantics

The target language is again uML. Since variables are mutable, we need to change the naming environment  $\rho$  to be a mapping from variable names to store locations. Each variable is bound to the identity of its reference cell.

The translation we define takes into account the fact that C has two kinds of values:

- lvalues, values which can appear in the left-hand side of an assignment; and
- rvalues, values that cannot appear on the LHS.

To capture this in the target language, for each expression  $e$  of C, we define  $\mathcal{L}\llbracket e \rrbracket \rho \sigma$  (the lvalue of  $e$ ) and  $\mathcal{R}\llbracket e \rrbracket \rho \sigma$  (the rvalue of  $e$ ). As before, we need to keep track of the state of the computation so the results of the translations are pairs. To summarize:

- $\mathcal{L}\llbracket e \rrbracket \rho \sigma = (\text{location of } e, \sigma')$
- $\mathcal{R}\llbracket e \rrbracket \rho \sigma = (\text{value of } e, \sigma')$ .

The translation is defined inductively on the structure of C programs. Note that some lvalue translations are not defined, corresponding to invalid expressions that a C compiler would catch. For instance,  $\mathcal{L}\llbracket n \rrbracket \rho \sigma$  is not defined. This prevents reassigning of, say, the integer 5 to have the value 3.

$$\begin{aligned} \mathcal{R}\llbracket n \rrbracket \rho \sigma &= (n, \sigma) \\ \mathcal{L}\llbracket x \rrbracket \rho \sigma &= (\rho \text{“}x\text{”}, \sigma) \\ \mathcal{R}\llbracket e \rrbracket \rho \sigma &= \mathbf{let} (\ell, \sigma') = \mathcal{L}\llbracket e \rrbracket \rho \sigma \mathbf{in} (\text{LOOKUP } \sigma' \ell, \sigma') \quad (\text{if } \mathcal{L}\llbracket e \rrbracket \text{ is defined}) \end{aligned}$$

For assignment:

$$\mathcal{R}\llbracket e_1 = e_2 \rrbracket \rho \sigma = \mathbf{let} (\ell, \sigma') = \mathcal{L}\llbracket e_1 \rrbracket \rho \sigma \mathbf{in} \mathbf{let} (v, \sigma'') = \mathcal{R}\llbracket e_2 \rrbracket \rho \sigma' \mathbf{in} (v, \text{UPDATE } \sigma'' \ell v)$$

Comments:

- This forces left-to-right evaluation.
- It allows a chain of assignments. For example,  $e_1 = e_2 = e_3 = e_4 = 5$  will assign 5 to each  $e_i$  in our translation, just as in C.
- An assignment statement is not an lvalue in C, nor is it here.

Addressing:

$$\begin{aligned}\mathcal{L}[\ast e]\rho\sigma &= \mathcal{R}[e]\rho\sigma \\ \mathcal{R}[\&e]\rho\sigma &= \mathcal{L}[e]\rho\sigma\end{aligned}$$

So the two operations,  $\ast$ ,  $\&$ , are inverses of each other; their function is to shift the view of a value rather than to do computation. For example, the semantics defined above let us write  $\ast\&\ast\&\ast\&x = 2$ . This will have the same effect as  $x = 2$ . Another example (written in C):

```
int y = 0;
int x = &y;
*x = 1;
```

This piece of code assigns 1 to  $y$ . Note that it would be illegal to include a line **&y = 0x1002**; because we can't change the address of a variable.

#### 4 Functions with mutable variables

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

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  $\ell$ . 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.

## 5 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:

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

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