

# CPSC 311: Definition of Programming Languages: State 20-state

Joshua Dunfield  
University of British Columbia

November 10, 2015

## 1 State

All of our Fun dialects have had only *immutable* bindings: During evaluation, once an identifier is bound to an expression, its meaning cannot change—it will have the same meaning as long as it is in scope (and not shadowed by another binding).

### 1.1 Classifying languages

Many languages have *mutable* state in some form:

- *By default, and idiomatic*: Fortran, Algol-60, Lisp, C, C++, Java, Smalltalk, . . .
- *By default, but less idiomatic*: Racket
- *Not by default*: Standard ML, OCaml
- *By simulation*: Haskell

The line between “functional” and “imperative” is fuzzy, but I think most people would draw it somewhere around Racket. The line between “purely functional” and “impurely functional”—*purity* meaning a “lack of side effects (such as state)”—is usually drawn between ML and Haskell. That line is also subject to debate, however.

Starting from the top of the list, in languages like Java, most features are mutable by default (unless `const` is given).

Racket occupies a strange position in this space: fundamental binding operations like **define** and **let** are mutable, but “good Racket style” discourages you from exploiting this. A few language features in Racket, including lists, are genuinely immutable by default (Racket also has mutable lists, but not by default; as we saw in our discussion of classifying languages, different languages often provide the same behaviours and differ only in which behaviour is the default).

The ML languages are fairly consistent in being immutable by default. A value bound by a `let` in SML or OCaml cannot be mutated. Both languages do have features similar to Racket’s boxes, but these must be used explicitly; the default is immutability. (An exception that puts OCaml slightly nearer the top of the page: strings are mutable.)

## §1 State

---

```
# let s = "abcd" ;;
val s : string = "abcd"
# String.set s 2 'r' ;;
- : unit = ()
# s ;;
- : string = "abrd"
```

Haskell is usually considered “pure” or “purely functional”, though there is debate about this too, partly because some people argue that nontermination is a side effect. In practice, Haskell has ample support (features like the appallingly named “monads”) for an imperative style of programming. (As an aside, the techniques used to *implement* Haskell are extremely imperative!)

### 1.2 Defining state

The particular form of state we’ll add to Fun is one you’re now (barely) familiar with: boxes. To avoid (and cause) confusion, I will use the ML terminology, *references* or *refs*: a reference is essentially a pointer to a *cell* (or “ref cell”) whose contents can be mutated.

Modelling boxes in a dynamic semantics requires significant changes. Now that we have an environment-based semantics, the environment *env* seems to be a logical place to store the current contents of ref cells. That turns out to be a bad idea—we want the state of ref cells to survive a lexical scope, but not the state of binders (see the question below)—so instead we’ll create yet another animal, a *store* *S*.

Stores  $S ::= \emptyset$       empty store  
          |  $\ell \triangleright v, S$     location  $\ell$  points to cell with contents  $v$ , followed by store  $S$

Let’s extend the concrete and abstract syntax.

Expressions  $\langle E \rangle ::= \dots$   
                  |  $\{\text{ref } \langle E \rangle\}$   
                  |  $\{\text{deref } \langle E \rangle\}$   
                  |  $\{\text{setref } \langle E \rangle \langle E \rangle\}$

The intended semantics is:

- $\{\text{ref } \langle E1 \rangle\}$  evaluates  $\langle E1 \rangle$  and returns the location of a new cell (think of a location as a pointer);
- $\{\text{deref } \langle E1 \rangle\}$  evaluates  $\langle E1 \rangle$ , which must evaluate to a location, and returns the contents of the cell at that location;
- $\{\text{setref } \langle E1 \rangle \langle E2 \rangle\}$  evaluates  $\langle E1 \rangle$ , which must evaluate to a location  $\ell$ , then evaluates  $\langle E2 \rangle$  and puts that resulting value into the cell at location  $\ell$ .

The abstract syntax is extended correspondingly:

```
(define-type E
  ...
  [ref (initial-contents E?)]
  [deref (loc-expr E?)])
```

## §1 State

---

```
[setref (loc-expr E?) (new-contents E?)]
```

We aren't quite done, though: a ref is how we *create* a new cell, not a *pointer* to a new cell. So we need one more variant:

```
(define-type E
  ...
  [ref (initial-contents E?)]
  [deref (loc-expr E?)]
  [setref (loc-expr E?) (new-contents E?)]
  [location (locsym symbol?)])
```

Racket has a built-in function called `gensym` that we'll use when we need a new location, to get a "fresh" symbol.

The point of a store is that its contents are mutable: evaluating an expression may change the store. So we need both an *input* store and an *output* store.

$env; S \vdash e \Downarrow v; S'$	Starting in environment $env$ and store $S$ , evaluating $e$ produces value $v$ and updated store $S'$
------------------------------------	--

$$\frac{env; S1 \vdash e \Downarrow v; S2}{env; S1 \vdash (ref\ e) \Downarrow (location\ \ell); \ell \triangleright v, S2} \text{??SEnv-ref}$$

What is  $\ell$ ? It really doesn't matter, as long as it isn't already in  $S1$ .

$$\frac{\ell \text{ fresh for } S1 \quad env; S1 \vdash e \Downarrow v; S2}{env; S1 \vdash (ref\ e) \Downarrow (location\ \ell); \ell \triangleright v, S2} \text{SEnv-ref}$$

When a new feature (like `ref`) leads us to introduce a new judgment form, we need to check two things:

- We can write rules for the new features.
- We can update the rules for old features.

We seem to have a rule for the new feature `ref`, so we should try to update an old rule. We'll do the rule for `pair`. We first need to update `Eval-pair` to use environments. Since pairs don't bind identifiers, this is straightforward:

$$\frac{env \vdash e1 \Downarrow v1 \quad env \vdash e2 \Downarrow v2}{env \vdash (pair\ e1\ e2) \Downarrow (pair\ v1\ v2)} \text{Env-pair (stateless version)}$$

$$\frac{env; S \vdash e1 \Downarrow v1; S1 \quad env; S1 \vdash e2 \Downarrow v2; S2}{env; S \vdash (pair\ e1\ e2) \Downarrow (pair\ v1\ v2); S2} \text{SEnv-pair}$$

This version of `Env-pair` says that, to evaluate a pair, we first evaluate  $e1$  under the given store  $S$ , producing a (possibly) changed store  $S1$ ; then, we evaluate  $e2$  under  $S1$ , producing another store  $S2$ , which is the store produced by the entire evaluation of `(pair  $e1$   $e2$ )`.

Following this pattern of passing stores along from premise to premise means that when we draw a derivation tree, and draw a line (really a curve) from the conclusion's starting store (to the left

## §1 State

of the turnstile  $\vdash$ ) to the conclusion's result (to the right of the semicolon), the line looks like a thread. The store is “threaded through” the derivation tree.

Unlike previous evaluation rules for `pair`, `SEnv-pair` specifies that an interpreter must evaluate the two expressions  $e_1$  and  $e_2$  in a particular order. The expression  $e_2$  *cannot* be evaluated before  $e_1$ , because we need to evaluate  $e_1$  to know what  $S_1$  is.

(Adding the full, tedious set of error-handling rules to the old big-step semantics, or to the environment-based semantics without state, would also enforce this order. But now, it is enforced within the non-error rule.)

Before we try to update all the old rules for this different evaluation judgment, we should make sure we can evaluate the other new features:

$$\frac{\ell \text{ fresh for } S_1 \quad \text{env}; S_1 \vdash e \Downarrow v; S_2}{\text{env}; S_1 \vdash (\text{ref } e) \Downarrow (\text{location } \ell); \ell \triangleright v, S_2} \text{SEnv-ref}$$

$$\frac{\text{env}; S_1 \vdash e \Downarrow (\text{location } \ell); S_2 \quad \text{lookup-loc}(S_2, \ell) = v}{\text{env}; S_1 \vdash (\text{deref } e) \Downarrow v; S_2} \text{SEnv-deref}$$

$$\frac{\text{env}; S \vdash e_1 \Downarrow (\text{location } \ell); S_1 \quad \text{env}; S_1 \vdash e_2 \Downarrow v_2; S_2 \quad \text{update-loc}(S_2, \ell, v_2) = S_3}{\text{env}; S \vdash (\text{setref } e_1 \ e_2) \Downarrow v_2; S_3} \text{SEnv-setref}$$

The idea of `update-loc` is that `update-loc`( $S_2, \ell, v_2$ ) =  $S_3$  where  $S_3$  is the same as  $S_2$ , but with  $\ell \triangleright \dots$  replaced by  $\ell \triangleright v_2$ .

**Question:** Why do we need a separate store? What goes wrong if we use the environment to store ref cells?

Because then we would end up with a similar problem as not doing a freeness check during substitution: we could access an identifier that should be out of scope.

$$e_{\text{Pair}} = \left( \text{pair } (\text{with } x \ (\text{num } 1) \ (\text{id } x)) \right. \\ \left. (\text{with } y \ (\text{num } 2) \ (\text{id } x)) \right)$$

Here, the second instance `(id x)` is not in the scope of `(with x ...)`. But (if the environment contains ref cells), we have to thread the environment through; otherwise, the second component of the following pair  $e_{\text{Pair}'}$  would be unable to see the effects of the first component. We expect  $e_{\text{Pair}'}$  to evaluate to `(pair (num 1) (num 1))` because the first component changes the contents of `r` from `(num 0)` to `(num 1)`:

$$e_{\text{Pair}'} = \left( \text{with } r \ (\text{ref } (\text{num } 0)) \ \left( \text{pair } (\text{with } x \ (\text{num } 1) \ (\text{setref } r \ (\text{num } 1))) \right. \right. \\ \left. \left. (\text{with } y \ (\text{num } 2) \ (\text{deref } r)) \right) \right)$$

If we thread the environment through, then in  $e_{\text{Pair}}$ , the binding of `(id x)` would survive and could be used outside its scope; if we don't thread the environment through,  $e_{\text{Pair}'}$  wouldn't behave as expected.

## §1 State

---

### 1.3 First implementation: `env-state.rkt`

We can **define-type** `Store` following the pattern of `Env`, and update `env-interp` to take *and* return a store. It's quite irritating, because Racket doesn't provide great support for returning pairs of things—I had to **define-type** `Res` to represent taking both an expression—the value  $v$  being returned in  $env; S1 \vdash e \Downarrow v; S2$ —and the output store  $S2$ . But it can be done, and there are no big surprises.

Some of this could be done more easily in ML or Haskell, because those languages have more general “pattern matching” than **type-case**, so adjacent **type-cases** can be combined in one. We still have to look up locations in the store, and update a cell by constructing a new store with different contents for that one cell.

■ **Question:** Racket has boxes! Why not just use those, instead of going to all this trouble?

Good question. One answer is that we want to know that our interpreter follows the rules (is “sound with respect to the rules”). The code is annoying to read, but the correspondence to the rules is clear. If we use Racket's boxes as our locations, the code becomes much simpler, but we have to trust that Racket's semantics for boxes matches our rules.

I'm pretty sure it does match the rules, which is why I wrote another version of the interpreter that *does* use Racket boxes (`env-state-direct.rkt`—the word “direct” refers to using Racket's boxes directly).

Another answer is that, while we can (I think!) use Racket's boxes to correctly represent Fun's refs, we are depending on Racket's idea of a store being the same as ours. What if we wanted to allow “time travel” in Fun, where we could “checkpoint” an old store and “rewind” to it later? Racket—as far as I know—doesn't have that feature. So we'd need to either figure out how to checkpoint Racket boxes, or use the `env-state.rkt` representation where we don't use boxes.

This leads us to...

### 1.4 Second implementation: `env-state-direct.rkt`

In this interpreter, the operations on the `Store` **define-type** are simply calls to Racket's `unbox` and `set-box!`. Instead of reflecting the “threading” of the store in our interpreter, we assume that Racket's store behaves in that same way. (Again, I'm pretty sure it does.)