

Program state refers to the ability to change program variables over time. We saw this feature in IMP (which was conspicuously missing functions) but then it went away in λ -calculus (which was conspicuously missing everything but functions). Even FL, which reintroduced most of the programming constructs, still did not have state. Once a variable was bound, there was no way to change its value as long as it was in scope.

State is not a necessary feature in a programming language—after all, λ -calculus is Turing complete despite no notion of state—but it is common in most languages and many programmers are accustomed to it.

Programming Paradigms. Two major programming paradigms are *functional* (stateless) and *imperative* (stateful). In a purely functional language, expressions resemble mathematical formulae. This structure allows programmers to reason equationally and not worry about where else a function is used or how many times something may be run. As a result, it avoids many of the pitfalls associated with a constantly changing execution environment.

Concurrency, in particular, is much simpler with functional programming. Confluence (the Church–Rosser property) means it does not matter which order operations execute, at the end the result will be the same. There can be no race conditions or inconsistent views of the world.

On the other hand, imperative programming more closely resembles the way we perceive the world and, especially, the operations happening inside of a real computer. There exists some underlying notion of state (of the world, of memory, of storage, etc), and that state can change over time.

1 Mutable References

Instead of allowing variable assignment directly, we will work with *mutable references* (aka *pointers*). These references can be updated in a way that cannot be handled by the simple substitution rules of their functional counterparts. They are somewhat more complicated than ordinary variable bindings because they introduce the extra complication of *aliasing*—the possibility of naming the same data value with different names.

For example, consider the following code.

```
let x = ref 1 in
let y = x in
  x := 2 ; !y
```

In this code, `ref 1` allocates a new reference pointing to the value 1 and returns the reference. So the first line creates this reference and assigns `x` to the reference. The second line assigns `y` to `x`, meaning both variables are bound to the same reference. Then we update the value pointed to by `x` to 2, and finally dereference `y` with `!y`. Because `x` and `y` point to the same place, modifying the value one points to also modifies the value the other points to, so this program will return 2. When you kick `x`, `y` jumps!

Reference should not be confused with mutable variables. A variable is *mutable* if its binding can change. The difference is subtle: variables are bound to values in an environment, and if the variable is mutable, it can be rebound to a different value. With references, the variable itself is bound to a *location*. The location is mutable—it can be rebound to a different value—but the variable itself is not. In IMP and imperative languages such as C, variables are typically mutable, whereas in functional languages such as FL, OCaml, and Haskell, they are typically not.

2 The FL! Language

To see how mutable references work in an otherwise-functional language, we will add them to the FL language we defined previously to create a new language FL!.¹ All FL expressions are also expressions of FL!, and there are a few more. Note that there is a set **Loc** of *memory locations* that we denote ℓ . The syntax is as follows.

$$e ::= \dots \mid \text{ref } e \mid e_1 := e_2 \mid !e \mid \ell \mid e_1 ; e_2$$

The “ $\dots$ ” is used to indicate that we are extending the BNF grammar from FL without having to write out everything again.

2.1 Small-Step Operational Semantics

Unlike in λ -calculus or functional FL, we can no longer define our semantics entirely by substitution.

Instead we will use a store to keep track of the mappings of memory locations for our references. This store is extremely similar to stores we saw previously in the semantics for IMP. It is a partial function $\sigma : \mathbf{Loc} \rightarrow \mathbf{Val}$ from memory locations to FL! values, and we require that it have a finite domain (that is, only finitely many location have mappings). We will again use the standard rebinding operator $\sigma[\ell \mapsto v]$. We now define the small-step operational semantics on *configurations*, $\langle e, \sigma \rangle$, just as we did for IMP.

We will again simplify our evaluation order rules by using evaluation contexts. Just as we extended the FL grammar for expressions, we extend the FL grammar for evaluation contexts.

$$E ::= \dots \mid \text{ref } E \mid E := e \mid v := E \mid !E \mid E ; e$$

Note that the hole $[\cdot]$ is already included with the $\dots$, so we do not need to write it again.

The operational semantic rule for using these contexts must change slightly. Before, our small step relation was over FL expressions, but now it is over entire configurations. Because of the nature of state, if the state changes when stepping an expression e , we need to retain that change when stepping e inside a larger context. The rule is therefore

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

We also need to add reduction rules for these new forms. Those rules are as follows.

$$\begin{array}{ll} \text{[NEW]} \frac{\ell \notin \text{dom}(\sigma)}{\langle \text{ref } v, \sigma \rangle \longrightarrow \langle \ell, \sigma[\ell \mapsto v] \rangle} & \text{[ASSIGN]} \frac{\ell \in \text{dom}(\sigma)}{\langle \ell := v, \sigma \rangle \longrightarrow \langle (), \sigma[\ell \mapsto v] \rangle} \\ \text{[DEREF]} \frac{\ell \in \text{dom}(\sigma)}{\langle !\ell, \sigma \rangle \longrightarrow \langle \sigma(\ell), \sigma \rangle} & \text{[SEQ]} \frac{}{\langle v ; e, \sigma \rangle \longrightarrow \langle e, \sigma \rangle} \end{array}$$

It can be shown that it is impossible to create dangling pointers in FL!. That is, if a program doesn't start with any pointers at all, then every pointer that appears will always be mapped in the store.

2.2 FL! in Action

To see how this operational semantics works, we look at the example from Section 1 and see how it executes.

$$\begin{aligned} \langle \text{let } x = \text{ref } 1 \text{ in let } y = x \text{ in } x := 2 ; !y, \emptyset \rangle &\longrightarrow \langle \text{let } x = \ell \text{ in let } y = x \text{ in } x := 2 ; !y, [\ell \mapsto 1] \rangle \\ &\longrightarrow \langle \text{let } y = \ell \text{ in } \ell := 2 ; !y, [\ell \mapsto 1] \rangle \\ &\longrightarrow \langle \ell := 2 ; !\ell, [\ell \mapsto 1] \rangle \end{aligned}$$

¹FL! is pronounced “FL bang.” The exclamation point, or *bang*, is a common symbol used to indicate a stateful operation in a primarily functional language.

We can now see that x and y have been substituted for the same location ℓ , and anything that happens to one of them will therefore be reflected in the other. In particular, this configuration continues to step as follows.

$$\langle \ell := 2 ; !\ell, [\ell \mapsto 1] \rangle \longrightarrow \langle () ; !\ell, [\ell \mapsto 2] \rangle \longrightarrow \langle !\ell, [\ell \mapsto 2] \rangle \longrightarrow \langle 2, [\ell \mapsto 2] \rangle$$

Exactly as expected, dereferencing y produced the value we assigned into x one operation earlier.

This behavior becomes particularly interesting when a function closes over a reference. For instance, consider the following code that produces a function that will ignore its argument and return how many times it has been called (including this one). The context of those calls is irrelevant, just the number of them.

let $cnt = \text{ref } 0$ in $(\lambda_. (cnt := !cnt + 1) ; !cnt)$

We could even do something weirder: we could build a pair of functions that each return how many times the *other* has been called.

let $cnt_1 = \text{ref } 0$ in
let $cnt_2 = \text{ref } 0$
in $((\lambda_. (cnt_1 := !cnt_1 + 1) ; !cnt_2),$
 $(\lambda_. (cnt_2 := !cnt_2 + 1) ; !cnt_1))$

Mutable references—and stateful computation in general—allows for some very odd behavior.

3 Translating FL! to FL

Even though we had to modify the structure of the small step operational semantic relation, this addition did not fundamentally add any computational power to the language. To show this, we will build an adequate translation from FL! to FL (though we will not prove its adequacy here).

To track the mutable state within the FL code, we construct an explicit representation of the heap that we call an *environment*, with the same type as a heap: a partial function with a finite domain from locations to values. We need a way to represent this environment in the target language of our translation, here FL. To do this, we will encode locations in the environment into values of the target language. We will write $\lceil \ell \rceil$ to denote the encoding of location ℓ into FL. The exact details of the encoding don't really matter as long as we can look up the value of a location given its encoding and update an environment with a new or modified binding. For instance, any structure that allows us to encode different locations differently and check equality of encoded variable names would work. Numbers are generally a good choice, but they are not the only one.

We do three things with heaps: allocate new locations, update existing locations, and loop up values, so we demand that if S is a representation of heap σ , the following three operations must exist.

1. $\text{alloc } S \ v = (\lceil \ell \rceil, S')$ where $\ell \notin \text{dom}(\sigma)$ and S' is a representation of $\sigma[\ell \mapsto v]$.
2. $\text{update } S \ \lceil \ell \rceil \ v$ is a representation of $\sigma[\ell \mapsto v]$
3. $\text{lookup } S \ \lceil \ell \rceil = \begin{cases} \sigma(\ell) & \text{if } \ell \in \text{dom}(\sigma) \\ \text{error} & \text{if } \ell \notin \text{dom}(\sigma) \end{cases}$

To simplify notation, despite a slight risk of confusion, we will use σ to represent both the heap and the representation of that heap in FL, so we will write $\text{alloc } \sigma \ v$, $\text{update } \sigma \ \lceil \ell \rceil \ v$, and $\text{lookup } \sigma \ \lceil \ell \rceil$.

The following translation maps FL! expressions to FL expression. Note that $\llbracket e \rrbracket$ represents a function that takes a store σ and produces an FL pair (v, σ') where v is an FL value and σ' is a store. These represent the output value (translated) and final store (representation) obtained by evaluating e . To simplify notation, we

will use the expression $\text{let } (x, \sigma') = \llbracket e \rrbracket \sigma \text{ in } e'$ as syntactic sugar for

$$\begin{aligned} & \text{let } p = \llbracket e \rrbracket \sigma \text{ in} \\ & \text{let } x = \#1 \ p \text{ in} \\ & \text{let } \sigma' = \#2 \ p \\ & \text{in } e' \end{aligned}$$

where $p \notin \text{FV}(e')$. Here is the translation for simple expressions. For simplicity of exposition, we skip tuples, and multi-argument functions, though those cases are not conceptually complicated.

$$\begin{aligned} \llbracket x \rrbracket \sigma &\triangleq (x, \sigma) \\ \llbracket n \rrbracket \sigma &\triangleq (n, \sigma) \\ \llbracket \text{true} \rrbracket \sigma &\triangleq (\text{true}, \sigma) \\ \llbracket \text{false} \rrbracket \sigma &\triangleq (\text{false}, \sigma) \\ \llbracket \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \rrbracket \sigma &\triangleq \text{let } (b, \sigma') = \llbracket e_0 \rrbracket \sigma \text{ in} \\ &\quad \text{if } b \text{ then } \llbracket e_1 \rrbracket \sigma' \text{ else } \llbracket e_2 \rrbracket \sigma' \\ \llbracket \lambda x. e \rrbracket \sigma &\triangleq (\lambda x \tau. \llbracket e \rrbracket \tau, \sigma) \\ \llbracket \text{rec } f(x) = e \rrbracket \sigma &\triangleq (\text{rec } f(x) = \lambda \tau. \llbracket e \rrbracket \tau, \sigma) \\ \llbracket e_1 \ e_2 \rrbracket \sigma &\triangleq \text{let } (f, \sigma_1) = \llbracket e_1 \rrbracket \sigma \text{ in} \\ &\quad \text{let } (x, \sigma_2) = \llbracket e_2 \rrbracket \sigma_1 \text{ in} \\ &\quad f \ x \ \sigma_2 \\ \llbracket \text{ref } e \rrbracket \sigma &\triangleq \text{let } (x, \sigma') = \llbracket e \rrbracket \sigma \text{ in } (\text{alloc } \sigma' \ x) \\ \llbracket e_1 := e_2 \rrbracket \sigma &\triangleq \text{let } (l, \sigma_1) = \llbracket e_1 \rrbracket \sigma \text{ in} \\ &\quad \text{let } (x, \sigma_2) = \llbracket e_2 \rrbracket \sigma_1 \text{ in} \\ &\quad ((), \text{update } \sigma_2 \ l \ x) \\ \llbracket !e \rrbracket \sigma &\triangleq \text{let } (l, \sigma') = \llbracket e \rrbracket \sigma \text{ in } ((\text{lookup } \sigma' \ l), \sigma') \\ \llbracket e_1 ; e_2 \rrbracket \sigma &\triangleq \text{let } (_, \sigma') = \llbracket e_1 \rrbracket \sigma \text{ in } \llbracket e_2 \rrbracket \sigma' \end{aligned}$$

Note that the translations of $\lambda x. e$ and $\text{rec } f(x) = e$ now take an extra argument for the new store that is current at the time of the call. That is because stores do not follow lexical scoping like variables, instead they follow a *dynamic scope* discipline whereby a location cares about the value of the store when it is used, not when it is defined. That is also why we need to track the store explicitly in the operational semantics, rather than doing using substitutions like we do for lexically-scoped variables.