Comp 411
Principles of Programming Languages
Lecture 15
Church and State: Supporting Assignment

Corky Cartwright

February 13, 2023

What Is Assignment?

- Assignment is rebinding (changing the value of) a variable in the current environment. This process is also called mutation since the environment is destructively changed.

- Nearly all practical programming languages include operations for mutating the values of program variables and data structures. Only plausible exception is Haskell, but is it really practical?

- To incorporate this feature in LC, we add an assignment operation to the language with syntax (taken from Scheme/Racket)

  ```
  (set! x M)
  ```

  and the abstract representation

  ```
  (define-struct (setter lhs rhs))
  ```

  where `x` is any lambda-bound identifier.

- Assignment (`set!`) enables us to model changing events in the real world.

How Do We Define the Semantics of Assignment Using a Meta-Interpreter?

Two common approaches:

1. Use mutation in the meta-language
2. Add another parameter to the eval function representing a *store* that maps *locations* to values. The *environment* maps assignable(mutable) variables [symbols] to locations. What is a *location*? An element of a specified denumerable set, typically the natural numbers (akin to machine addresses!). Such an interpreter is called *store-passing*. It has additional (a constant multiplicative factor) overhead

3.

Implications:

- Trade-offs: the second approach is pure but ugly and inefficient (if used in an actual implementation). It makes interpreters look like compilers where symbol lookup is performed on every access; identifiers stand for addresses.
- Conclusion: assignment is inherently ugly from a semantic perspective. Store-passing is used in denotational semantics for imperative (mutation supporting) languages but it gives no insight on how to build an efficient implementation or how to reason about programs written in such languages. Meta-interpreters that rely on mutation are much closer to efficient language implementations because all modern computers support mutation (updating the value at an address in memory!)

Two Different Formulations of Assignment

1. Assignable Variables
   - Mutate (change) bindings in environment.
   - The semantics of assignment critically depend on the fact that all extensions of an environment share that environment.
   - Racket Scheme/C++/Java/C#/Scala/Swift use this formulation.

2. Mutable Cells
   - Cells (boxes) are *values*. In essence, the data domain is augmented by a new unary constructor called **box** or **ref**. The *only* mutable data cells are these cells.
   - ML/Haskell uses this formulation.

Note: OO Characterization of Assignment
- Assignable variables are objects with two methods: a getter and a setter. In functional languages, immutable variables are objects only equipped only with getters.
- Mutable cells are value objects belonging to a class named **Box/Ref**; this form of value includes a setter. No other value does. Neither do variables.

Comparing the Two Different Formulations

Pros and Cons of Assignable Variables

- Simpler notation for common usage but the *cells/boxes* holding the values of variables typically are not conventional data values/objects which forces extra machinery (typically a prefix operator like **&** in C or a parameter attribute like **var**) in parameter passing (such as call-by-reference) and a distinction between *left-hand* and *right-hand* evaluation.
- If cells are data objects (*e.g.*, pointers as in C) the internals of the language implementation are exposed (as in C). This convention inhibits modular reasoning about program behavior.

Pros and Cons of Mutable Cells

- Mutable cells are simply a special form of data.
- The design of the language is unaffected otherwise.
- Simulating *call-by-reference* is trivial. The implementation of other parameter passing methods such as *call-by-name* and *call-by-need* are not directly affected by the addition of mutation.

Using Mutation to Define Mutation

Key intuition: implementing mutation in a language supporting mutation is easy, *provided* that environment sharing-relationships are modeled correctly. A nested environment shares its parent environment representation!

Observation: in the absence of a complex store-passing semantics, there is no straightforward way to support assignment if environments are represented as functions. Why? Assignment must update shared bindings but functions do not directly support sharing relationships. Linked lists (and other concrete mutable data structures) foster sharing!

To change the value associated with a variable **x**, we must bind a different value to the variable **x**. We can accomplish this by including a clause in the **eval** case-split of the form:

```
((setter? M) <change the environment>)
```

But how do we do this?

Using Mutation to Define Mutation cont.

- To make variables assignable, we need support modifying the values they stand for.

- Mutable variables cannot be directly associated with values; rather, they must be associated with an object which can be modified to hold a different value.

- What kind of object can we use?

- In Scheme, a particularly apt choice is to use a *box* (a built-in *struct* with a single field to hold the value of each variable. Then we can use mutation on Scheme boxes to change the value of the field. Note that we can also use closures and mutate local variables in these closures, which is an important trick featured in *SICP* (*Structure and Interpretation of Progams*).

- In a Java meta-interpreter, the value field in a **Binding** object simply has to be mutable.

- Moral: variables should stand for boxes (mutable cells).

- Comment: assignment languages like Java implicitly use boxes almost everywhere, but these boxes are *not* objects. They cannot be passed as values. They are not visible data values except via the Java reflection facility which is ugly and often breaks portability and backward compatibility.

Revising Our Meta-Interpreter

We must revise the clause that binds new variables (which in LC are only introduced in λ-expressions):

```
((app? M)
 (apply
   (eval (app-rator M) env)      ;; head of the application
   (box (eval (app-rand M) env)))) ;; box is a constructor
```

Since variables are now bound to boxes containing values, we must change the code that for evaluating variables:

```
((var? M)
 (unbox (lookup (var-name M) env))) M)
```

We are finally ready to add the clause for assignment:

```
((setter? M)
 (set-box! (lookup (setter-lhs M) env)
           (eval (setter-rhs M) env)))
```

Can Boxes Be Values?

- Yes. Many languages support some formulation of this concept. But the details can be delicate because we must know from context whether a variable **x** means either its value or the enclosing box. In ML, it is trivial.

- Traditional *context-based* approach: support references by distinguishing *left-hand* and *right-hand* contexts and using a different definition of evaluation in these two situations. In a meta-interpreter for such a language, *left-hand* evaluation is primary and has a simple recursive definition. In such an interpreter, *right-hand* evaluation has a trivial definition; it simply extracts the contents of a cell when *left-hand* evaluation yields a cell; otherwise it is the identity. In practice, right-hand-evaluation is far more common because the set of left-hand evaluation sites is sparse. (The canonical example is simple assignment, *e.g.*, in **x = x + 1** the first occurrence of **x** is *left-hand* evaluated and the second is *right-hand* evaluated. In function/procedure/method declarations, attaching a *reference* or **var** attribute to a parameter stipulates that the corresponding argument expression is *left-hand* evaluated. Examples: **var** parameters in Pascal, reference (tagged with **&**) parameters in C++. Arguments passed by reference are interpreted differently using *left-hand* evaluation rather than *right-hand* evaluation. Ordinary values typically cannot be cells/boxes (check out Algol 68 for insight).

- Cleaner *comprehensive* approach: treat boxes as ordinary values (as in ML) or, in lower-level languages, pointers as ordinary values (as in C). But there is a conceptual cost: these boxes/pointers must be explicitly dereferenced to get the associated values. In C, the data model is ultimately machine memory and explicit use of pointers is perilous; tiny mistakes can cause catastrophic behavior (screen of blue death [SOBD]) In contrast, the ML convention is much simpler than the traditional *context-based* approach, *but* it requires explicit dereferencing (using the unary prefix operator **!**) to designate the value of a mutable variable. Type-checking detects nearly all dereferencing failures, but code involving mutation looks messy.