Programming Languages and Techniques (CIS120)

Lecture 13 – Part 2

Mutable State, Aliasing, and the Abstract Stack Machine Chapters 14 & 15

Mutable State

Mutable Record Fields

- By default, all record fields are *immutable*—once initialized, they can never be modified.
- OCaml supports *mutable* fields that can be imperatively updated by the "set" command: record.field <- val

note the 'mutable' keyword

```
type point = {mutable x:int; mutable y:int}
let p0 = {x=0; y=0}
(* set the x coord of p0 to 17 *)
;; p0.x <- 17
;; print_endline ("p0.x = " ^ (string_of_int p0.x))
p0.x = 17</pre>
```

Record Update

- Functions can assign to mutable record fields
- Note that the return type of '<-' is unit
 - i.e., it is a command

```
type point = {mutable x:int; mutable y:int}
(* a command to shift a point by dx,dy *)
let shift (p:point) (dx:int) (dy:int) : unit =
    p.x <- p.x + dx;
    p.y <- p.y + dy</pre>
```

- Note that the result type of shift is also unit
 - i.e., shift is a user-defined command

Why Use Mutable State?

- Action at a distance
 - allow remote parts of a program to communicate / share information without threading the information through all the points in between
- Data structures with explicit sharing
 - e.g. graphs
 - without mutation, it is only possible to build trees – no cycles
- Efficiency/Performance
 - A few data structures have imperative implementations with better asymptotic efficiency than the best declarative version
- Re-using space (in-place update)
- Random-access data (arrays)
- Direct manipulation of hardware
 - device drivers, displays, etc.



Different views of imperative programming

Java (and C, C++, C#)

- Code is a sequence of statements (a.k.a. commands) that do something, sometimes using expressions to compute values.
- References are **mutable** by default, must be explicitly declared to be constant

OCaml (and Haskell, etc.)

- Code is an expression that has a value. Sometimes computing that value has other effects.
- References are immutable by default, must be explicitly declared to be mutable

```
What answer does the following function produce when called?
     type point = {mutable x:int; mutable y:int}
     let f (p1:point) : int =
       p1.x <- 17;
       p1.x
1. 17
2. something else
3. sometimes 17 and sometimes something else
```

4. f is ill typed

ANSWER: 1



ANSWER: 3

The Challenge of Mutable State: Aliasing

What does this function return?

```
let f (p1:point) (p2:point) : int =
    p1.x <- 17;
    p2.x <- 42;
    p1.x</pre>
```

```
(* Consider this call to f: *)
let p0 = {x=0; y=0} in
    f p0 p0
```

Two identifiers are said to be *aliases* if they both name the *same* mutable record. Inside f, the identifiers p1 and p2 might or might not be aliased, depending on which arguments are passed in.

SEE THE COURSE NOTES FOR MORE ON THIS EXAMPLE

Opening a Whole New Can of Worms*



*t-shirt courtesy of ahrefs.com

The Challenge of Mutable State: Aliasing

What does this function return?

```
let f (p1:point) (p2:point) : int =
    p1.x <- 17;
    p2.x <- 42;
    p1.x</pre>
```

```
(* Consider this call to f: *)
let p0 = {x=0; y=0} in
    f p0 p0
```

Two identifiers are said to be *aliases* if they both name the *same* mutable record. Inside f, the identifiers p1 and p2 might or might not be aliased, depending on which arguments are passed in.

SEE THE COURSE NOTES (ch 14) FOR MORE ON THIS EXAMPLE

The Abstract State Machine

Location, Location, Location!

We need a new Computation Model

- The simple substitution model works well for value-oriented programming
 - "Observable" behavior of a value is *completely* determined by its structure
 - Pure functions are *referentially transparent*: two different calls to the same function with the same arguments yield the same results
 - These properties justify "replace equals by equals" reasoning



- With mutable state...
 - The *location* of values matters, not just their structure
 - Results returned by functions are not fully determined by their arguments (can also depend on "hidden" mutable state)

Abstract Stack Machine

Three "spaces"

- workspace
 - the expression the computer is currently simplifying
- stack
 - temporary storage for local variables and saved work
- heap
 - storage area for large data structures

Workspace	Stack	Неар

Abstract stack machine

Abstract Stack Machine

Initial state:

- workspace contains whole program
- stack and heap are empty

Machine operation:

- In each step, choose "next part" of the workspace expression and simplify it
- (Sometimes this will change the stack and/or heap)
- Stop when there are no more simplifications to be done

Workspace	Stack	Неар

Abstract stack machine

Values and References

A *value* is either:

- a *primitive value* like an integer, or,
- a *reference* to a location in the heap
- A reference value is the *address (location)* of data in the heap. We draw a reference value as an "arrow"
 - The arrow "points" to a box or cell located at this address
 - Where we are storing this value also matters:



References as an Abstraction

- In a real computer, the memory consists of an array of 32-bit words, numbered 0 ... 2³²-1 (for a 32-bit machine)
 - A reference is just an address that tells you where to look up a value
 - Data structures are usually laid out in contiguous blocks of memory
 - Constructor tags are just numbers chosen by the compiler
 e.g. Nil = 42 and Cons = 120120120



The ASM: Simplifying variables, operators, let expressions, and if expressions









Workspace

let y = 2 + x in if x > 23 then 3 else 4







let y = 2 + 22 in
 if x > 23 then 3 else 4





let $y = \frac{2 + 22}{10}$ in if x > 23 then 3 else 4



Workspace

let y = 24 in
 if x > 23 then 3 else 4



Workspace

 $\frac{\text{let } y = 24 \text{ in}}{\text{if } x > 23 \text{ then } 3 \text{ else } 4}$



Workspace

if x > 23 then 3 else 4



WorkspaceStackHeapif $\underline{x} > 23$ then 3 else 4y 24y 24yyyy

Looking up x in the stack proceeds from most recent entries to the least recent entries. Note that the "top" (most recent part) of the stack is drawn toward the bottom of the diagram.

Workspace

if 22 > 23 then 3 else 4



Workspace

if 22 > 23 then 3 else 4



Workspace

if false then 3 else 4



Workspace

if false then 3 else 4





What does the <u>Stack</u> look like after simplifying the following code on the workspace?

let z = 20 in let w = 2 + z in w



What does the <u>Stack</u> look like after simplifying the following code on the workspace?

let z = 20 in let z = 2 + z in z



Mutable Records

- The reason for introducing the ASM model is to make heap locations and sharing *explicit*.
 - Now we can say what it means to mutate a heap value *in place*.

```
type point = {mutable x:int; mutable y:int}
let p1 : point = {x=1; y=1;}
let p2 : point = p1
let ans : int = (p2.x <- 17; p1.x)</pre>
```

- We draw a record in the heap like this:
 - The doubled outlines indicate that those cells are mutable



A point record in the heap.

- Everything else is immutable

Allocate a Record













Note: p1 and p2 are references to the *same* heap record. They are *aliases* – two different names for the *same thing*.

Look Up 'p2'

Look Up 'p2'

Assign to x field

Assign to x field

This is the step in which the 'imperative' update occurs. The mutable field x has been modified in place to contain the value 17.

Sequence ';' Discards Unit

Look Up 'p1'

Look Up 'p1'

Project the 'x' field

Project the 'x' field

Let Expression

Push ans

Answer: 1

What do the <u>Stack</u> and <u>Heap</u> look like after simplifying the following code on the workspace?

Answer: 1