CMSC 330: Organization of Programming Languages

Objects vs. Functional Programming

OOP vs. FP

- Object-oriented programming (OOP)
  - Computation as interactions between objects
  - Objects encapsulate mutable data (state)
    - Accessed / modified via object's public methods
- Functional programming (FP)
  - Computation as evaluation of functions
    - Mutable data used to improve efficiency
  - Higher-order functions implemented as closures
    - Closure = function + environment

An Integer "Stack" Abstraction in Java

```
class Stack {
  class Node {
    Integer val; Node next;
    Node(Integer v, Node n) { val = v; next = n; }
  };
  private Node theStack;
  void push(Integer v) {
    theStack = new Node(v, theStack);
  }
  Integer pop() {
    if (theStack == null)
      throw new NoSuchElementException();
    Integer temp = theStack.val;
    theStack = theStack.next;
    return temp;
  }
}
```


A "Stack" Abstraction in OCaml

```
module type STACK =
  sig
    type 'a stack
    val new_stack : unit -> 'a stack
    val push : 'a stack -> 'a -> unit
    val pop : 'a stack -> 'a
  end

module Stack : STACK =
  struct
    type 'a stack = 'a list ref
    let new_stack () = ref []
    let push s x = s := (x::!s)
    let pop s = match !s with
      [] -> failwith "Empty stack"
    | (h::t) -> s := t; h
  end
```


Another "Stack" Abstraction in OCaml

```
let new_stack () =
  let this = ref [] in
    let push x = this := (x::!this)
    and pop () = match !this with
        [] -> failwith "Empty stack"
      | (h::t) -> this := t; h
    in
    (push, pop)
```

```
# let s = new_stack ();;
val s : ('_a -> unit) * (unit -> '_a) = (<fun>, <fun>)
# Pervasives.fst s 3;;  (* applies 1st part of s to 3  *)
- : unit = ()
# Pervasives.snd s ();; (* applies 2nd part of s to () *)
- : int = 3
```


Two OCaml Stack Implementations

- 1st implementation (OOP style)
  - Based on modules
  - Specifies methods for
    - Creating stack
    - Pushing value onto stack parameter
    - Popping value from stack parameter
- 2nd implementation (FP style)
  - Based on closures
  - Creating stack returns tuple containing
    - Closure for pushing value onto created stack
    - Closure for popping value from created stack

Relating Objects and Closures

- An object...
  - Is a collection of fields (data)
  - ...and methods (code)
  - When a method is invoked
    - Method has implicit this parameter that can be used to access fields of object
- A closure...
  - Is a pointer to an environment (data)
  - ...and a function body (code)
  - When a closure is invoked
    - Function has implicit environment that can be used to access variables

Relating Objects and Closures (cont.)

```
class C {
  int x = 0;
  void set_x(int y) { x = y; }
  int get_x() { return x; }
}
```

```
let make () =
  let x = ref 0 in
    ( (fun y -> x := y),
      (fun () -> !x) )
```

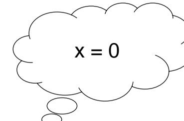

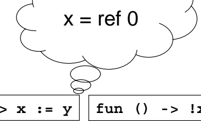


```
fun y -> x := y
```

```
fun () -> !x
```

```
C c = new C();
c.set_x(3);
int y = c.get_x();
```

```
let (set, get) = make ();;
set 3;;
let y = get ();;
```


Encoding Objects with Functions

- We can apply this transformation in general

```
class C { f1 ... fn; m1 ... mn; }
```

  - becomes

```
let make () =
 let f1 = ...
 ...
 and fn = ... in
 ( fun ... , (* body of m1 *)
   ...
   fun ...,   (* body of mn *)
 )
```

Tuple containing closures

  - make ( ) is like the constructor
  - The closure environment contains the fields

Recall a Useful Higher-Order Function

```
let rec map f = function
   [] -> []
 | (h::t) -> (f h)::(map f t)
```

- Map applies an arbitrary function f
  - To each element of a list
  - And returns the resulting modified list
- Can we encode this in Java?
  - Using object oriented programming

A Map Method for Stack

- Problem – Write a map method in Java
  - Must pass a function into another function
- Solution
  - Can be done using an object with a known method
  - Use interface to specify what method must be present

```
public interface Function {
  Integer eval(Integer arg);
}
```


A Map Method for Stack (cont.)

- Examples
  - Two classes which both implement Function interface

```
class AddOne implements Function {
  Integer eval(Integer arg) {
    return new Integer(arg + 1);
  }
}
```

```
class MultTwo implements Function {
  Integer eval(Integer arg) {
    return new Integer(arg * 2);
  }
}
```

The New Stack Class

```
class Stack {
  class Node {
    Integer val; Node next;
    Node (Integer v, Node n) { val = v; next = n; }
    Entry map(Function f) {
      if (next == null)
        return new Node(f.eval(val), null);
      else return new Node(f.eval(val), next.map(f));
    }
  }
  Node theStack;
  ...
  Stack map(Function f) {
    Stack s = new Stack();
    s.theStack = theStack.map(f);
    return s;
  }
}
```


Applying Map To A Stack

- Then to apply the function, we just do

```
Stack s = ...;
Stack t = s.map(new AddOne());
Stack u = s.map(new MultTwo());
```

  - We make a new object
    - That has a method that performs the function we want
  - This is sometimes called a callback
    - Because map "calls back" to the object passed into it
  - But it's really just a higher-order function
    - Written more awkwardly

Relating Closures and Objects

```
let app f x = f x
```

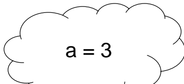


```
fun b -> a + b
```

```
let add a b = a + b;;
let f = add 3;;
app f 4;;
```

```
interface F {
  Integer eval(Integer y);
}
class C {
  static Integer app(F f, Integer x) {
    return f.eval(x);
  }
}
```

```
class G implements F {
  Integer a;
  G(Integer a) { this.a = a; }
  Integer eval(Integer y) {
    return new Integer(a + y);
  }
}
```

```
F adder = new G(3);
C.app(adder, 4);
```

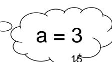

Encoding Functions with Objects

- We can apply this transformation in general

```
...(fun x -> (* body of fn *)) ...
let h f ... = ...f y...
```

  - becomes

```
interface F { Object eval(Object x); }
class G implements F {
  Object eval(Object x) { /* body of fn */ }
}
class C {
  Typ h(F f, ...) {
    ...f.eval(y)...
  }
}
```

  - F is the interface to the callback
  - G represents the particular function

Code as Data

- Closures and objects are related
  - Both of them allow
    - Data to be associated with higher-order code
    - Pass code around the program
- The key insight in all of these examples
  - Treat code as if it were data
    - Allowing code to be passed around the program
    - And invoked where it is needed (as callback)
- Approach depends on programming language
  - Higher-order functions (OCaml, Ruby, Lisp)
  - Function pointers (C, C++)
  - Objects with known methods (Java)

Code as Data (cont.)

- This is a powerful programming technique
  - Solves a number of problems quite elegantly
    - Create new control structures (e.g., Ruby iterators)
    - Add operations to data structures (e.g., visitor pattern)
    - Event-driven programming (e.g., observer pattern)
  - Keeps code separate
    - Clean division between higher & lower-level code
  - Promotes code reuse
    - Lower-level code supports different callbacks

CMSC 330: Organization of Programming Languages

Garbage Collection

Memory Attributes

- Memory to store data in programming languages has several attributes
  - Persistence (or lifetime)
    - How long the memory exists
  - Allocation
    - When the memory is available for use
  - Recovery
    - When the system recovers the memory for reuse

Memory Attributes (cont.)

- Most programming languages are concerned with some subset of the following 4 memory classes
  1. Fixed (or static) memory
  2. Automatic memory
  3. Programmer allocated memory
  4. Persistent memory

Memory Classes

- Static memory – Usually a fixed address in memory
  - Persistence – Lifetime of execution of program
  - Allocation – By compiler for entire execution
  - Recovery – By system when program terminates
- Automatic memory – Usually on a stack
  - Persistence – Lifetime of method using that data
  - Allocation – When method is invoked
  - Recovery – When method terminates

Memory Classes (cont.)

- Allocated memory – Usually memory on a heap
  - Persistence – As long as memory is needed
  - Allocation – Explicitly by programmer
  - Recovery – Either by programmer or automatically (when possible and depends upon language)

Memory Classes (cont.)

- Persistent memory – Usually the file system
  - Persistence – Multiple execution of a program (e.g., files or databases)
  - Allocation – By program or user, often outside of program execution
  - Recovery – When data no longer needed
  - Note
    - Dealing with persistent memory → databases (CMSC 424)

Memory Management in C

- Local variables live on the stack
  - Allocated at function invocation time
  - Deallocated when function returns
  - Storage space reused after function returns
- Space on the heap allocated with malloc()
  - Must be explicitly freed with free()
  - Called explicit or manual memory management
    - Deletions must be done by the user

Memory Management Mistakes

- May forget to free memory (memory leak)

```
{ int *x = (int *) malloc(sizeof(int)); }
```

- May retain ptr to freed memory (dangling pointer)

```
{ int *x = ...malloc();
  free(x);
  *x = 5; /* oops! */
}
```

- May try to free something twice

```
{ int *x = ...malloc(); free(x); free(x); }
```

  - This may corrupt the memory management data structures
    - E.g., the memory allocator maintains a free list of space on the heap that's available

Ways to Avoid Mistakes

- Don't allocate memory on the heap
  - Often impractical
  - Leads to confusing code (e.g., alloca())
- Never free memory
  - OS will reclaim process's memory anyway at exit
  - Memory is cheap; who cares about a little leak?
  - LISP model – System halts program and reclaims unused memory when there is no more available
- Use a garbage collector
  - E.g., conservative Boehm-Weiser collector for C

Memory Management in Ruby

- Local variables live on the stack
  - Storage reclaimed when method returns
- Objects live on the heap
  - Created with calls to Class.new
- Objects never explicitly freed
  - Ruby uses automatic memory management
    - Uses a garbage collector to reclaim memory

Memory Management in OCaml

- Local variables live on the stack
- Tuples, closures, and constructed types live on the heap
  - `let x = (3, 4)` (* heap-allocated *)
  - `let f x y = x + y in f 3` (* result heap-allocated *)
  - `type 'a t = None | Some of 'a`
  - `None` (* not on the heap–just a primitive *)
  - `Some 37` (* heap-allocated *)
- Garbage collection reclaims memory

Memory Management in Java

- Local variables live on the stack
  - Allocated at method invocation time
  - Deallocated when method returns
- Other data lives on the heap
  - Memory is allocated with new
  - But never explicitly deallocated
    - Java uses automatic memory management

Fragmentation

- Another memory management problem
- Example sequence of calls
  - allocate(a);
  - allocate(x);
  - allocate(y);
  - free(a);
  - allocate(z);
  - free(y);
  - allocate(b);
  - ⇒ Not enough contiguous space for b

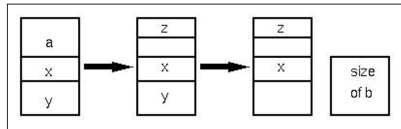

Garbage Collection Goal

- Process to reclaim memory
  - Also solve fragmentation

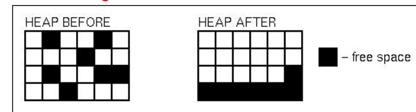

- Algorithm: You can do garbage collection and memory compaction if you know where every pointer is in a program. If you move the allocated storage, simply change the pointer to it.
- This is true in LISP, OCAML, Java, Prolog
- Not true in C, C++, Pascal, Ada

Garbage Collection (GC)

- At any point during execution, can divide the objects in the heap into two classes
  - Live objects will be used later
  - Dead objects will never be used again
    - They are "garbage"
- Idea
  - Can reuse memory from dead objects (recycling!)
- Goals
  - Reduce memory leaks
  - Make dangling pointers impossible

Many GC Techniques

- In most languages we can't know for sure which objects are really live or dead
  - Undecidable, like solving the halting problem
- Thus we need to make an approximation
  - OK if we decide something is live when it's not
  - But we'd better not deallocate an object that will be used later on

Reachability

- An object is reachable if it can be accessed by chasing pointers from live data
- Safe policy: delete unreachable objects
  - An unreachable object can never be accessed again by the program
    - The object is definitely garbage
  - A reachable object may be accessed in the future
    - The object could be garbage but will be retained anyway
    - Could lead to memory leaks

Roots

- At a given program point, we define liveness as being data reachable from the root set
  - Global variables
    - What are these in Java? Ruby? OCaml?
  - Local variables of all live method activations
    - I.e., the stack
- At the machine level
  - Also consider the register set
    - Usually stores local or global variables
- Next
  - Techniques for pointer chasing

Reference Counting

- Old technique (1960)
- Each object has count of number of pointers to it from other objects and from the stack
  - When count reaches 0, object can be deallocated
- Counts tracked by either compiler or manually
- To find pointers, need to know layout of objects
  - In particular, need to distinguish pointers from ints
- Method works mostly for reclaiming memory
  - Doesn't handle fragmentation problem

Reference Counting Example

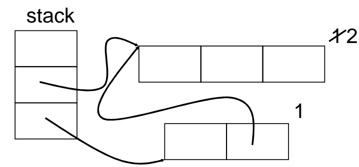

Reference Counting Example (cont.)

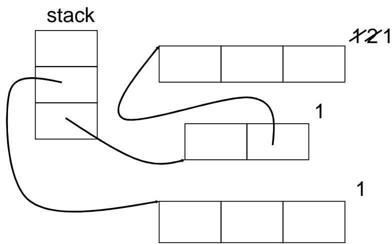

Reference Counting Example (cont.)

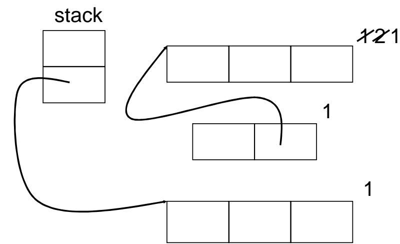

Reference Counting Example (cont.)

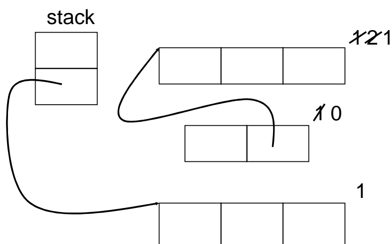

Reference Counting Example (cont.)

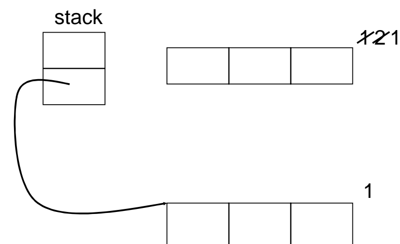

Reference Counting Example (cont.)

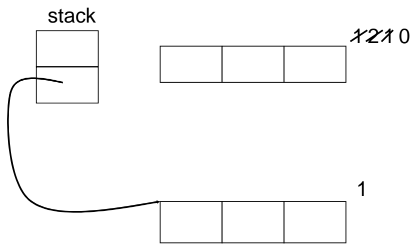

Reference Counting Example (cont.)

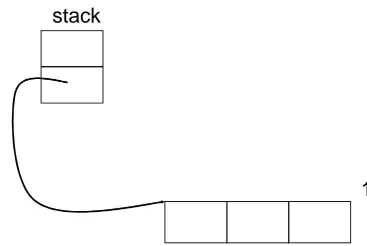

Reference Counting Tradeoffs

- Advantage
  - Incremental technique
    - Generally small, constant amount of work per memory write
    - With more effort, can even bound running time
- Disadvantages
  - Cascading decrements can be expensive
  - Requires extra storage for reference counts
  - Can't collect cycles, since counts never go to 0

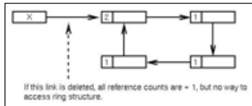

Mark and Sweep GC

- Idea
  - Only objects reachable from stack can be live
- Every so often, stop the world and do GC
  - Mark all objects on stack as live
  - Mark object reachable from live object as live
    - Repeat until no more reachable objects
  - Deallocate any non-reachable objects
- This is a tracing garbage collector
  - Does not handle fragmentation problem

Mark and Sweep Example

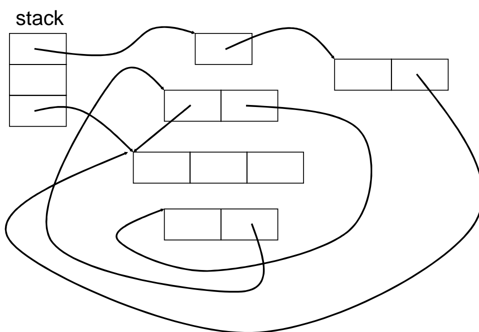

Mark and Sweep Example (cont.)

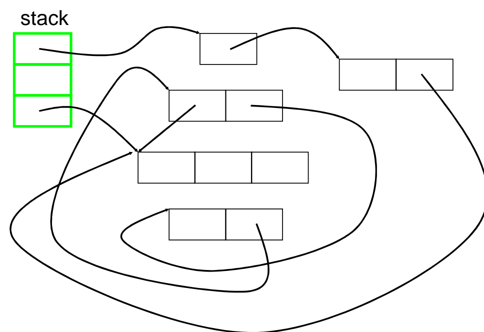

Mark and Sweep Example (cont.)

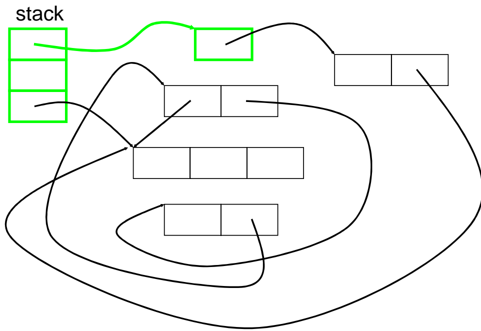

Mark and Sweep Example (cont.)

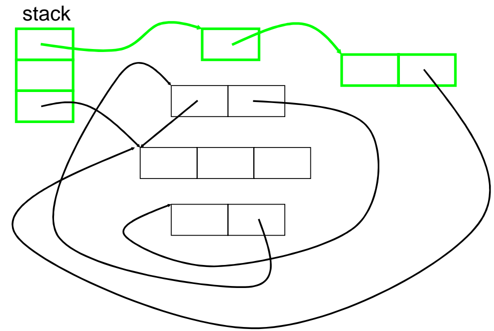

Mark and Sweep Example (cont.)

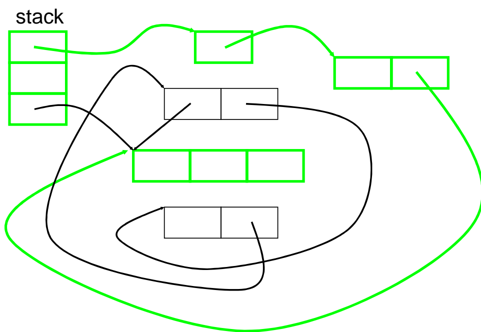

Mark and Sweep Example (cont.)

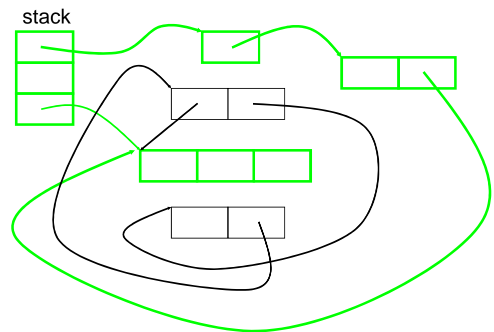

Mark and Sweep Example (cont.)

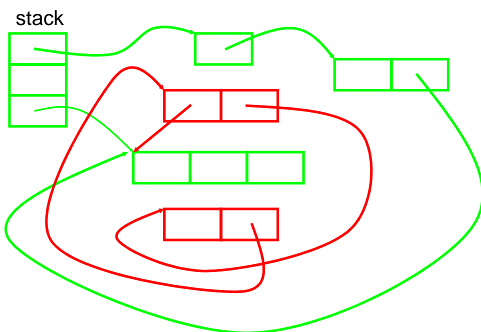

Mark and Sweep Tradeoffs

- Advantages
  - No problem with cycles
  - Memory writes have no cost

Mark and Sweep Tradeoffs (cont.)

- Disadvantages
  - Fragmentation
    - Available space broken up into many small pieces
      - Thus many mark-and-sweep systems may also have a compaction phase (like defragmenting your disk)
  - Cost proportional to heap size
    - Sweep phase needs to traverse whole heap – it touches dead memory to put it back on to the free list
  - Not appropriate for real-time applications
    - Bad if your car's braking system performs GC while you are trying to stop at a busy intersection

Stop and Copy GC

- Like mark and sweep, but only touches live objects
  - Divide heap into two equal parts (semispaces)
  - Only one semispace active at a time
  - At GC time, flip semispaces
    1. Trace the live data starting from the stack
    2. Copy live data into other semispace
    3. Declare everything in current semispace dead
    4. Switch to other semispace

Stop and Copy Example

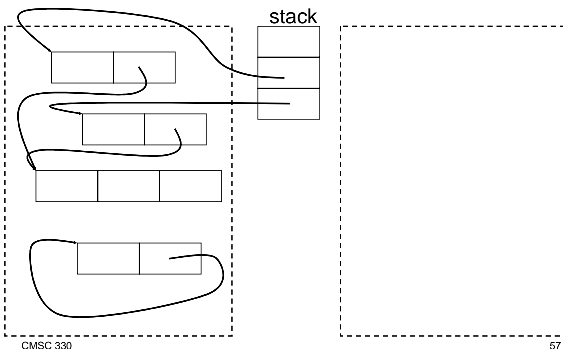

Stop and Copy Example (cont.)

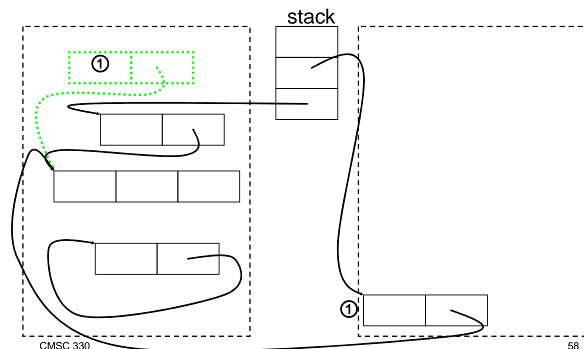

Stop and Copy Example (cont.)

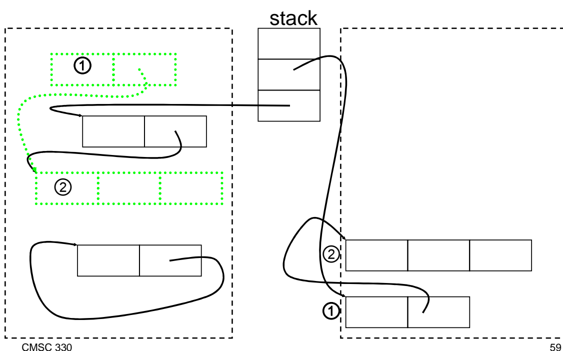

Stop and Copy Example (cont.)

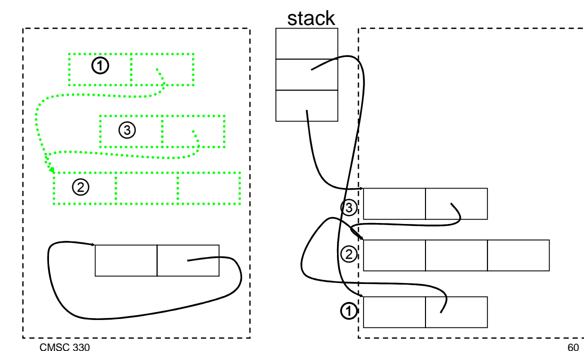

Stop and Copy Tradeoffs

- Advantages
  - Only touches live data
  - No fragmentation (automatically compacts)
    - Will probably increase locality
- Disadvantages
  - Requires twice the memory space
  - Like mark and sweep, need to "stop the world"
    - Program must stop running to let garbage collector move around data in the heap

The Generational Principle

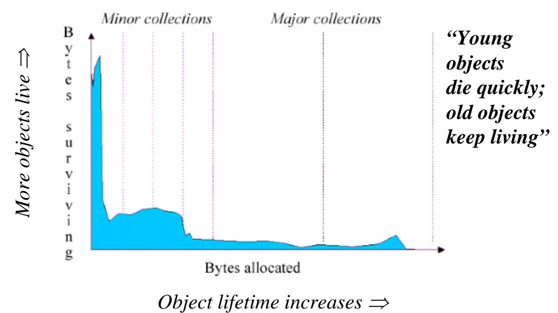

"Young objects die quickly; old objects keep living"

Generational Collection

- Long lived objects get copied over and over
  - Idea: Have more than one semispace, divide into generations
    - Older generations collected less often
    - Objects that survive many collections get pushed into older generations
    - Need to track pointers from old to young generations to use as roots for young generation collection
- One popular setup
  - Generational stop and copy

Java HotSpot SDK 1.4.2 Collector

- Multi-generational, hybrid collector
  - Young generation
    - Stop and copy collector
  - Tenured generation
    - Mark and sweep collector
  - Permanent generation
    - No collection
- Questions
  - Why does using a copy collector for the youngest generation make sense?
  - What apps will be penalized by this setup?

More Issues in GC (cont.)

- Stopping is world is a big hit
  - Unpredictable performance
    - Bad for real-time systems
  - Need to stop all threads
    - Without a much more sophisticated GC
- One-size fits all solution
  - Sometimes, GC just gets in the way
  - But correctness comes first

What Does GC Mean to You?

- Ideally, nothing
  - GC should make programming easier
  - GC should not affect performance (much)
- Usually bad idea to manage memory yourself
  - Using object pools, free lists, object recycling, etc…
  - GC implementations have been heavily tuned
    - May be more efficient than explicit deallocation
- If GC becomes a problem, hard to solve
  - You can set parameters of the GC
  - You can modify your program

Increasing Memory Performance

- Don't allocate as much memory
  - Less work for your application
  - Less work for the garbage collector
- Don't hold on to references
  - Null out pointers in data structures
  - Example

    Object a = new Object;
    …use a…
    a = null; // when a is no longer needed

Find the Memory Leak

```
class Stack {
  private Object[] stack;
  private int index;
  public Stack(int size) {
    stack = new Object[size];
  }
  public void push(Object o) {
    stack[index++] = o;
  }
  public void pop() {
    stack[index] = null;  // null out ptr
    return stack[index--];
  }
}
```

From Haggar, Garbage Collection and the Java Platform Memory Model

Answer: pop() leaves item on stack array; storage not reclaimed