CS577 W'04 Lecture Notes
Lecture 10

Basic Garbage Collection

Garbage Collection (GC) is the automatic reclamation of heap records that will never again be accessed by the program.

GC is universally used for languages with closures and complex data structures that are **implicitly** heap-allocated.

GC may be useful for any language that supports heap allocation, because it obviates the need for explicit deallocation, which is tedious, error-prone, and often non-modular.

GC technology is increasingly interesting for "conventional" language implementation, especially as users discover that `free` isn't free. I.e., explicit memory management can be costly too.

We view GC as part of an **allocation service** provided by the runtime environment to the user program, usually called the **mutator**. When the mutator needs heap space, it calls an allocation routine, which in turn performs garbage collection activities if needed.

Simple Heap Model

For simplicity, consider a heap containing "cons" cells.

Heap consists of **two-word cells** and each element of a cell is a **pointer** to another cell. (We'll deal with distinguishing pointers from non-pointers later.)

There may also be pointers into the heap from the stack and global variables; these constitute the **root set**.

At any given moment, the system's **live data** are the heap cells that can be reached by **some** series of pointer traversals starting from a member of the root set.

Garbage is the heap memory containing non-live cells. (Note that this is a slightly conservative definition.)

Reference Counting

The most straightforward way to recognize garbage and make its space reusable for new cells is to use **reference counts**.

We augment each heap cell with a **count field** that records the total number of pointers in the system that point to the cell. Each time we create or copy a pointer to the cell, we increment the count; each time we destroy a pointer, we decrement the count.

If the reference count ever goes to 0, we can reuse the cell by placing it on a **free list**.

When allocating a new cell, we first try the free list (before extending the heap).

Pros:
Conceptually simple;
Immediate reclamation of storage

Cons:
Extra space;
Extra time (every pointer assignment has to change/check count)
Can't collect "cyclic garbage"

Mark and Sweep

There's no real need to remove garbage as long as unused memory is available. So GC is typically deferred until the **allocator** fails due to lack of memory. The collector then takes control of the processor, performs a collection—hopefully freeing enough memory to satisfy the allocation request—and returns control to the mutator. This approach is known generically as "stop and collect."

There are several options for the collection algorithm. Perhaps the simplest is called **mark and sweep**, which operates in two phases:

• First, **mark** each live data cell by tracing all pointers starting with the root set.

• Then, **sweep** all unmarked cells onto the free list (also unmarking the marked cells).

```
struct cell {
   int mark:1;
   struct cell *c[2];}

struct cell *free;

struct cell heap[HEAPSIZE];

struct cell *roots[ROOTS];
```

```
/* Initially all cells are on free list.
   Use c[0] to link members of free list. */
void init_heap() {
  for (i=0; i < HEAPSIZE-1; i++)
    heap[i].c[0] = &(heap[i+1]);
  heap[HEAPSIZE-1].c[0] = 0;
  free = &(heap[0]);
}


struct cell *allocate() {
  struct cell *a;
  if (!free) { /* no more room => */
    gc();        /*    try gc */
    if (!free) /* still no more room */
      die();
  };
  a = free;
  free = free->c[0];
  return a;
}
```

```
void gc() {
   for (i = 0; i < ROOTS; i++)
      mark(roots[i]);
   sweep();
}

void mark(struct cell *cell) {
   if (!cell->mark) {
      cell->mark = 1;
      mark(cell->c[0]);
      mark(cell->c[1]);
    }
}

void sweep() {
   for (i = 0; i < HEAPSIZE; i++)
      if (heap[i].mark)
        heap[i].mark = 0;
      else {
        heap[i].c[0] = free;
        free = &(heap[i]);
      }
}
```

Here `mark` traverses the live data graph in depth-first order, and
potentially uses **lots** of stack! A standard trick called **pointer
reversal** can be used to avoid needing extra space during the
traversal.

Copying Collection

Mark and sweep has several problems:

● It does work proportional to the size of the entire heap.

● It leaves memory fragmented.

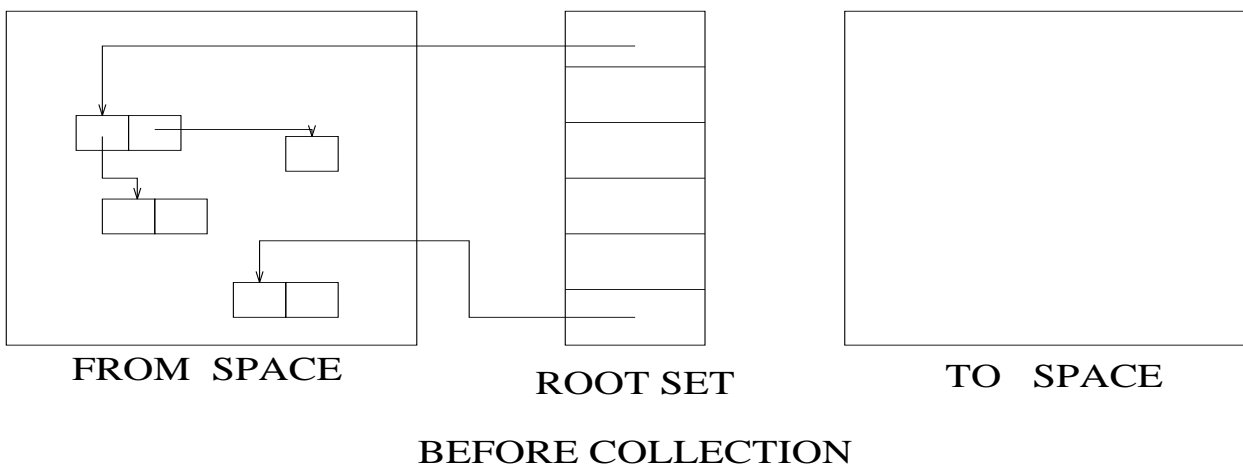● It doesn't cope well with non-uniform cell sizes.

An alternative that solves these problems is **copying collection**. The idea is to divide the available heap into 2 **semi-spaces**. Initially, the allocator uses just one space; when it fills up, the collector **copies** the live data (only) into the other space, and reverses the role of the spaces.
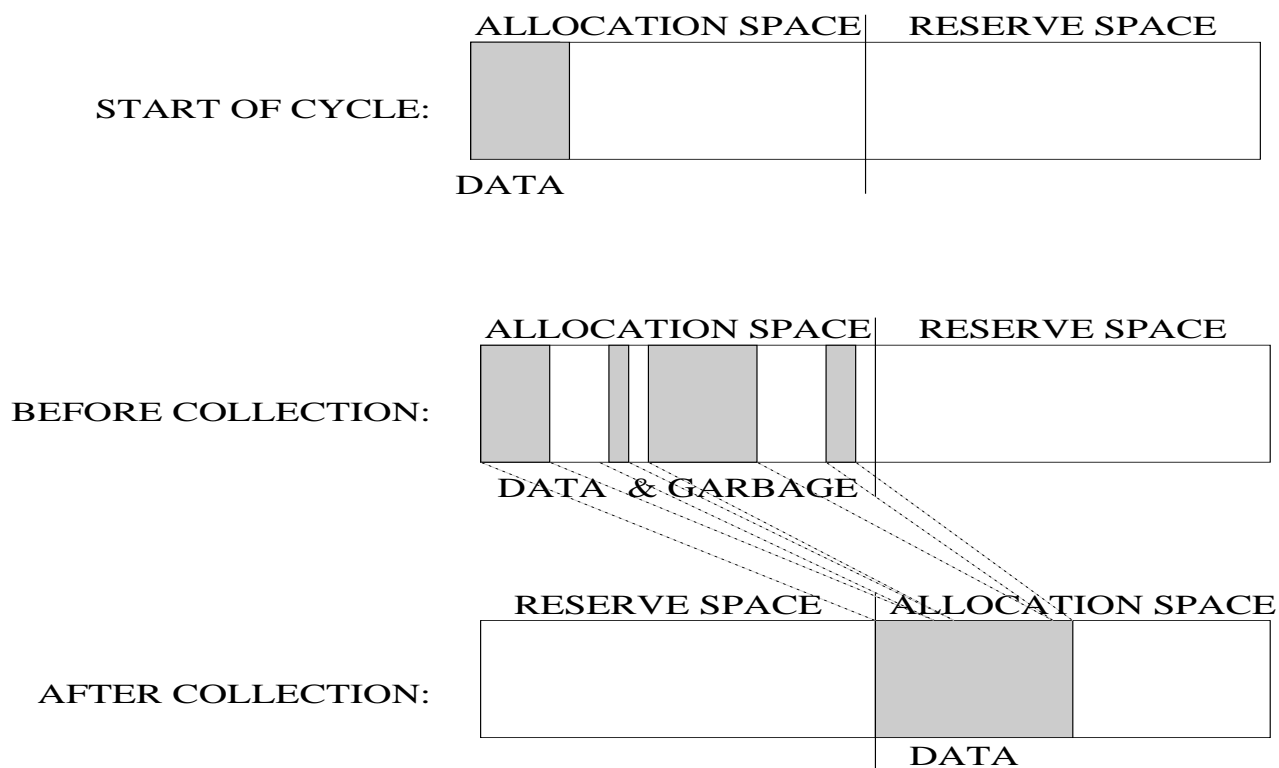
Copying collection must fix up **all** pointers to copied data. To do this, it leaves a **forwarding pointer** in the "from" space after the copy is made.

A copying collector typically traverses the live data graph **breadth first**, using "to" space itself as the search "queue."

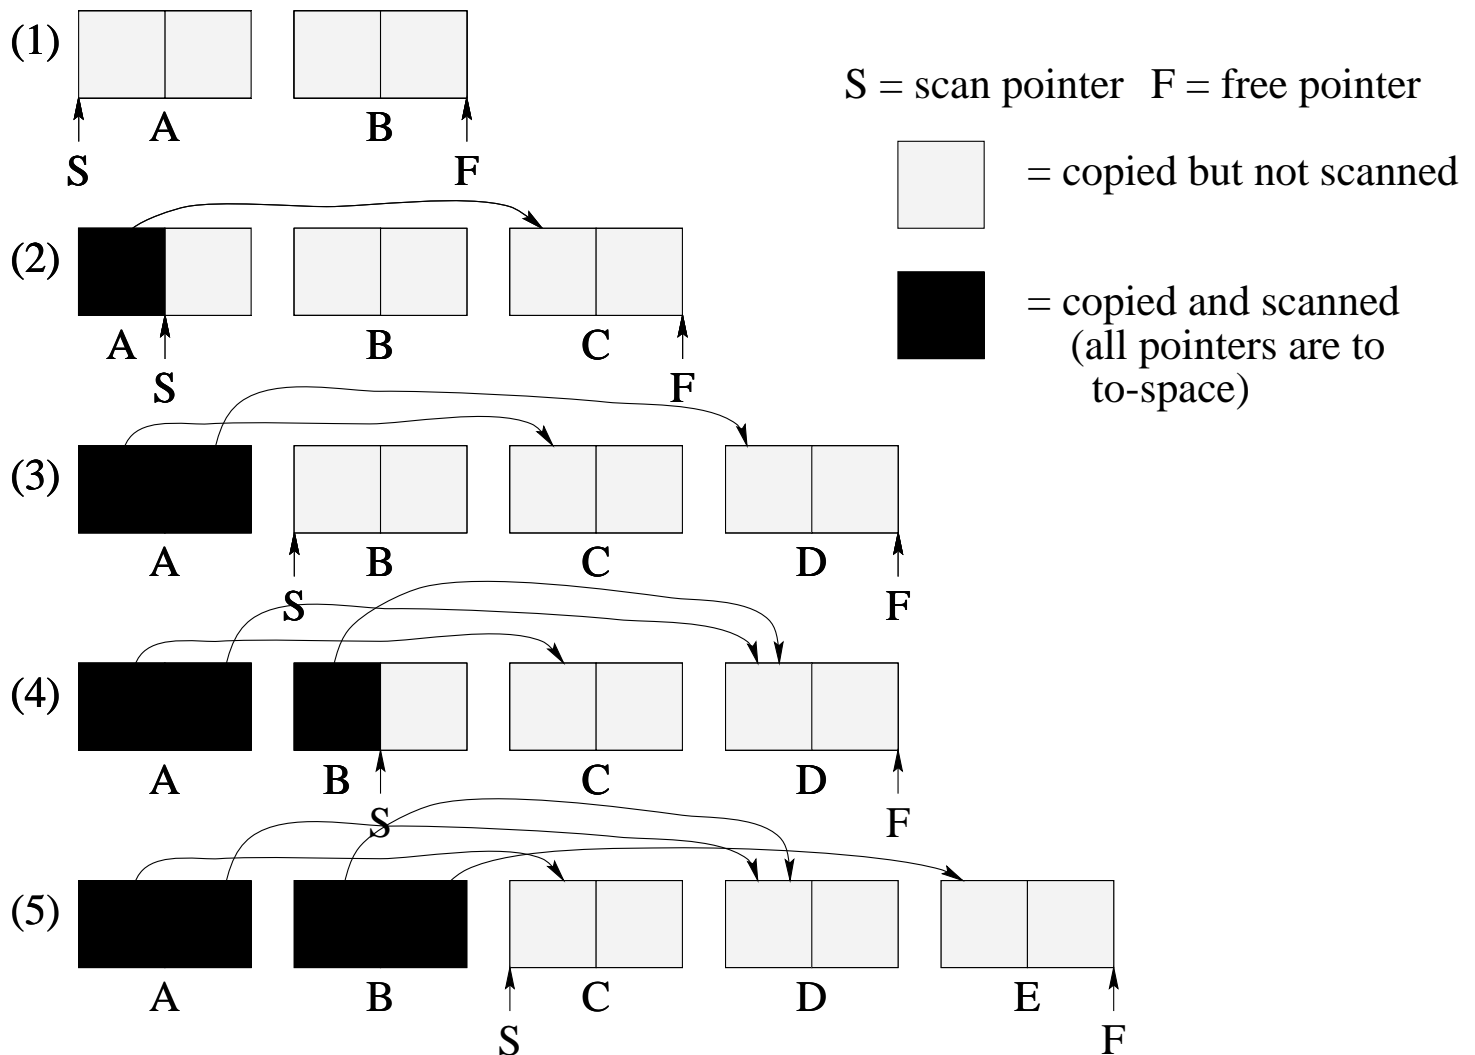Copying **compacts** live data, which improves locality and may be good for virtual memory and caches.

FROM SPACE   ROOT SET   TO SPACE

BEFORE COLLECTION

FROM SPACE   ROOT SET   TO SPACE

AFTER COLLECTION

Copying Collection Details

GRAPH

Root
Set

A

B

C

D

E

TO-SPACE

(1)

A   B

S   F

S = scan pointer   F = free pointer

= copied but not scanned

(2)

A   B   C

S   F

= copied and scanned
(all pointers are to
to-space)

(3)

A   B   C   D

S   F

(4)

A   B   C   D

S   F

(5)

A   B   C   D   E

S   F

```
struct cell {
  struct cell *c[2];
}

struct cell space[2][HALFSIZE];

struct cell *roots[ROOTS];

struct cell *free = &(space[0][0]);
struct cell *end = &(space[0][HALFSIZE]);

int from_space = 0;
int to_space = 1;

struct cell *allocate() {
  if (free == end) { /* no room */
    gc();
    if (free == end) /* still no room */
      die();
  };
  return free++;
}
```

```
gc() {
  int i;
  struct cell *scan = &(space[to_space][0]);
  free = scan;
  for (i = 0 ; i < ROOTS; i++)
    roots[i] = forward(roots[i]);
  while (scan < free) {
    scan->c[0] = forward(scan->c[0]);
    scan->c[1] = forward(scan->c[1]);
    scan++;
  };
  from_space = 1-from_space;
  to_space = 1-to_space;
  end = *(space[from_space][HALFSIZE]);
}

struct cell *forward(struct cell *p) {
  if (p >=&(space[from_space][0]) &&
        p < &(space[from_space][HALFSIZE]))
      if (p->c[0] >= &(space[to_space][0]) &&
           p->c[0] < &(space[to_space][HALFSIZE]))
         return p->c[0];
      else {
        *free = *p;
        p->c[0] = free++;
        return p->c[0];
      }
  else return p;
}
```

Comparison

Copying collector does work proportional to amount of **live** data. Asymptotically, this means it does less work than mark and sweep. Let

A = amount of live data
M = total memory size before a collection.

After the collection, there is M-A space left for allocation before the next collection. We can calculate the **amortized cost** per allocated byte as follows:

$$C_{M\&S} = \frac{c_1 A + c_2 M}{M - A} \qquad\qquad \text{for some } c_1, c_2$$

$$C_{COPY} = \frac{c_3 A}{\frac{M}{2} - A} \qquad\qquad \text{for some } c_3$$

As $M \to \infty$, $C_{COPY} \to 0$, while $C_{M\&S} \to c_2$.

Of course, real memories aren't infinite, so the values of c_1, c_2, c_3 matter, especially if a significant percentage of data are live at collection (since generally $c_3 > c_1$).