

15-213

"The course that gives CMU its Zip!"

Memory Management II: Dynamic Storage Allocation Mar 7, 2000

Topics

- Segregated free lists
 - Buddy system
 - Garbage collection
 - Mark and Sweep
 - Copying
 - Reference counting

class15.ppt

CS 213 S'00

- 2 -

Basic allocator mechanisms

Sequential fits (implicit or explicit single free list)

- best fit, first fit, or next fit placement
 - splitting thresholds
 - immediate or deferred coalescing

Segregated free lists

- simple segregated storage -- separate heap for each size class
 - segregated fits -- separate linked list for each size class
 - buddy systems

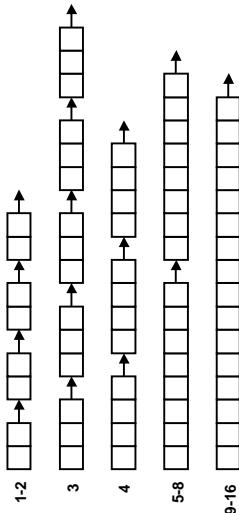
class15.ppt

CS 213 S'00

- 2 -

Segregate Storage

Each size "class" has its own collection of blocks



- Often have separate collection for every small size (2,3,4,...)
- For larger sizes typically have a collection for each power of 2

class15.ppt

CS 213 S'00

- 3 -

Simple segregated storage

Separate heap and free list for each size class

No splitting

To allocate a block of size n:

- if free list for size n is not empty,
 - allocate first block on list (note, list can be implicit or explicit)
- if free list is empty,
 - get a new page
 - create new free list from all blocks in page
 - allocate first block on list

constant time

To free a block:

- Add to free list
 - If page is empty, return the page for use by another size (optional)

Tradeoffs:

- fast, but can fragment badly
- class15.ppt
- CS 213 S'00

- 4 -

Segregated fits

Array of free lists, each one for some size class

To allocate a block of size n:

- search appropriate free list for block of size $m > n$
- if an appropriate block is found:
 - split block and place fragment on appropriate list (optional)
 - if no block is found, try next larger class
 - repeat until block is found

To free a block:

- coalesce and place on appropriate list (optional)

Tradeoffs

- faster search than sequential fits (i.e., log time for power of two size classes)
- controls fragmentation of simple segregated storage
- coalescing can increase search times
 - deferred coalescing can help

class15.ppt

CS 213 S'00

- 5 -

Buddy systems

Special case of segregated fits.

- all blocks are power of two sizes

Basic idea:

- Heap is 2^m words
 - Maintain separate free lists of each size 2^k , $0 \leq k \leq m$.
 - Requested block sizes are rounded up to nearest power of 2.
 - Originally, one free block of size 2^m .

- 6 -

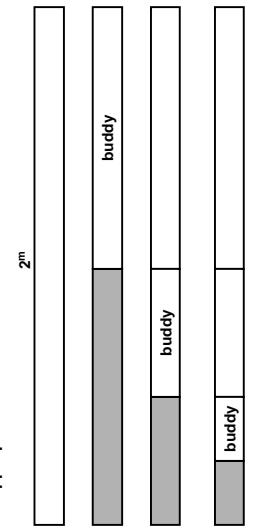
class15.ppt

CS 213 S'00

Buddy systems (cont)

To allocate a block of size 2^k :

- Find first available block of size 2^j s.t. $k \leq j \leq m$.
- if $j == k$ then done.
- otherwise recursively split block until $j == k$.
- Each remaining half is called a "buddy" and is placed on the appropriate free list



class15.ppt

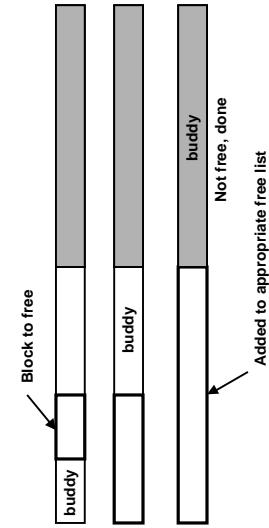
CS 213 S'00

- 7 -

Buddy systems (cont)

To free a block of size 2^k

- continue coalescing with buddies while the buddies are free



CS 213 S'00

- 8 -

Buddy systems (cont)

Key fact about buddy systems:

- given the address and size of a block, it is easy to compute the address of its buddy
 - e.g., block of size 32 with address `xxxx...x00000` has buddy `xxxx...x10000`

Tradeoffs:

- fast search and coalesce
- subject to internal fragmentation

class15.ppt

CS 213 S'00

- 9 -

Internal fragmentation

Internal fragmentation is wasted space inside allocated blocks:

- minimum block size larger than requested amount
 - e.g., due to minimum free block size, free list overhead
- policy decision not to split blocks
 - e.g., buddy system
 - Much easier to define and measure than external fragmentation.

class15.ppt

CS 213 S'00

- 10 -

Implicit Memory Management

Garbage collector

Garbage collection: automatic reclamation of heap-allocated storage -- application never has to free

```
void foo() {  
    int *p = malloc(128);  
    return /* p block is now garbage */  
}
```

Common in functional languages, scripting languages, and modern object oriented languages:

- Lisp, ML, Java, Perl, Mathematica, and C++

- Cannot collect all garbage

class15.ppt

CS 213 S'00

- 11 -

Garbage Collection

How does the memory manager know when memory can be freed?

- In general we cannot know what is going to be used in the future since it depends on conditionals
 - But we can tell that certain blocks cannot be used if there are no pointers to them

Need to make certain assumptions about pointers

- Memory manager can distinguish pointers from non-pointers
 - All pointers point to the start of a block
 - Cannot hide pointers (e.g. by coercing them to an int, and then back again)

class15.ppt

CS 213 S'00

- 12 -

Classical GC algorithms

Mark and sweep collection (McCarthy, 1960)

- Does not move blocks (unless you also “compact”)
- Reference counting (Collins, 1960)

Copying collection (Minsky, 1963)

- Does not move blocks
- Moves blocks

For more information see *Jones and Lin, “Garbage Collection: Algorithms for Automatic Dynamic Memory”*, John Wiley & Sons, 1996.

class15.ppt

CS 213 S'00

- 13 -

class15.ppt

CS 213 S'00

- 14 -

class15.ppt

CS 213 S'00

- 15 -

class15.ppt

CS 213 S'00

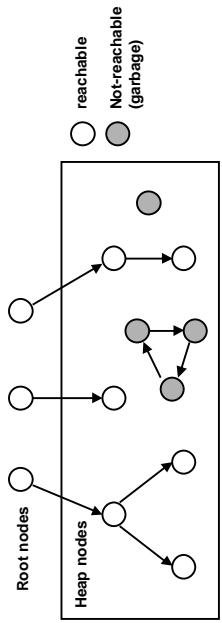
- 16 -

CS 213 S'00

Memory as a graph

We view memory as a directed graph

- Each block is a node in the graph
- Each pointer is an edge in the graph
- Locations not in the heap that contain pointers into the heap are called root nodes (e.g., registers, locations on the stack, global variables)



A node (block) is reachable if there is a path from any root to that node.
Non-reachable nodes are garbage (never needed by the application)

class15.ppt

CS 213 S'00

- 14 -

class15.ppt

CS 213 S'00

- 15 -

class15.ppt

CS 213 S'00

- 16 -

CS 213 S'00

Assumptions for this lecture

Application

- new(n): returns pointer to new block with all locations cleared
- read(b, i): read location i of block b into register v
- write(b, i, v): write v into location i of block b

Each block will have a header word

- addressed as $b[-1]$, for a block b
- Used for different purposes in different collectors

Instructions used by the Garbage Collector

- is_ptr(p): determines whether p is a pointer
- length(b): returns the length of block b, not including the header
- get_roots(): returns all the roots

class15.ppt

CS 213 S'00

- 15 -

class15.ppt

CS 213 S'00

- 16 -

CS 213 S'00

- 17 -

class15.ppt

CS 213 S'00

- 18 -

CS 213 S'00

Mark and sweep (cont.)

Mark using depth-first traversal of the memory graph

```
ptr mark(ptr p) {
    if (!is_ptr(p)) return; // do nothing if not pointer
    if (markBitSet(p)) return; // check if already marked
    setMarkBit(p); // set the mark bit
    for (i=0; i < length(p); i++) // mark all children
        mark(p[i]);
    return;
}
```

Sweep using lengths to find next block

```
ptr sweep(ptr p, ptr end) {
    while (p < end) {
        if (markBitSet(p))
            clearMarkBit();
        else if (allocateBitSet(p))
            free(p);
        p += length(p);
    }
}
```

class15.ppt

- 17 -

CS 213 S'00

Mark and sweep in C

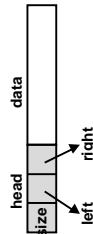
A C Conservative Collector

- `is_ptr()` can determine if a word is a pointer by checking if it points to an allocated block of memory.
- But, in C pointers can point to the middle of a block.



So how do we find the beginning of the block

- Can use balanced tree to keep track of all allocated blocks where the key is the location
- Balanced tree pointers can be stored in head (use two additional words)



class15.ppt

- 18 -

CS 213 S'00

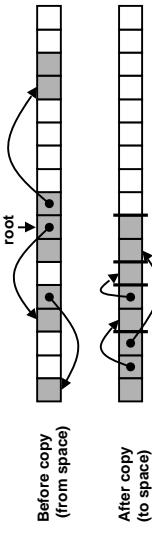
Copying collection (new)

Keep two equal-sized spaces, from-space and to-space

Repeat until application finishes

- Application allocates in one space contiguously until space is full.
- Stop application and copy all reachable blocks to contiguous locations in the other space.

- Flip the roles of the two spaces and restart application.



Copy does not necessarily keep the order of the blocks

Has the effect or removing all fragments

class15.ppt

- 19 -

CS 213 S'00

Copying collection (new)

```
ptr new (int n) {
    if (free+n > top) flip();
    newblock = free;
    free += (n+1);
    for (i=0; i < n+1; i++)
        newblock[i] = 0;
    return newblock+n;
}
```

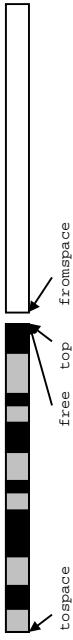
- All new blocks are allocated in tospace, one after the other
- An extra word is allocated for the header
- The Garbage-Collector starts (flips), when we reach top

class15.ppt

- 20 -

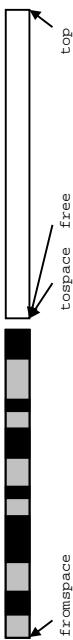
CS 213 S'00

Copying collection (flip)

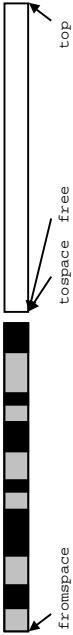


```
void flip() {
    swap(&fromspace, &tospace);
    top = tospace + size;
    free = tospace;
    for (r in roots)
        r = copy(r);
}
```

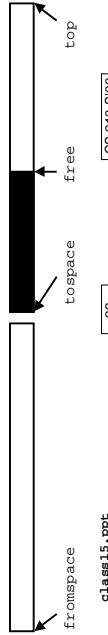
After the first three lines of flip (before the copy).



Copying collection (copy)



```
ptr copy(ptr p) {
    if (!is_ptr(p)) return p; // do nothing if not pointer
    if (p[-1] != 0) return p[-1]; // check if already forwarded
    new = free+1; // location for the copy
    p[-1] = new;
    new[-1] = 0;
    free += length(p)+1; // clear forward in new copy
    for (i=0; i < length(p); i++)
        new[i] = copy(p[i]); // increment free pointer
    for (i=0; i < length(p); i++)
        copy_all_children_and // copy all children and
        update_pointers_to_new_loc
    return new;
}
```



class15.ppt - 22 -

CS 213 S'00

Reference counting

Basic algorithm

- Keeps count on each block of how many pointers point to the block
- When a count goes to zero, the block can be freed

Data structures

- Can be built on top of an existing explicit allocator
- Add an additional header word for the "reference count"



- Keeping the count updated requires that we modify every read and write (can be optimized out in some cases)

class15.ppt - 23 -

CS 213 S'00

Reference counting

Set reference count to one when creating a new block

```
ptr new (int n) {
    newblock = allocate(n+1);
    newblock[0] = 1;
    return newblock+1;
}
```

When writing a value increment its reference counter

```
val read(ptr b, int i) {
    v = b[i];
    if (is_ptr(v)) v[-1]++;
    return v;
}
```

When writing decrement the old value and increment the new value

```
void write(ptr b, int i, val v) {
    decrement(b[i]);
    if (is_ptr(v)) v[-1]++;
    b[i] = v;
}
```

class15.ppt - 24 -

CS 213 S'00

Reference counting

Decrement

- If counter decrements to zero then the block can be freed
- When freeing a block, the algorithm must decrement the counters of everything pointed to by the block – this might in turn recursively free more blocks

```
void decrement(ptr p) {
    if (!is_ptr(p)) return;
    p[-1]--;
    if (p[-1] == 0) {
        for (i=0; i<length(p); i++)
            decrement(p[i]);
        free(p-1);
    }
}
```

class15.ppt

CS 213 S'00

CS 213 S'00

-26-

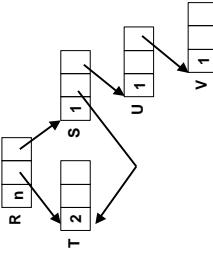
class15.ppt

CS 213 S'00

-26-

Reference counting example

Initially



Now consider write(R,1,NULL)

- This will execute a: decrement(s)

CS 213 S'00

-26-

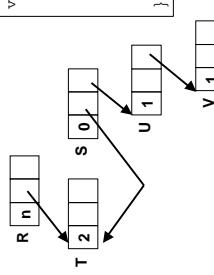
class15.ppt

CS 213 S'00

-26-

Reference counting example

After counter on S is decremented



```
void decrement(ptr p) {
    if (!is_ptr(p)) return;
    p[-1]--;
    if (p[-1] == 0) {
        for (i=0; i<length(p); i++)
            decrement(p[i]);
        free(p-1);
    }
}
```

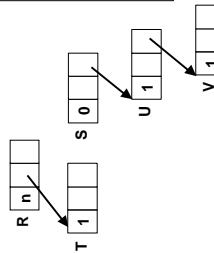
class15.ppt

CS 213 S'00

-27-

Reference counting example

After: decrement(s[0])



```
void decrement(ptr p) {
    if (!is_ptr(p)) return;
    p[-1]--;
    if (p[-1] == 0) {
        for (i=0; i<length(p); i++)
            decrement(p[i]);
        free(p-1);
    }
}
```

CS 213 S'00

-26-

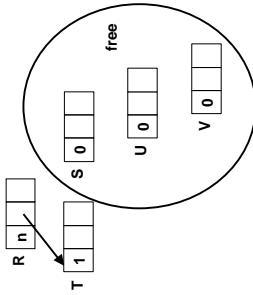
class15.ppt

CS 213 S'00

-26-

Reference counting example

After: decrement ($s[1]$)



class15.ppt

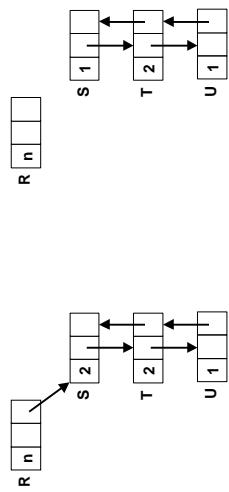
- 29 -

CS 213 S'00

Reference counting cyclic data structures

write ($R, 1, \text{NULL}$)

Before

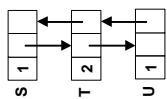


class15.ppt

- 30 -

CS 213 S'00

After



CS 213 S'00

Garbage Collection Summary

Copying Collection

- Pros: prevents fragmentation, and allocation is **very cheap**
- Cons: requires twice the space (from and to), and stops allocation to collect

Mark and Sweep

- Pros: requires little extra memory (assuming low fragmentation) and does not move data
- Cons: allocation is somewhat slower, and all memory needs to be scanned when sweeping

Reference Counting

- Pros: requires little extra memory (assuming low fragmentation) and does not move data
- Cons: reads and writes are more expensive and difficult to deal with cyclic data structures

Some collectors use a combination (e.g. copying for small objects and reference counting for large objects)

class15.ppt

- 31 -

CS 213 S'00