

Dynamic Memory Allocation: Advanced Concepts

Kai Shen

1

Keeping Track of Free Blocks

- Method 1: *Implicit list* using length—links all blocks



- Method 2: *Explicit list* among the free blocks using pointers

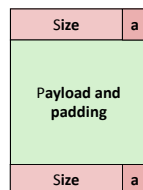


- Method 3: *Segregated free list*
 - Different free lists for different size classes

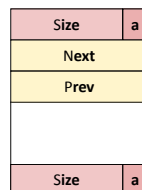
2

Explicit Free Lists

Allocated (as before)



Free

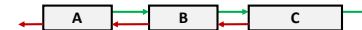


- Maintain list(s) of *free* blocks, not *all* blocks
 - The “next” free block could be anywhere
 - So we need to store forward/back pointers, not just sizes
 - Luckily we track only free blocks, the space usage is pretty much free
 - Still need boundary tags for coalescing

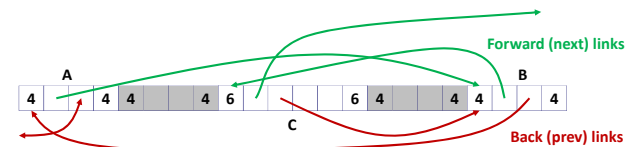
3

Explicit Free Lists

- Logically:



- Physically: blocks can be in any order



4

Allocating from Explicit Free Lists

- First fit
- Next fit
- Best fit

5

Allocating from Explicit Free Lists

conceptual graphic

Before

After (with splitting)

= malloc(...)

6

Freeing to Front

conceptual graphic

Before

Root

free()

- Insert the freed block at the root of the list

After

Root

- Simple and constant time

7

Freeing to Front (Coalescing)

conceptual graphic

Before

Root

free()

- Splice out predecessor block, coalesce both memory blocks, and insert the new block at the root of the list

After

Root

8

Explicit List Summary

- Comparison to implicit list:
 - Allocate is linear time in number of **free** blocks instead of **all** blocks
 - Much faster** when most of the memory is full
 - Slightly more complicated allocation and free since needs to splice blocks in and out of the list
 - Some extra space for the links (only in free blocks)
- But linear allocation time is still bad!

9

Keeping Track of Free Blocks

- Method 1: **Implicit list** using length—links all blocks



- Method 2: **Explicit list** among the free blocks using pointers

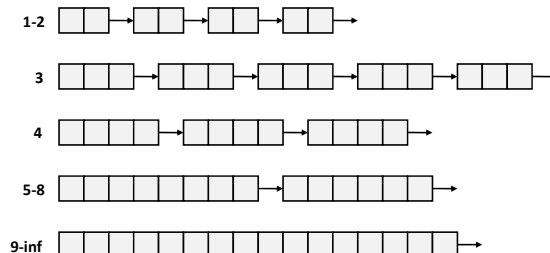


- Method 3: **Segregated free list**
 - Different free lists for different size classes

10

Segregated List Allocators

- Each **size class** of blocks has its own free list



11

Segregated List Allocator

- To allocate a block of size n :
 - Search appropriate free list for block of size $m \geq n$
 - If an appropriate block is found:
 - Split block and place fragment on appropriate list
 - If no block is found, try next larger class
 - Repeat until block is found
- To free a block:
 - Coalesce and place on appropriate list
- Advantages of segregated list allocators
 - High speed
 - Linear on the number of lists
 - log time for power-of-two size classes
 - Good memory utilization
 - Approximates a best-fit search of entire heap.

12

Implicit Memory Management: Garbage Collection

- **Garbage collection:** automatic reclamation of heap-allocated storage—application never has to free
- How does the memory manager know when memory can be freed?
 - If we can tell that programs will have no way to access certain heap objects

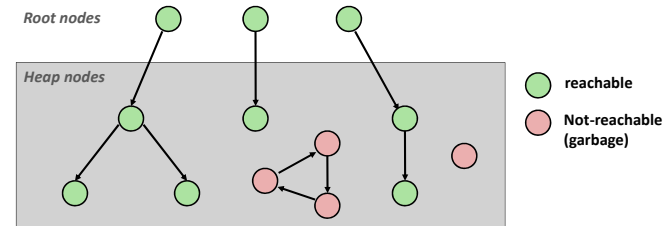
```
void foo() {
    String[] p = new String[2];
    return; /* p block is now garbage */
}
```

- All heap accesses must be made through known “references” (think of Java)

13

Memory as a Graph

- We view memory as a directed graph
 - Each data object is a node in the graph
 - Each reference is an edge in the graph
 - Registers, global variables, locations on the stack are always reachable



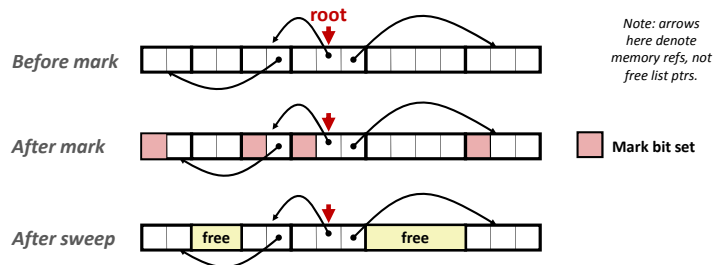
A node (object) is **reachable** if there is a path from any root to that node.

Non-reachable nodes are **garbage** (cannot be accessed by the program)

14

Mark and Sweep Collecting

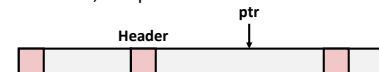
- Garbage Collection
 - Use extra **mark bit** in the head of each block
 - **Mark:** Start at roots and set mark bit on each reachable block
 - **Sweep:** Scan all blocks and free blocks that are not marked



15

Garbage Collection for C

- Problem for C garbage collection
 - References (pointers) are not strongly regulated
 - For instance, can point to middle of an allocated block



- C types can be cast back and forth
 - I cast **ptr** into an integer, subtract 1024 and save it
 - Later I add 1024 to the saved value and cast it back to a pointer, then dereference it ...

16

Memory-Related Perils and Pitfalls

- Reading uninitialized memory
- Overwriting memory
- Misunderstanding of pointer arithmetic
- Referencing nonexistent variables
- Referencing freed blocks
- Failing to free blocks

17

Reading Uninitialized Memory

- Assuming that heap data is initialized to zero

```
/* return y = Ax */
int *matvec(int **A, int *x) {
    int *y = malloc(N*sizeof(int));
    int i, j;

    for (i=0; i<N; i++)
        for (j=0; j<N; j++)
            y[i] += A[i][j]*x[j];

    return y;
}
```

18

Overwriting Memory

```
int **p;

p = malloc(N*sizeof(int *));

for (i=0; i<=N; i++) {
    p[i] = malloc(M*sizeof(int));
}
```

- Off-by-one error

19

Overwriting Memory

- Not checking the max string size

```
char s[8];

gets(s); /* reads "123456789" from stdin */
```

- Basis for classic buffer overflow attacks

20

Misunderstanding of Pointer Arithmetic

```
int *search(int *p, int val) {  
    while (*p && *p != val)  
        p += sizeof(int);  
  
    return p;  
}
```

21

Referencing Nonexistent Variables

- Forgetting that local variables disappear when a function returns

```
int *foo () {  
    int val;  
  
    return &val;  
}
```

22

Referencing Freed Blocks

```
x = malloc(N*sizeof(int));  
<manipulate x>  
free(x);  
  
... ..  
  
y = malloc(M*sizeof(int));  
for (i=0; i<M; i++)  
    y[i] = x[i]++;
```

23

Failing to Free Blocks (Memory Leaks)

- Slow, long-term killer!

```
foo() {  
    int *x = malloc(N*sizeof(int));  
    ...  
    return;  
}
```

24

Failing to Free Blocks (Memory Leaks)

- Freeing only part of a data structure

```
struct list {
    int val;
    struct list *next;
};

foo() {
    struct list *head = malloc(sizeof(struct list));
    head->val = 0;
    head->next = NULL;
    <create and manipulate the rest of the list>
    ...
    free(head);
    return;
}
```

25

Dealing With Memory Bugs

- Conventional debugger (**gdb**)
 - Not able to do much
- Debugging **malloc**
 - Wrapper around conventional **malloc**
 - Add canary values on allocation boundaries and do sanity checks sometimes
- Some malloc implementations contain checking code
 - Linux glibc malloc: **setenv MALLOC_CHECK_ 2**
- Binary translator: valgrind, Purify
 - Maintain memory map/status structure for each byte
 - Instrument data accesses to check for errors

26

Disclaimer

These slides were adapted from the CMU course slides provided along with the textbook of "Computer Systems: A programmer's Perspective" by Bryant and O'Hallaron. The slides are intended for the sole purpose of teaching the computer organization course at the University of Rochester.

27