

UNIVERSITY of WASHINGTON L25: Memory Allocation II CSE351, Autumn 2017

Memory Allocation II

CSE 351 Autumn 2017

Instructor:
Justin Hsia

Teaching Assistants:
Lucas Wotton Michael Zhang Parker DeWilde Ryan Wong
Sam Gehman Sam Wolfson Savanna Yee Vinny Palaniappan

UNIVERSITY of WASHINGTON L25: Memory Allocation II CSE351, Autumn 2017

Administrivia

- ❖ Homework 5 due Friday (12/1)
- ❖ Lab 5 due 12/8
- ❖ **Final Exam:** Wed, Dec. 13 @ 12:30pm in KNE 120
 - Same seating chart as Midterm
 - Review Session: Mon, Dec. 11 @ 5:00pm in EEB 105
 - Cumulative (midterm clobber policy applies)
 - You get TWO double-sided handwritten 8.5x11" cheat sheets
 - Recommended that you reuse or remake your midterm cheat sheet

2

UNIVERSITY of WASHINGTON L25: Memory Allocation II CSE351, Autumn 2017

Peer Instruction Question

- ❖ Which allocation strategy and requests remove *external* fragmentation in this Heap? B3 was the last fulfilled request.
 - <http://PollEv.com/justinh>

(A) Best-fit:
`malloc(50), malloc(50)`

(B) First-fit:
`malloc(50), malloc(30)`

(C) Next-fit:
`malloc(30), malloc(50)`

(D) Next-fit:
`malloc(50), malloc(30)`

3

UNIVERSITY of WASHINGTON L25: Memory Allocation II CSE351, Autumn 2017

Implicit List: Allocating in a Free Block

- ❖ Allocating in a free block: *splitting*
 - Since allocated space might be smaller than free space, we might want to split the block

Assume `ptr` points to a *free* block and has *unscaled* pointer arithmetic

```
void split(ptr b, int bytes) { // bytes = desired block size
    int newsize = ((bytes+7) >> 3) << 3; // round up to multiple of 8
    int oldsize = *b; // why not mask out low bit?
    *b = newsize; // initially unallocated
    if (newsize < oldsize) // set length in remaining
        *(b+newsize) = oldsize - newsize; // part of block (UNSCALED +)
```

```
malloc(12):
    ptr b = find(12+4)
    split(b, 12+4)
    allocate(b)
```

4

UNIVERSITY of WASHINGTON L25: Memory Allocation II CSE351, Autumn 2017

Implicit List: Freeing a Block

- ❖ Simplest implementation just clears "allocated" flag
 - `void free(ptr p) {*(p-BOX) &= -2;}`
 - But can lead to "false fragmentation"

`free(p)`

`malloc(20)`

Oops! There is enough free space, but the allocator won't be able to find it

5

UNIVERSITY of WASHINGTON L25: Memory Allocation II CSE351, Autumn 2017

Implicit List: Coalescing with Next

- ❖ Join (*coalesce*) with next block if also free

```
void free(ptr p) { // p points to data
    ptr b = p - BOX; // b points to block
    *b &= -2; // clear allocated bit
    ptr next = b + *b; // find next block (UNSCALED +)
    if ((*next & 1) == 0) // if next block is not allocated,
        *b += *next; // add its size to this block
}
```

- ❖ How do we coalesce with the *previous* block?

6

Implicit List: Bidirectional Coalescing

- ❖ **Boundary tags** [Knuth73]
 - Replicate header at "bottom" (end) of free blocks
 - Allows us to traverse backwards, but requires extra space
 - Important and general technique!

Format of allocated and free blocks:

Header	size	a
	payload and padding	
Boundary tag (footer)	size	a

a = 1: allocated block
a = 0: free block

size: block size (in bytes)

payload: application data (allocated blocks only)

7

Constant Time Coalescing

Block being freed →

8

Constant Time Coalescing

Case 1: Allocated blocks m1, m1, n, m2, m2. Free blocks n, m2. Result: Allocated blocks m1, m1, n, m2, m2. Free blocks n+m2.

Case 2: Allocated blocks m1, m1, n, m2, m2. Free blocks n, m2. Result: Allocated blocks m1, m1, n+m2, m2.

Case 3: Allocated blocks m1, m1, n, m2, m2. Free blocks n, m2. Result: Allocated blocks m1, n+m1, m2, m2. Free blocks n+m1.

Case 4: Allocated blocks m1, m1, n, m2, m2. Free blocks n, m2. Result: Allocated blocks m1, n+m1+m2, m2. Free blocks n+m1+m2.

10

Implicit Free List Review Questions

- ❖ What is the block header? What do we store and how?
- ❖ What are boundary tags and why do we need them?
- ❖ When we coalesce free blocks, how many neighboring blocks do we need to check on either side? Why is this?
- ❖ If I want to check the size of the *n*-th block forward from the current block, how many memory accesses do I make?

10

Keeping Track of Free Blocks

□ = 4-byte box (free)
■ = 4-byte box (allocated)

- Implicit free list** using length – links all blocks using math
 - No actual pointers, and must check each block if allocated or free
- Explicit free list** among only the free blocks, using pointers
- Segregated free list**
 - Different free lists for different size "classes"
- Blocks sorted by size**
 - Can use a balanced binary tree (e.g. red-black tree) with pointers within each free block, and the length used as a key

11

Explicit Free Lists

Allocated block:

size	a
payload and padding	
size	a

(same as implicit free list)

Free block:

size	a
next	
prev	
size	a

- ❖ Use list(s) of **free** blocks, rather than implicit list of **all** blocks
 - The "next" free block could be anywhere in the heap
 - So we need to store next/previous pointers, not just sizes
 - Since we only track free blocks, so we can use "payload" for pointers
 - Still need boundary tags (header/footer) for coalescing

12

Doubly-Linked Lists

- ❖ **Linear**
 - Needs head/root pointer
 - First node prev pointer is NULL
 - Last node next pointer is NULL
 - Good for first-fit, best-fit
- ❖ **Circular**
 - Still have pointer to tell you which node to start with
 - No NULL pointers (term condition is back at starting point)
 - Good for next-fit, best-fit

Explicit Free Lists

- ❖ **Logically:** doubly-linked list
- ❖ **Physically:** blocks can be in any order

Allocating From Explicit Free Lists

Note: These diagrams are not very specific about *where inside a block* a pointer points. In reality we would always point to one place (e.g. start/header of a block).

Before

After (with splitting)

= malloc(...)

Allocating From Explicit Free Lists

Note: These diagrams are not very specific about *where inside a block* a pointer points. In reality we would always point to one place (e.g. start/header of a block).

Before

After (fully allocated)

= malloc(...)

Freeing With Explicit Free Lists

- ❖ **Insertion policy:** Where in the free list do you put the newly freed block?
 - **LIFO (last-in-first-out) policy**
 - Insert freed block at the beginning (head) of the free list
 - Pro: simple and constant time
 - Con: studies suggest fragmentation is worse than the alternative
 - **Address-ordered policy**
 - Insert freed blocks so that free list blocks are always in address order:
 $address(previous) < address(current) < address(next)$
 - Con: requires linear-time search
 - Pro: studies suggest fragmentation is better than the alternative

Coalescing in Explicit Free Lists

	Case 1	Case 2	Case 3	Case 4
Block being freed	Allocated	Allocated	Free	Free
	Allocated	Free	Allocated	Free

- ❖ Neighboring free blocks are *already part of the free list*
 - 1) Remove old block from free list
 - 2) Create new, larger coalesced block
 - 3) Add new block to free list (insertion policy)
- ❖ How do we tell if a neighboring block is free?

UNIVERSITY of WASHINGTON L25: Memory Allocation II CSE351, Autumn 2017

Keeping Track of Free Blocks

□ = 4-byte box (free)
■ = 4-byte box (allocated)

- Implicit free list** using length – links all blocks using math
 - No actual pointers, and must check each block if allocated or free
- Explicit free list** among only the free blocks, using pointers
 - Diagram: A row of 10 boxes. Boxes at indices 20, 16, 24, and 8 are shaded. Solid arrows connect the length of each free block to the next free block. The free blocks are at indices 24-20, 20-16, 16-8, and 8-0.
- Segregated free list**
 - Different free lists for different size “classes”
- Blocks sorted by size**
 - Can use a balanced binary tree (e.g. red-black tree) with pointers within each free block, and the length used as a key

25

UNIVERSITY of WASHINGTON L25: Memory Allocation II CSE351, Autumn 2017

Segregated List (SegList) Allocators

- Each *size class* of blocks has its own free list
- Organized as an array of free lists

Size class (in bytes)

- Often have separate classes for each small size
- For larger sizes: One class for each two-power size

26

UNIVERSITY of WASHINGTON L25: Memory Allocation II CSE351, Autumn 2017

Allocation Policy Tradeoffs

- Data structure of blocks on lists
 - Implicit (free/allocated), explicit (free), segregated (many free lists) – others possible!
- Placement policy: first-fit, next-fit, best-fit
 - Throughput vs. amount of fragmentation
- When do we split free blocks?
 - How much internal fragmentation are we willing to tolerate?
- When do we coalesce free blocks?
 - Immediate coalescing:** Every time `free` is called
 - Deferred coalescing:** Defer coalescing until needed
 - e.g. when scanning free list for `malloc` or when external fragmentation reaches some threshold

30

UNIVERSITY of WASHINGTON L25: Memory Allocation II CSE351, Autumn 2017

More Info on Allocators

- D. Knuth, “*The Art of Computer Programming*”, 2nd edition, Addison Wesley, 1973
 - The classic reference on dynamic storage allocation
- Wilson et al, “*Dynamic Storage Allocation: A Survey and Critical Review*”, Proc. 1995 Int’l Workshop on Memory Management, Kinross, Scotland, Sept, 1995.
 - Comprehensive survey
 - Available from CS:APP student site (csapp.cs.cmu.edu)

31