

Memory Allocation III

CSE 351 Winter 2018

Instructor:

Mark Wyse

Teaching Assistants:

Kevin Bi


Parker DeWilde

Emily Furst

Sarah House

Waylon Huang

Vinny Palaniappan



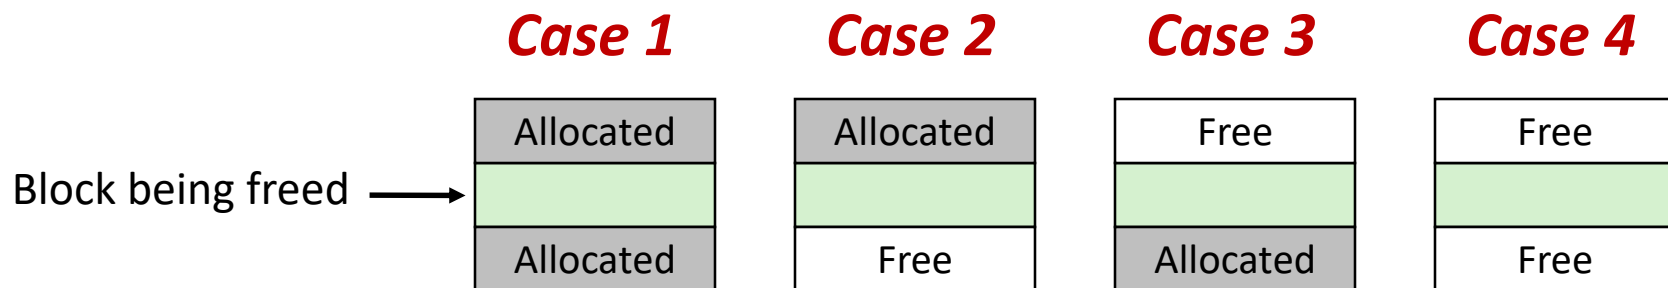
<https://xkcd.com/835/>

Administrivia

- ❖ Homework 5 due tonight
- ❖ Lab 5 due Saturday (3/10)
 - Recommended that you watch the Lab 5 helper videos
- ❖ Course Evals – due Sun, March 11

- ❖ **Final Exam:** Wed, March 14 @ 2:30pm in KNE 110
 - Review Packet posted
- ❖ **Final Review:** Mon, March 12 @ 4:30pm in SIG 134
 - Go over problems from review packet
 - 4:30 – 6:30 pm


Coalescing in Explicit Free Lists




- ❖ 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?
boundary tags

Freeing with LIFO Policy (Case 1)

Boundary tags not shown, but don't forget about them!




❖ Insert the freed block at the root of the list




Freeing with LIFO Policy (Case 2)

Boundary tags not shown, but don't forget about them!

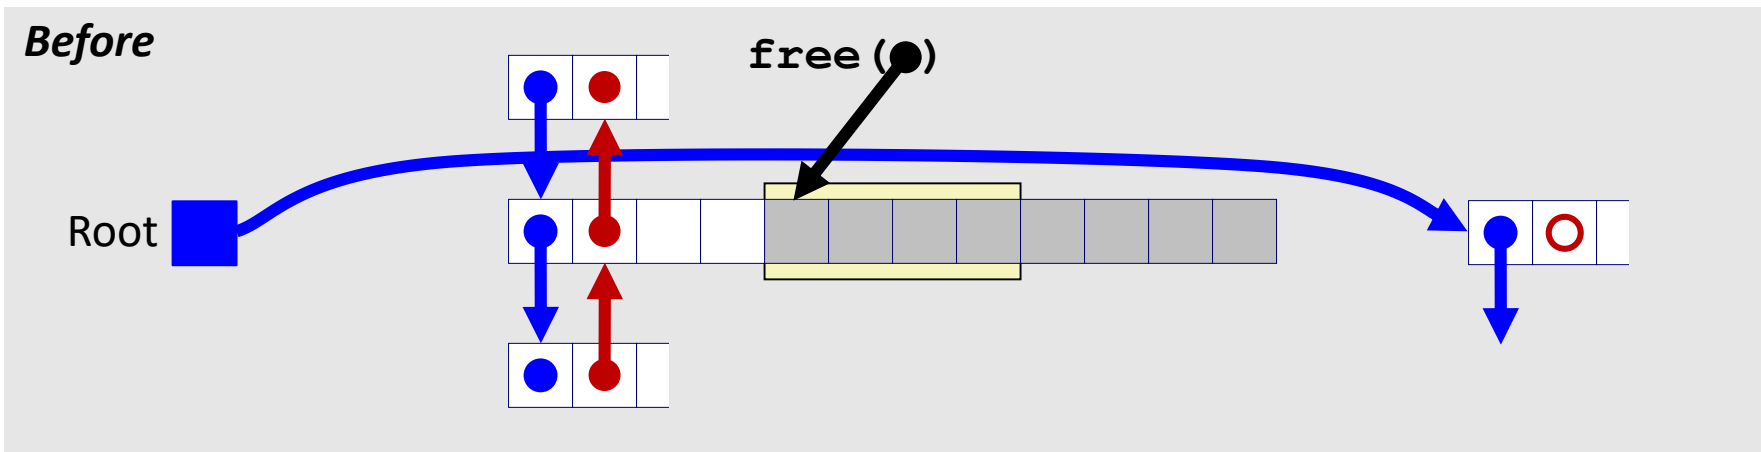


- ❖ Splice successor block out of list, coalesce both memory blocks, and insert the new block at the root of the list




Freeing with LIFO Policy (Case 3)

Boundary tags not shown, but don't forget about them!




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




Freeing with LIFO Policy (Case 4)

Boundary tags not shown, but don't forget about them!



- ❖ Splice predecessor and successor blocks out of list, coalesce all 3 memory blocks, and insert the new block at the root of the list

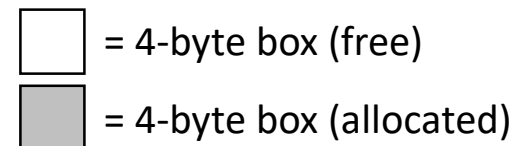


Explicit List Summary

- ❖ Comparison with implicit list:
 - Block allocation is linear time in number of *free* blocks instead of *all* blocks
 - *Much faster* when most of the memory is full
 - Slightly more complicated allocate and free since we need to splice blocks in and out of the list
 - Some extra space for the links (2 extra pointers needed for each free block)
 - Increases minimum block size, leading to more internal fragmentation


- ❖ Most common use of explicit lists is in conjunction with *segregated free lists*
 - Keep multiple linked lists of different size classes, or possibly for different types of objects

Keeping Track of Free Blocks




1) *Implicit free list* using length – links all blocks using math

- No actual pointers, and must check each block if allocated or free



2) *Explicit free list* among only the free blocks, using pointers



3) *Segregated free list*

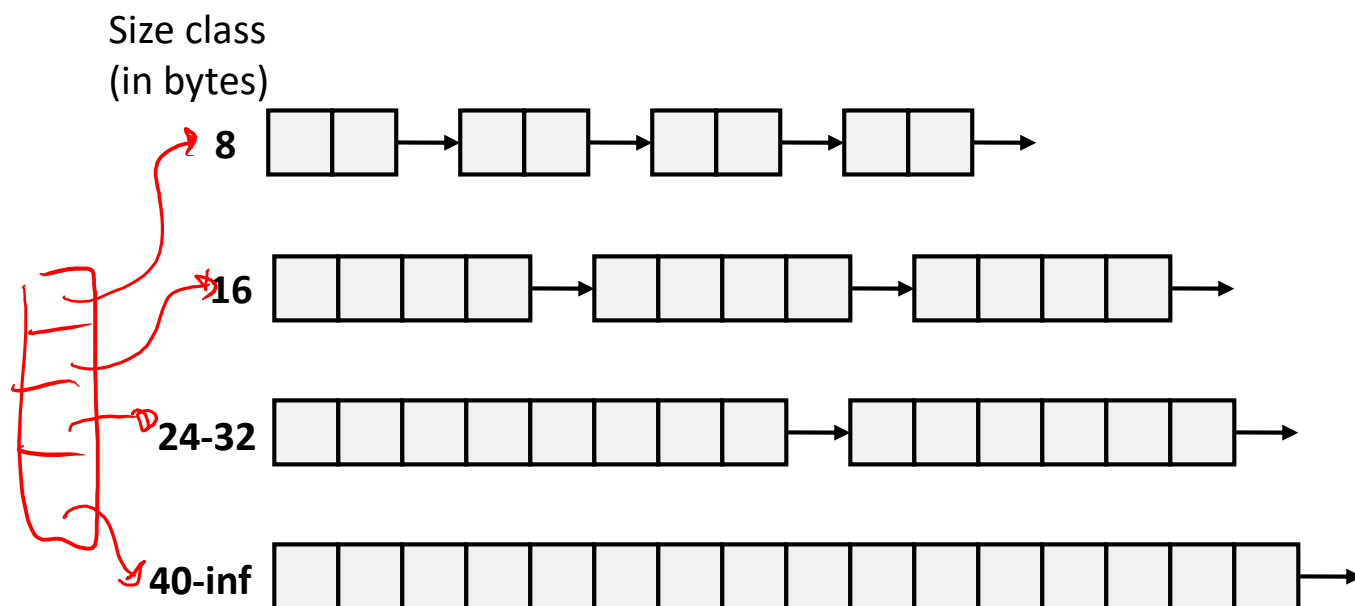
- Different free lists for different size “classes”

4) *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

Segregated List (SegList) Allocators

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



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

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

Memory Allocation

- ❖ Dynamic memory allocation
 - Introduction and goals
 - Allocation and deallocation (free)
 - Fragmentation
- ❖ Explicit allocation implementation
 - Implicit free lists
 - Explicit free lists (Lab 5)
 - Segregated free lists
- ❖ **Implicit deallocation: garbage collection**
- ❖ **Common memory-related bugs in C**

Wouldn't it be nice...

- ❖ If we never had to free memory?
- ❖ Do you free objects in Java?
 - Reminder: *implicit* allocator

Garbage Collection (GC)

(Automatic Memory Management)

- ❖ *Garbage collection*: automatic reclamation of heap-allocated storage – application never explicitly frees memory

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


- ❖ Common in implementations of functional languages, scripting languages, and modern object oriented languages:
 - Lisp, Racket, Erlang, ML, Haskell, Scala, Java, C#, Perl, Ruby, Python, Lua, JavaScript, Dart, Mathematica, MATLAB, many more...
- ❖ Variants (“conservative” garbage collectors) exist for C and C++
 - However, cannot necessarily collect all garbage

Garbage Collection

- ❖ How does the memory allocator 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 they are *unreachable* (via pointers in registers/stack/globals)
- ❖ Memory allocator needs to know what is a pointer and what is not – how can it do this?
 - Sometimes with help from the compiler

Memory as a Graph

- ❖ We view memory as a directed graph
 - Each allocated heap 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, stack locations, global variables)



A node (block) is **reachable** if there is a path from any root to that node
Non-reachable nodes are **garbage** (cannot be needed by the application)

Garbage Collection



- ❖ Dynamic memory allocator can free blocks if there are no pointers to them
- ❖ How can it know what is a pointer and what is not?
- ❖ We'll make some *assumptions* about pointers:
 - Memory allocator can distinguish pointers from non-pointers
 - All pointers point to the start of a block in the heap
 - Application cannot hide pointers (e.g. by coercing them to an `int`, and then back again)

Classical GC Algorithms

- ❖ Mark-and-sweep collection (McCarthy, 1960)
 - Does not move blocks (unless you also “compact”)
- ❖ Reference counting (Collins, 1960)
 - Does not move blocks (not discussed)
- ❖ Copying collection (Minsky, 1963)
 - Moves blocks (not discussed)
- ❖ Generational Collectors (Lieberman and Hewitt, 1983)
 - Most allocations become garbage very soon, so focus reclamation work on zones of memory recently allocated.
- ❖ For more information:
 - Jones, Hosking, and Moss, *The Garbage Collection Handbook: The Art of Automatic Memory Management*, CRC Press, 2012.
 - Jones and Lin, *Garbage Collection: Algorithms for Automatic Dynamic Memory*, John Wiley & Sons, 1996.

Mark and Sweep Collecting

- ❖ Can build on top of `malloc/free` package
 - Allocate using `malloc` until you “run out of space”
- ❖ When out of space:
 - Use extra **mark bit** in the header 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



Arrows are NOT free list pointers

Mark bit set

Memory-Related Perils and Pitfalls in C

		Slide	Prog stop Possible?	Security Flaw?
A)	Bad order of operations	32	Y	N
B)	Bad pointer arithmetic	31	Y	N
C)	Dereferencing a non-pointer	26	Y	Y
D)	Freed block – access again	35	Y	N
E)	Freed block – free again	34	Y	N
F)	Memory leak – failing to free memory	36	N	N
G)	No bounds checking	30	Y	Y
H)	Off-by-one error	29	Y	N
I)	Reading uninitialized memory	27	N	N
J)	Referencing nonexistent variable	33	N	Y
K)	Wrong allocation size	28	Y	N

Find That Bug! (Slide 26)

❖ The classic scanf bug

■ `int scanf(const char *format)`

```
int val;  
...  
scanf("%d", val);
```

→ stores int into address val

Error
Type:

C

Prog stop
Possible?

Y

Security flaw
Possible?

Y

Fix:

&val

if val
contains stack
addr

Find That Bug! (Slide 27)

```
/* return y = Ax */
int *matvec(int **A, int *x) {
    int *y = (int *)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;
}
```

- A is NxN matrix, x is N-sized vector (so product is vector of size N)
- N defined elsewhere (#define)

*Reading
Uninit. Mem*

Error
Type: I

Prog stop
Possible? N

Security flaw
Possible? N

Fix: *calloc()*

Find That Bug! (Slide 28)

```

int **p;

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

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

```

- N and M defined elsewhere (#define)

wrong
allocation
size

Error
Type: K

run off end
of alloc. block

Prog stop
Possible? Y

heap memory

Security flaw
Possible? N

Fix:

$N * \text{sizeof}(int^*)$

Find That Bug! (Slide 29)

```
int **p;

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

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

N+1 accesses

off by one error

Error
Type:

H

Prog stop
Possible?

Y

Security flaw
Possible?

N

Fix:

i < N

Find That Bug! (Slide 30)

```
char s[8];  
int i;  
  
gets(s); /* reads "123456789" from stdin */
```

*no bounds
checking*


Error
Type:



Prog stop
Possible?



Security flaw
Possible?



Fix:

fgets(s)

Find That Bug! (Slide 31)

```

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

```

no bounds checking

p += sizeof(int);
p += 4 → stride 16B

bad pointer arithmetic

if val not found,
will run off end
of array

only reading

Error
Type:

B

Prog stop
Possible?

Y

Security flaw
Possible?

N

Fix:

p++
add end condition

Find That Bug! (Slide 32)

```

int* getPacket(int** packets, int* size) {
    int* packet;
    packet = packets[0];
    packets[0] = packets[*size - 1];
    *size--; // what is happening here?
    reorderPackets(packets, *size);
    return packet;
}

```

*Trying to
decrement
size*

❖ ' -- ' happens first

*order of
operations*

*if you don't have
access to memory
just below size*

*just
reading*

Error
Type:

A

Prog stop
Possible?

Y

Security flaw
Possible?

N


Fix:

*(*size) --*

Find That Bug! (Slide 33)

```
int* foo() {
    int val;

    return &val;
}
```



ref. nonexistent variables

valid address on stack

if stack frame gets allocated w/ ret. addr at &val

pass by reference to foo or use malloc instead

Error Type: J

Prog stop Possible? N

Security flaw Possible? Y

Fix:

Find That Bug! (Slide 34)

```
x = (int*)malloc( N * sizeof(int) );  
    <manipulate x>  
free(x);  
  
...  
  
y = (int*)malloc( M * sizeof(int) );  
    <manipulate y>  
free(x); → y?
```

free again

*process
exits*

Error
Type:

E

Prog stop
Possible?

Y

Security flaw
Possible?

N

Fix:

free(y)
→
prob. a typo

Find That Bug! (Slide 35)

```

x = (int*)malloc( N * sizeof(int) );
    <manipulate x>
free(x);

    ...

y = (int*)malloc( M * sizeof(int) );
for (i=0; i<M; i++)
    y[i] = x[i]++;
  
```

→ already freed!

*access freed
memory*

Error
Type:

D

*process
exits*

Prog stop
Possible?

Y

Security flaw
Possible?

N

Fix:

*free(x) later
(at bottom)*

Find That Bug! (Slide 36)

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

void foo() {
    list *head = (list *) malloc( sizeof(list) );
    head->val = 0;
    head->next = NULL;
    <create and manipulate the rest of the list>
    ...
    free(head); ← only frees first node!
    return;
}
```

memory
leak

Error
Type: F

Prog stop
Possible? N

Security flaw
Possible? N

Fix: recursive/iterative
free over list

Dealing With Memory Bugs

- ❖ Conventional debugger (`gdb`)
 - Good for finding bad pointer dereferences
 - Hard to detect the other memory bugs
- ❖ Debugging `malloc` (UToronto CSRI `malloc`)
 - Wrapper around conventional `malloc`
 - Detects memory bugs at `malloc` and `free` boundaries
 - Memory overwrites that corrupt heap structures
 - Some instances of freeing blocks multiple times
 - Memory leaks
 - Cannot detect all memory bugs
 - Overwrites into the middle of allocated blocks
 - Freeing block twice that has been reallocated in the interim
 - Referencing freed blocks

Dealing With Memory Bugs (cont.)

- ❖ Some `malloc` implementations contain checking code
 - Linux glibc malloc: `setenv MALLOC_CHECK_ 2`
 - FreeBSD: `setenv MALLOC_OPTIONS AJR`
- ❖ Binary translator: **valgrind** (Linux), Purify
 - Powerful debugging and analysis technique
 - Rewrites text section of executable object file
 - Can detect all errors as debugging `malloc`
 - Can also check each individual reference at runtime
 - Bad pointers
 - Overwriting
 - Referencing outside of allocated block

What about Java or ML or Python or ...?

- ❖ In *memory-safe languages*, most of these bugs are impossible
 - Cannot perform arbitrary pointer manipulation
 - Cannot get around the type system
 - Array bounds checking, null pointer checking
 - Automatic memory management
- ❖ But one of the bugs we saw earlier is possible. Which one?

Memory Leaks with GC

- ❖ Not because of forgotten `free` — we have GC!
- ❖ Unneeded “leftover” roots keep objects reachable
- ❖ *Sometimes* nullifying a variable is not needed for correctness but is for performance
- ❖ Example: Don't leave big data structures you're done with in a static field

