

15-213 Recitation

Malloc Lab (Part II)

Your TAs

Friday, October 25th

Reminders

- **malloc** Deadlines:
 - *Checkpoint: **October 29th (Tuesday)***
 - *Final: **November 5th***
 - 7% of final grade (+4% for Checkpoint)
- Watch your email for Checkpoint Code Review sign-ups!
- *Bootcamp 5: Post Checkpoint Malloc* will be in-person!
 - **12PM-3PM, October 27th (Sunday)**
 - **NSH 3305**

Agenda

- **Review:**
 - **Heap Layout**
 - **Guide to Malloc Checkpoint**
- **Debugging**
 - **Finding errors with contracts and gdb**
 - **Instrumentation**
- **Style**
- **If time: Malloc Final Overview**


Review: Malloc Checkpoint

Suggested Roadmap: Checkpoint

- First: read the write-up!
 - “Roadmap to Success” section
- 0. Start writing your heap checker!
- 1. Implement **coalesce_block()** *first*.
- 2. Implement an *explicit free list*.
- 3. Implement *segregated lists*!

Starter Code

- Working implementation of an *implicit list* with boundary tags.
- No coalescing.



Implicit List


- Let's recap how we lay out and structure blocks in the heap!

Implicit List with Boundary Tags

Header and footer store:

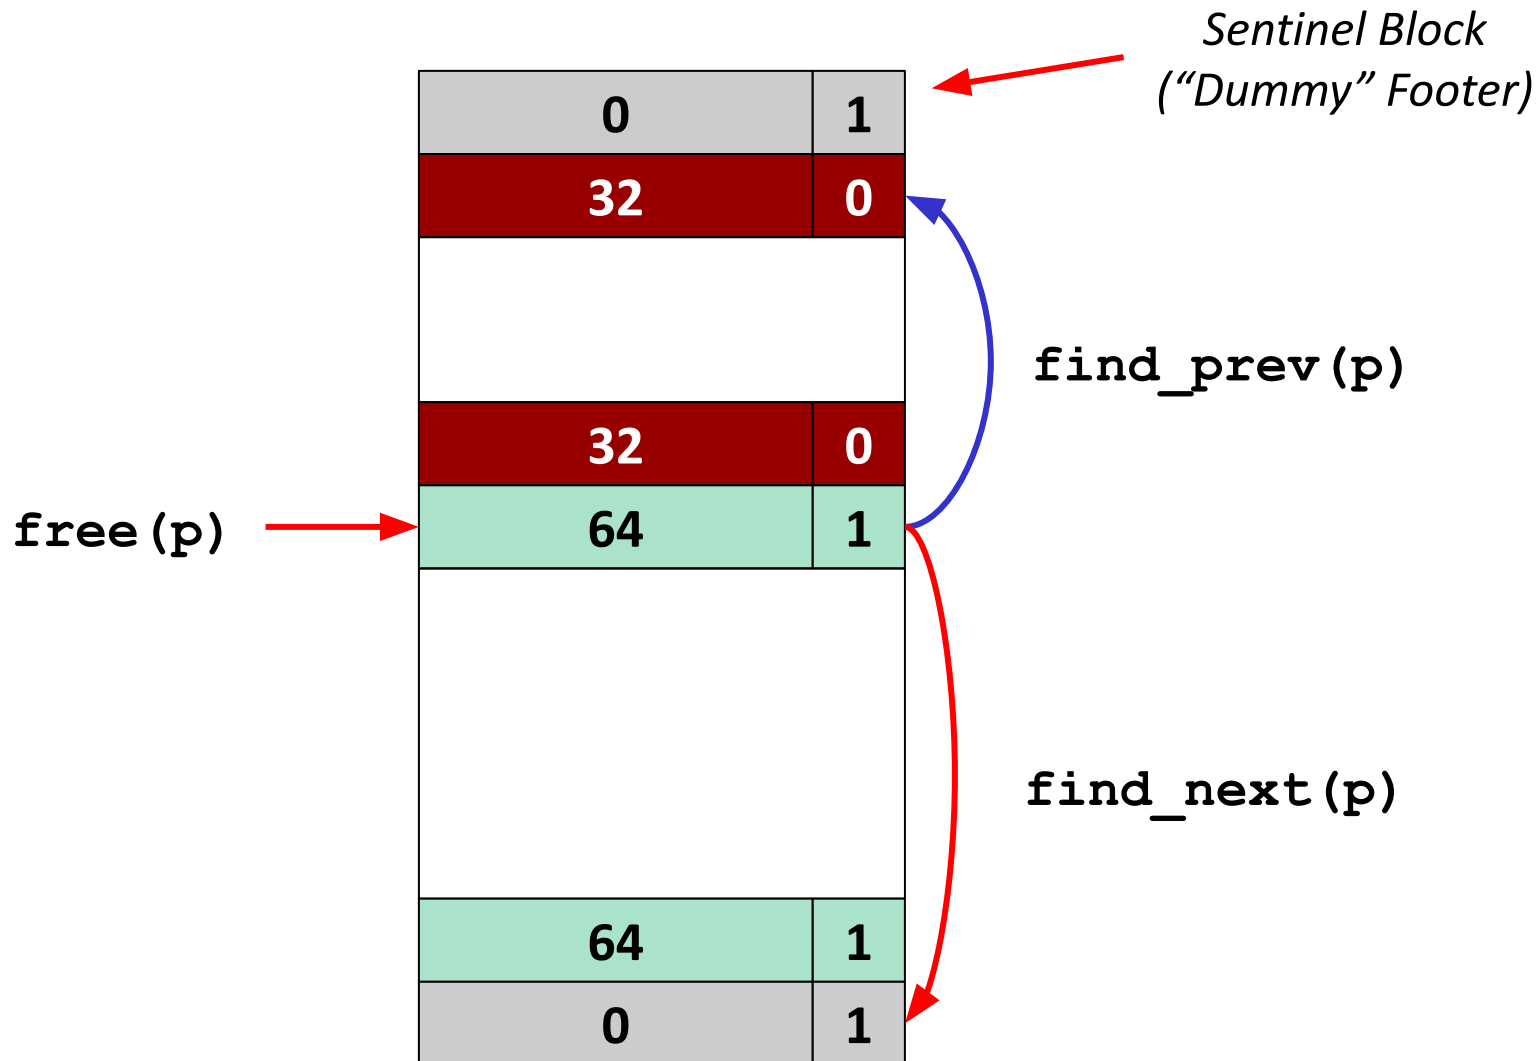
- Block size (including overhead)
- Allocation bit flag

```
typedef struct block {
    word_t header;
    char payload[0];
} block_t;
```




Payloads are 16-byte
aligned

Step 1: Coalescing



Step 2: Explicit Lists

- Idea: only search *free* blocks rather than the entire heap.
- Explicit list: free blocks explicitly point to other blocks, like in a linked list.
- You have 128 bytes of writable globals to work with:
 - Can maintain free list “head”. Remember to initialize in `mm_init()`!




`free_list_start`


(not necessarily first free block in heap)

Explicit Lists: Implementation

- Use free payload space to store pointers.
 - Unions!
- Write *helper functions* for inserting/removing.
- When do you need to insert? When do you need to remove?
- Update `find_fit` to scan your new list!




Allocated (as before)



Free

Step 3: Seglists!



- Now store *multiple* lists, one for each size class.
- Can maintain a global *array* of free list heads
 - Make sure to initialize these in `mm_init()`!
- Update insert/remove functions to insert to the correct bucket.
- Update `find_fit()` to scan the appropriate buckets.

Debugging

What does “Garbled Bytes” mean?

1. Your `malloc` returns a block pointer to satisfy a request.
2. `mdriver` writes bytes into payload
3. Later, `mdriver` checks that those bytes are intact:
 - If bytes have been overwritten, your `malloc` is overwriting data in an allocated block!

Now what?

- Double check your heap invariants. Are they exhaustive?
- If that doesn't help, use `gdb` to watch writes to the address getting garbled.

Debugging: Overview

- Refer to last week's recitation for common errors, and what they mean.
- *Use tools: **gdb** breakpoints and watchpoints.*
- *Write a heap checker! We'll be grading it in the next code review!*
 - Add new heap invariants as you add new features.
- ***Today: debugging activity!***
 - Garbled bytes
 - Using contracts

Debugging Activity

- Log into a Shark machine, then type:

```
$ wget http://www.cs.cmu.edu/~213/activities/rec9.tar
$ tar -xvf rec9.tar
$ cd rec9
```

- `mm.c` is a fake implicit list implementation, based on the starter code.
- It is buggy. Let's try and find the bugs!

Debugging Activity

- What happens if we run the program normally?

```
$ ./mdriver -c ./traces/syn-struct-short.rep

ERROR [trace ./traces/syn-struct-short.rep, line 16]: block 1 (at
0x8000000a0) has 8 garbled bytes, starting at byte 16
ERROR [trace ./traces/syn-struct-short.rep, line 21]: block 4 (at
0x800000180) has 8 garbled bytes, starting at byte 16

correctness check finished, by running tracefile
"traces/syn-struct-short.rep".
=> incorrect.

Terminated with 2 errors
```

Not very helpful...

Debugging Activity: Using Watchpoints

- Now let's try again with watchpoints!

```
$ gdb --args ./mdriver-dbg1 -c ./traces/syn-struct-short.rep

(gdb) watch *0x8000000a0
(gdb) run

// Keep continuing through the breaks:
// write_block()
// 4 x memcpy
Hardware watchpoint 1: *0x8000000a0

Old value = 129
New value = 32
write_block() at mm.c:333
```

- Now we know to take a closer look at `write_block()`!

Debugging Activity: Using Contracts

- Now let's run a version of the file that uses *contracts*:

```
$ ./mdriver-dbg2 -c ./traces/syn-struct-short.rep  
  
mdriver-dbg: mm.c:331: void write_block(block_t *, size_t, _Bool):  
Assertion `(unsigned long)footerp < ((long)block + size)' failed.  
Aborted (core dumped)
```

- This version had a contract in place to check that the footer is where we expect it to be.
- Writing effective contracts can save a lot of debugging time!

Debugging: Miscellaneous Tips

■ `mdriver`

- Use `-D` option to detect garbled bytes as soon as possible
- Use `-V` for verbose mode to find out which trace caused the error

- If the error happens in the first few allocations, can set breakpoints on `mm_malloc` and `mm_free` and step through line by line.

Instrumentation

Common Problems

- *Throughput is very low*
 - Which operation is likely the most costly? Where is the program likely to spend most of its time?
- *Utilization is very low / Out of Memory*
 - Which operation can cause you to allocate more memory than you may need?
- We can use *instrumentation* to investigate both problems!

Adding Instrumentation

- Instrumentation: add *temporary* code that collects measurements for metrics you're interested in.
 - e.g. how often are certain functions called?
 - You can always remove the code afterwards.
 - Can temporarily go over 128 byte writable global limit!
- These measurements can guide your development process:
 - Develop insights into performance before you spend time on implementation.

Instrumentation Example: Low Throughput

- Program is likely to spend most of its time in `find_fit()`'s loops.
- How efficient is your fit algorithm? How might you find out?

```
static block_t *find_fit(size_t asize)
{
    block_t *block; call_count++
    for (block = heap_listp; get_size(block) > 0;
         block = find_next(block))
    { block_count++
        if (!(get_alloc(block)) && (asize <= get_size(block)))
        {
            return block;
        }
    }
    return NULL; // no fit found
}
```

Instrumentation: Other Metrics

- What are the most common request sizes?
 - How many are 8 bytes or less?
 - How many are 16 bytes or less?
 - How might this inform your design?
- What other things might we want to measure?

Style

Style

- Checkpoint Code Review: Heap Checker Quality
- Final Code Review: Code Style
- Remember the style guidelines!
 - Modularity: use helper functions (e.g., for linked lists)!
 - Documentation
 - *File header*: have you described all your design decisions (block structure, fit algorithm, etc.)?

Wrapping Up

- `malloc` Deadlines:
 - *Checkpoint: **October 29th (Tuesday)***
 - *Final: **November 5th***
- Come to `malloc-final` bootcamp!
 - **12PM-3PM, October 27th (Sunday)**
 - **NSH 3305**


Further Content: Malloc Final

What are we trying to do?

- In Checkpoint, you dramatically improved the *throughput* of your allocator.
- For Final, you will need to greatly improve *utilization* while maintaining a high throughput.
- We will cover:
 1. Footer Removal in Allocated Blocks
 2. Decreasing Minimum Block Size
- For further content, see **malloc** bootcamp!


Footer Removal: Example

- Let's say we call `malloc(24)`. Can our block of size 32 satisfy the request?
- Add on overhead:
 - Header: +8 bytes = 32 bytes
 - Footer: +8 bytes = 40 bytes
- Round to multiple of 16 => 48 bytes
- Doesn't fit!



Footer Removal: Example

- What if we had no footer?
- Add on overhead:
 - Header: +8 bytes = 32 bytes
 - ~~○ Footer: +8 bytes = 40 bytes~~
- Round to multiple of 16 => 32 bytes
- Now it fits!
 - We have reduced *internal fragmentation*.



Footer Removal: Implementation

- What do we need footers for?
 - Coalescing
 - **Key observation:** do we need to know the size or position of the previous block if we're not going to coalesce with it?
- We just need some way to determine whether the block before us is allocated.

0	1
32	1
32	1
64	0
64	0
0	1

Footer Removal: Implementation

- We want to store one piece of information: is the previous block allocated?
- Where can we store a little *bit* of information for each block?

0	1
32	1
64	0
64	0
0	1

No footers in allocated blocks!

Free blocks still need footers.

Decreasing Minimum Block Size

- Currently, minimum block size is 32:
 - 8 byte header
 - 16 byte payload (min.)
 - 8 byte footer
- If we do `malloc(5)`, there's a lot of wasted space...
- Can we cut some of the fields for smaller blocks?

32	0
next	
prev	
32	0

32	1
5 bytes	
11 bytes wasted	
32	1

The End