

CSE251: System Programming

18. Dynamic Memory Allocation (2)

Seongil Wi

Recap: Fragmentation




`p1 = malloc(32)`



`p2 = malloc(40)`



`p3 = malloc(48)`



`free(p2)`



`p4 = malloc(64)`




Recap: Fragmentation




p1 = malloc(32)



p2 = malloc(40)



p3 = malloc(48)



free(p2)

Oops! There is enough aggregate heap memory, but no single free block is large enough



p4 = malloc(64)



Recap: Fragmentation




p1 = malloc(32)



p2 = malloc(40)



p3 = malloc(48)




free(p2)

Fragmentation (단편화, 파편화)



p4 = malloc(64)




Motivating Question



p1 = malloc(32)



p2 = malloc(40)



p3 = malloc(48)




free(p2)

- How to avoid fragmentation?
- How to maximize throughput and peak memory utilization?



p4 = malloc(64)



Today's Topic: Managing Free Blocks

1 square = 1 word = 8 bytes
□: Free word
■: Allocated word


- We need to efficiently track free blocks and select one for allocation



1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Today's Topic: Managing Free Blocks

- We need to efficiently track free blocks and select one for allocation




Keep track of free blocks as the list (e.g., linked list)

1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Today's Topic: Managing Free Blocks


- We need to efficiently track free blocks and select one for allocation



1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Today's Topic: Managing Free Blocks

- We need to efficiently track free blocks and select one for allocation



How we implement this one?
What implementation issues should we consider?

1. How do we know how much memory to free given just a pointer?

1. How do we know how much memory to free given just a pointer?

```
p0 = malloc(32)
```

```
⋮
```


```
free(p0)
```

How does free know the size of the block pointed to by p0?

1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Introduction to Header Field


- Store the block size, in bytes, in the first word of each block
 - This word is often called the *header field* or *header*
 - The block size includes the header



1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Introduction to Header Field


- Store the block size, in bytes, in the first word of each block
 - This word is often called the *header field* or *header*
 - The block size includes the header



1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Introduction to Header Field

- Store the block size, in bytes, in the first word of each block
 - This word is often called the *header field* or *header*
 - The block size includes the header




1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Introduction to Header Field

- Store the block size, in bytes, in the first word of each block
 - This word is often called the *header field* or *header*
 - The block size includes the header


Read the header at $p_0 - 1$ to check the block size

```
p0 = malloc(32)
```



⋮

```
free(p0)
```



1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Introduction to Header Field

- Store the block size, in bytes, in the first word of each block
 - This word is often called the *header field* or *header*
 - The block size includes the header

```
p0 = malloc(32)
```

⋮




```
free(p0)
```

1. How do we know how much memory to free given just a pointer?
→ ***The header field stores the size of a block***

1. How do we know how much memory to free given just a pointer?

→ ***The header field stores the size of a block***


2. How do we keep track of the free blocks?



1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Keeping Track of Free Blocks


- Method #1: **Implicit list** using length. It links all blocks, regardless of whether they are allocated or free




1 square = 1 word = 8 bytes
 □ : Free word
 ■ : Allocated word

Keeping Track of Free Blocks

- Method #1: **Implicit list** using length. It links all blocks, regardless of whether they are allocated or free



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




Q. In terms of throughput, which one is good?
 Q. Which one is more likely to cause internal fragmentation?


1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Keeping Track of Free Blocks

- Method #1: **Implicit list** using length. It links all blocks, regardless of whether they are allocated or free



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




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

1 square = 1 word = 8 bytes
 □: Free word
 ■: Allocated word


Keeping Track of Free Blocks

Today, we focus on implicit lists

- Method #1: **Implicit list** using length. It links all blocks regardless of whether they are allocated or free



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




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

We will cover these later

1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Overview of The Implicit Free List

- Method #1: **Implicit list** using length. It links all blocks, regardless of whether they are allocated or free



How to traverse this list?

```
block = start;  
while (block != end) {  
    ...  
    block_size = p[0]  
    offset = block_size / 8  
    block = block + offset  
}
```


Move to the next block

Overview of The Implicit Free List

1 square = 1 word = 8 bytes

□: Free word


■: Allocated word




We need to know that whether the block is **allocated** or **free**

Overview of The Implicit Free List

1 square = 1 word = 8 bytes
□: Free word
■: Allocated word



We need to know that whether the block is **allocated** or **free**




Introduce the tag field!


- 1: Allocated block
- 0: Free block

Overview of The Implicit Free List

1 square = 1 word = 8 bytes
□: Free word
■: Allocated word



We need to know that whether the block is **allocated** or **free**



Introduce the tag field!

- 1: Allocated block
- 0: Free block



Any problems?

For each block, we need both size and allocation status

Memory consumption

(internal fragmentation – 2 words per block)!

Standard Trick



1 square = 1 word = 8 bytes
□ : Free word
■ : Allocated word

Solution: Express both meanings in a single word



How It Can Be Possible?

1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

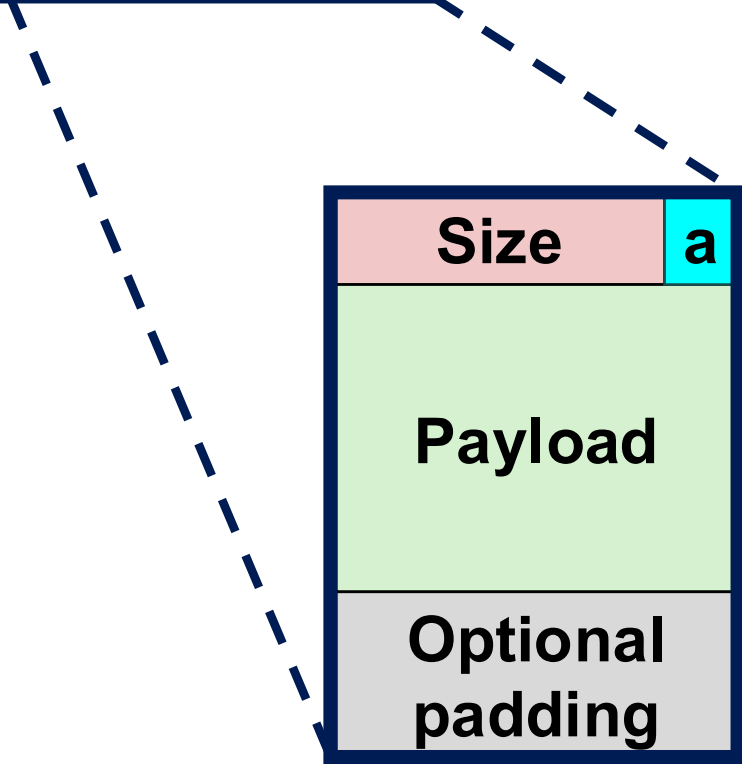



- Remember that a block is aligned to a 16 byte-boundary
- Possible block size:
 - 16 bytes = 010000 (2)
 - 32 bytes = 100000 (2)
 - 48 bytes = 110000 (2)

Observation: 4 low-order bits are always 0

How It Can Be Possible? *

1 square = 1 word = 8 bytes
 □ : Free word
 ■ : Allocated word



- Remember that a block is aligned to a 16 byte-boundary
- Possible block size:
 - 16 bytes = 010000 (2)
 - 32 bytes = 100000 (2)
 - 48 bytes = 110000 (2)


Observation: 4 low-order bits are always 0

Idea: use this bit as 'flag'

- 1: Allocated block
- 0: Free block


Problems?

1 square = 1 word = 8 bytes
□: Free word
■: Allocated word



Problems?

1 square = 1 word = 8 bytes
□: Free word
■: Allocated word




Address of the payload (malloc-returned pointer) is not aligned to 16 bytes ☹️

On x86-64, malloc payloads are typically aligned to 16 bytes because this satisfies data-type alignment requirements and makes block management simpler

Deployment of Unused Blocks

1 square = 1 word = 8 bytes
□: Free word
■: Allocated word




Introduce start block

Payloads are aligned to 16 bytes! 😊

Deployment of Unused Blocks

1 square = 1 word = 8 bytes
□: Free word
■: Allocated word



Introduce start block

We need to indicate end block

Final Version of The Implicit List


1 square = 1 word = 8 bytes

□: Free word

■: Allocated word

heap_start

heap_end



How to traverse this list?

```
block = heap_start;

while (block != heap_end) {
    header = block[0];

    block_size = header & ~0xF;
    allocated = header & 0x1;

    if (allocated == 0) {
        ... // Free block!
    }

    block = block + block_size / 8;
}
```

1. How do we know how much memory to free given just a pointer?
→ ***The header field stores the size of a block***
2. How do we keep track of the free blocks?
→ ***Managed by the linked list (implicit list, explicit list, segregated free list)***

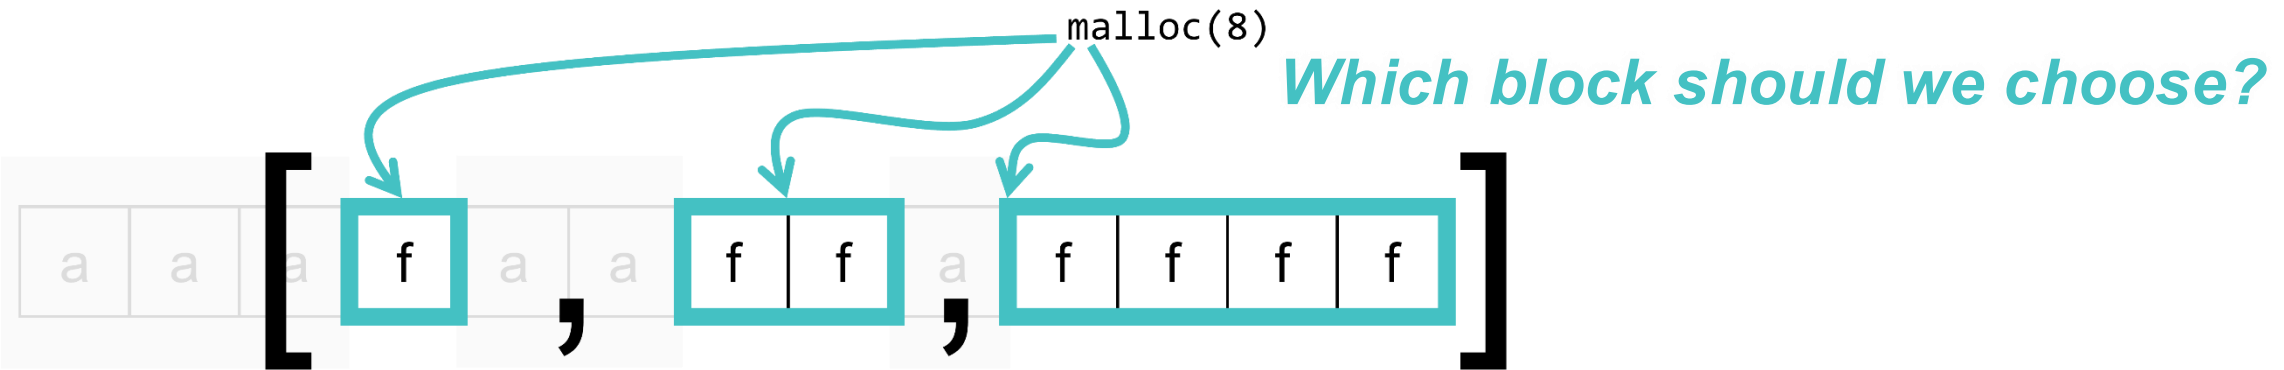
1. How do we know how much memory to free given just a pointer?

→ ***The header field stores the size of a block***

2. How do we keep track of the free blocks?

→ ***Managed by the linked list (implicit list, explicit list, segregated free list)***

3. How do we select one of the free block for allocation?



Placement Policy



- **First fit**
- **Next fit**
- **Best fit**

Placement Policy




- **First fit**
 - Search list from beginning, choose *first* free block that fits
- **Next fit**
- **Best fit**



1 square = 1 word = 8 bytes
□ : Free word
■ : Allocated word

Placement Policy – First Fit


- Search list from beginning, choose **first** free block that fits



malloc(40)
→ 5 words



malloc(5)
→ 3 words



1 square = 1 word = 8 bytes
□: Free word
■: Allocated word


Placement Policy – First Fit

- Search list from beginning, choose **first** free block that fits




Assume a simplified heap (no headers, no end...)


malloc(40)
→ 5 words



Limitation #1: Can take linear time in total number of blocks (allocated and free)




malloc(5)
→ 3 words



1 square = 1 word = 8 bytes
□: Free word
■: Allocated word

Placement Policy – First Fit

- Search list from beginning, choose **first** free block that fits



Assume a simplified heap (no headers, no end...)

Limitation #2: Can cause many small free block at beginning

Limitation #1: Can take linear time in total number of blocks (allocated and free)

malloc(5)
→ 3 words



Placement Policy



- **First fit**
 - Search list from beginning, choose *first* free block that fits
- **Next fit**
- **Best fit**


Placement Policy



- **First fit**
 - Search list from beginning, choose **first** free block that fits
- **Next fit**
 - Like first fit, but search list ***starting where previous search finished***
- **Best fit**


Placement Policy – Next Fit

- Like first fit, but search list *starting where previous search finished*




Assume a simplified heap (no headers, no end...)

malloc(40)
→ 5 words




malloc(5)
→ 3 words



Placement Policy – Next Fit

- Like first fit, but search list *starting where previous search finished*



Assume a simplified heap (no headers, no end...)

Limitation: Some research suggests that fragmentation is worse

Advantage: Often be faster than first fit (avoids rescanning unhelpful blocks)

malloc
→

malloc(5)
→ 3 words



Placement Policy



- **First fit**
 - Search list from beginning, choose *first* free block that fits
- **Next fit**
 - Like first fit, but search list starting where previous search finished
- **Best fit**


Placement Policy



- **First fit**
 - Search list from beginning, choose *first* free block that fits
- **Next fit**
 - Like first fit, but search list starting where previous search finished
- **Best fit**
 - Search the list, choose the *best* free block: fits, with fewest bytes left over



Placement Policy – Best Fit

- Search the list, choose the **best** free block: fits, with fewest bytes left over




Assume a simplified heap (no headers, no end...)

malloc(40)
→ 5 words




malloc(5)
→ 3 words



Placement Policy – Best Fit

- Search the list, choose the **best** free block: fits, with fewest bytes left over



Assume a simplified heap (no headers, no end...)

Limitation: Typically run slower than first fit

Advantage: Keep fragments small (improve memory utilization)

malloc(5)
→ 3 words

malloc(5)
→ 3 words



Placement Policy



- **First fit**

Search list from beginning, choose **first** free block that fits

Which one is better?

Each strategy has its own trade-off.

There is no universally best placement policy.

The final choice should be made based on research-based results.

- **Best fit**

- Search the list, choose the **best** free block: fits, with fewest bytes left over

Summary



1. How do we know how much memory to free given just a pointer?
→ ***The header field stores the size of a block***
2. How do we keep track of the free blocks?
→ ***Managed by the linked list (implicit list, explicit list, segregated free list)***
3. How do we select one of the free block for allocation?
→ ***Placement policy (first fit, next fit, best fit)***

Question?