

CS 107

Lecture 22:

Managing the Heap II

Friday, March 1, 2024

Computer Systems

Winter 2024

Stanford University

Computer Science Department

Reading: Course Reader: x86-64 Assembly

Language, Textbook: Chapter 3.1-3.4

Lecturer: Chris Gregg

```
malloc()  
calloc()  
realloc()  
free()
```



Today's Topics

- Reading: Chapter 9.9
- Programs from class: `/afs/ir/class/cs107/samples/lect21`
- Logistics
 - Bank vault — how is it going?
 - This week's lab: work on A5
- Managing the Heap
 - Tracing the heap
 - How do we track heap allocations?
 - Placement: first-fit, next-fit, best-fit (throughput -vs- utilization)
 - Two different free lists: implicit and explicit
 - Splitting / Coalescing



Tracing the Heap (possible implementation)

```
void *a, *b, *c, *d, *e;
```

← All allocated on the stack:

```
a = malloc(16);
```

```
b = malloc(8);
```

```
c = malloc(24);
```

```
d = malloc(16);
```

```
free(a);
```

```
free(c);
```

```
e = malloc(8);
```

```
b = realloc(b, 24);
```


```
e = realloc(e, 24);
```

```
void *f = malloc(24);
```

	Address	Value
e	0xfffffe820	0x0
d	0xfffffe818	0xabcde
c	0xfffffe810	0xf0123
b	0xfffffe808	0x0
a	0xfffffe800	0xbeef

heap

← 96 bytes →



Tracing the Heap (possible implementation)

```
void *a, *b, *c, *d, *e;
```

← All allocated on the stack:

```
a = malloc(16);
```

```
b = malloc(8);
```

```
c = malloc(24);
```

```
d = malloc(16);
```

```
free(a);
```

```
free(c);
```

```
e = malloc(8);
```

```
b = realloc(b, 24);
```

```
e = realloc(e, 24);
```


```
void *f = malloc(24);
```

	Address	Value
e	0xfffffe820	0x0
d	0xfffffe818	0xabcde
c	0xfffffe810	0xf0123
b	0xfffffe808	0x0
a	0xfffffe800	0xbeef

heap

96 bytes

Each section
represents 4
bytes



Tracing the Heap (possible implementation)

```
void *a, *b, *c, *d, *e;
```

← All allocated on the stack:

```
a = malloc(16);
```

```
b = malloc(8);
```

```
c = malloc(24);
```

```
d = malloc(16);
```

```
free(a);
```

```
free(c);
```

```
e = malloc(8);
```

```
b = realloc(b, 24);
```

```
e = realloc(e, 24);
```


```
void *f = malloc(24);
```

	Address	Value
e	0xfffffe820	0x0
d	0xfffffe818	0xabcde
c	0xfffffe810	0xf0123
b	0xfffffe808	0x0
a	0xfffffe800	0xbeef

heap

96 bytes

Each section
represents 4
bytes



Tracing the Heap (possible implementation)

`void *a, *b, *c, *d, *e;` ← All allocated on the stack:

`a = malloc(16);`

`b = malloc(8);`

`c = malloc(24);`

`d = malloc(16);`

`free(a);`

`free(c);`


`e = malloc(8);`

`b = realloc(b, 24);`

`e = realloc(e, 24);`

`void *f = malloc(24);`

	Address	Value
e	0xfffffe820	0x0
d	0xfffffe818	0xabcde
c	0xfffffe810	0xf0123
b	0xfffffe808	0x0
a	0xfffffe800	0x100



Tracing the Heap (possible implementation)

`void *a, *b, *c, *d, *e;` ← All allocated on the stack:

`a = malloc(16);`

`b = malloc(8);`

`c = malloc(24);`

`d = malloc(16);`

`free(a);`

`free(c);`


`e = malloc(8);`

`b = realloc(b, 24);`

`e = realloc(e, 24);`

`void *f = malloc(24);`

	Address	Value
e	0xfffffe820	0x0
d	0xfffffe818	0xabcde
c	0xfffffe810	0xf0123
b	0xfffffe808	0x110
a	0xfffffe800	0x100



Tracing the Heap (possible implementation)

`void *a, *b, *c, *d, *e;` ← All allocated on the stack:

`a = malloc(16);`

`b = malloc(8);`

`c = malloc(24);`

`d = malloc(16);`

`free(a);`

`free(c);`

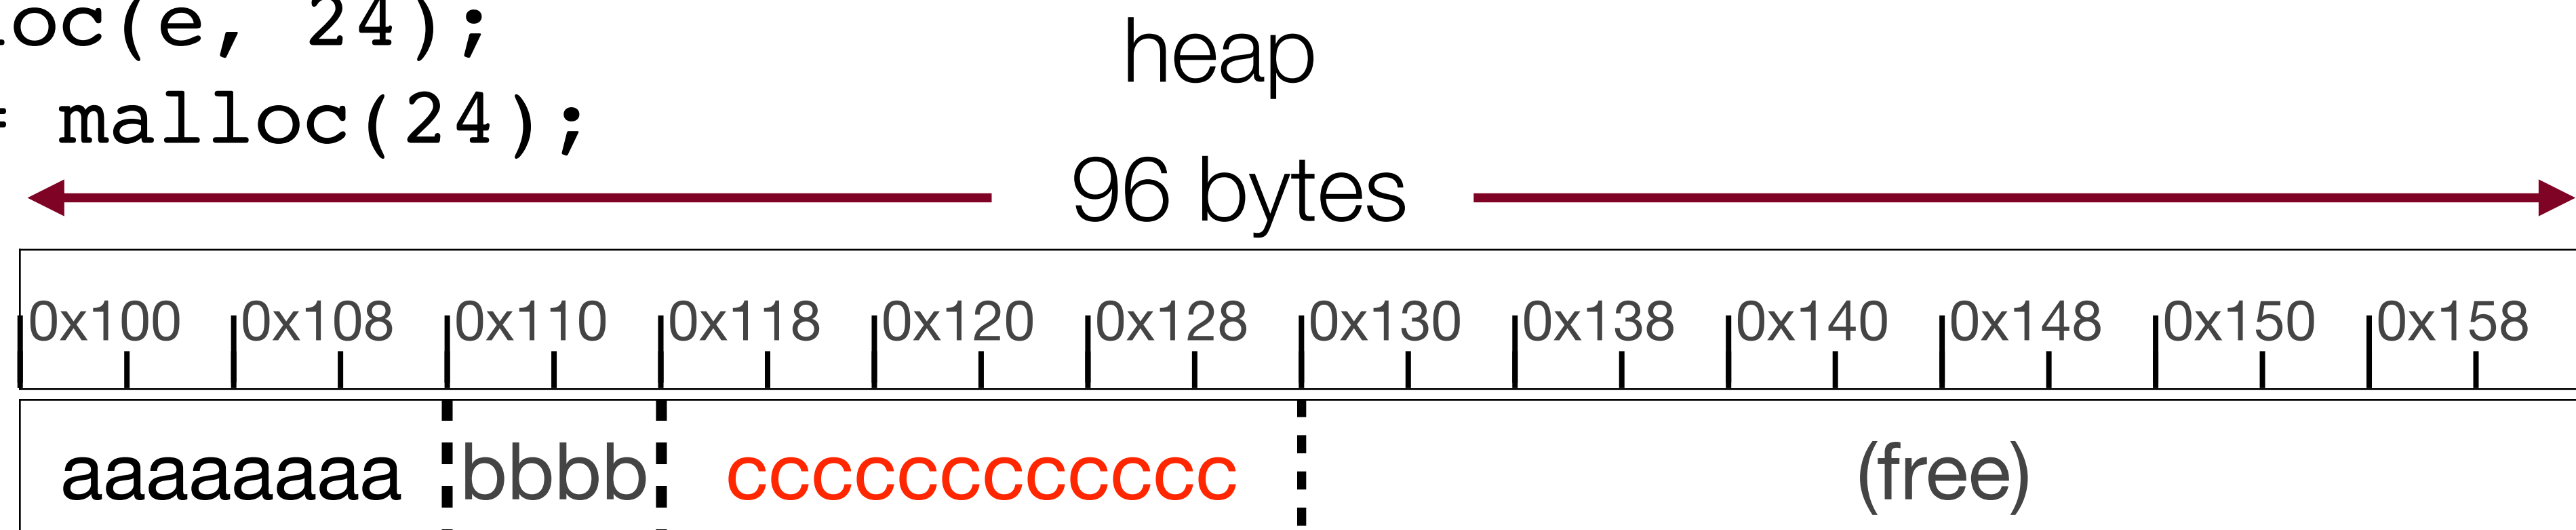
`e = malloc(8);`

`b = realloc(b, 24);`

`e = realloc(e, 24);`

`void *f = malloc(24);`

	Address	Value
e	0xfffffe820	0x0
d	0xfffffe818	0xabcde
c	0xfffffe810	0x118
b	0xfffffe808	0x110
a	0xfffffe800	0x100



Tracing the Heap (possible implementation)

`void *a, *b, *c, *d, *e;` ← All allocated on the stack:

`a = malloc(16);`

`b = malloc(8);`

`c = malloc(24);`

`d = malloc(16);`

`free(a);`

`free(c);`


`e = malloc(8);`

`b = realloc(b, 24);`

`e = realloc(e, 24);`

`void *f = malloc(24);`

	Address	Value
e	0xfffffe820	0x0
d	0xfffffe818	0x130
c	0xfffffe810	0x118
b	0xfffffe808	0x110
a	0xfffffe800	0x100



Tracing the Heap (possible implementation)

`void *a, *b, *c, *d, *e;` ← All allocated on the stack:

`a = malloc(16);`

`b = malloc(8);`

`c = malloc(24);`

`d = malloc(16);`

`free(a);`

`free(c);`


`e = malloc(8);`

`b = realloc(b, 24);`

`e = realloc(e, 24);`

`void *f = malloc(24);`

	Address	Value
e	0xfffffe820	0x0
d	0xfffffe818	0x130
c	0xfffffe810	0x118
b	0xfffffe808	0x110
a	0xfffffe800	0x100



Tracing the Heap (possible implementation)

`void *a, *b, *c, *d, *e;` ← All allocated on the stack:

`a = malloc(16);`

`b = malloc(8);`

`c = malloc(24);`

`d = malloc(16);`

`free(a);`

`free(c);`


`e = malloc(8);`

`b = realloc(b, 24);`

`e = realloc(e, 24);`

`void *f = malloc(24);`

	Address	Value
e	0xfffffe820	0x0
d	0xfffffe818	0x130
c	0xfffffe810	0x118
b	0xfffffe808	0x110
a	0xfffffe800	0x100



Tracing the Heap (possible implementation)

`void *a, *b, *c, *d, *e;` ← All allocated on the stack:

`a = malloc(16);`

`b = malloc(8);`

`c = malloc(24);`

`d = malloc(16);`

`free(a);`

`free(c);`


`e = malloc(8);`

`b = realloc(b, 24);`

`e = realloc(e, 24);`

`void *f = malloc(24);`

	Address	Value
e	0xfffffe820	0x100
d	0xfffffe818	0x130
c	0xfffffe810	0x118
b	0xfffffe808	0x110
a	0xfffffe800	0x100



Tracing the Heap (possible implementation)

`void *a, *b, *c, *d, *e;` ← All allocated on the stack:

`a = malloc(16);`

`b = malloc(8);`

`c = malloc(24);`

`d = malloc(16);`

`free(a);`

`free(c);`


`e = malloc(8);`

`b = realloc(b, 24);`

`e = realloc(e, 24);`

`void *f = malloc(24);`

	Address	Value
e	0xfffffe820	0x100
d	0xfffffe818	0x130
c	0xfffffe810	0x118
b	0xfffffe808	0x110
a	0xfffffe800	0x100



Tracing the Heap (possible implementation)

`void *a, *b, *c, *d, *e;` ← All allocated on the stack:

`a = malloc(16);`

`b = malloc(8);`

`c = malloc(24);`

`d = malloc(16);`

`free(a);`

`free(c);`


`e = malloc(8);`

`b = realloc(b, 24);`

`e = realloc(e, 24);`

`void *f = malloc(24);`

	Address	Value
e	0xfffffe820	0x140
d	0xfffffe818	0x130
c	0xfffffe810	0x118
b	0xfffffe808	0x110
a	0xfffffe800	0x100



Tracing the Heap (possible implementation)

`void *a, *b, *c, *d, *e;` ← All allocated on the stack:

`a = malloc(16);`

`b = malloc(8);`

`c = malloc(24);`

`d = malloc(16);`

`free(a);`

`free(c);`

`e = malloc(8);`

`b = realloc(b, 24);`

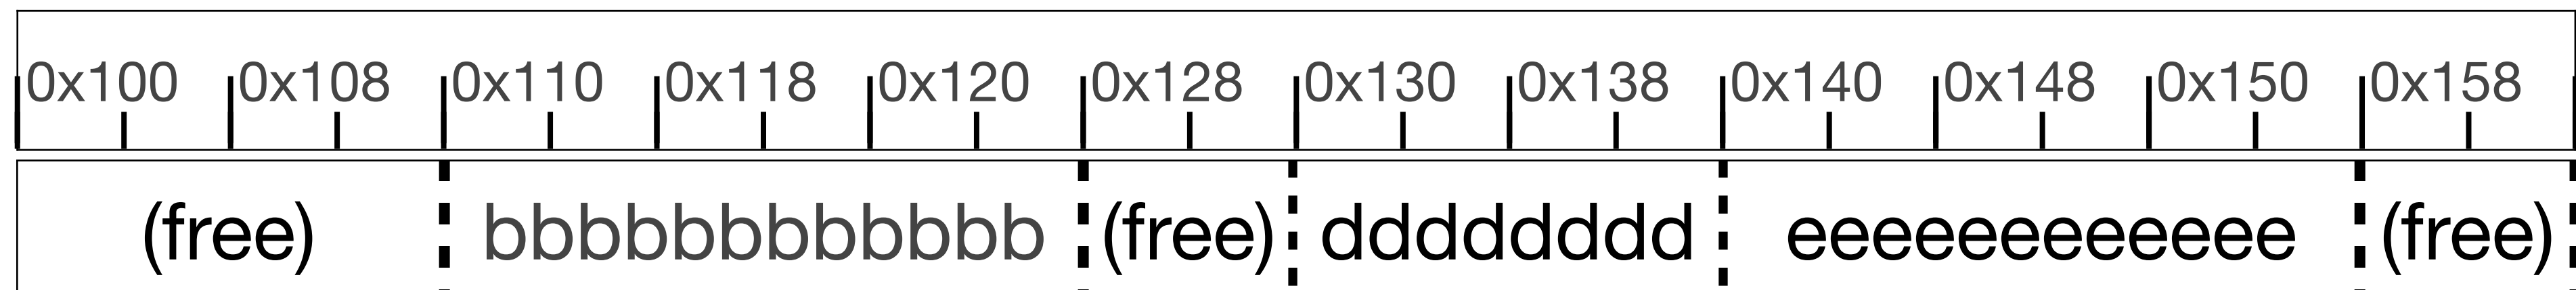
`e = realloc(e, 24);`

`void *f = malloc(24);`

Returns NULL

heap

96 bytes



	Address	Value
e	0xfffffe820	0x140
d	0xfffffe818	0x130
c	0xfffffe810	0x118
b	0xfffffe808	0x110
a	0xfffffe800	0x100
f	0xfffffe7f0	0x0



Heap Allocator Implementation Issues

- How do we track the information in a block?
 - Remember, `free()` is only given a pointer, not a size
- How do we organize/find free blocks?
- How do we pick which free block from available options?
- What do we do with excess space when allocating a block?
- How do we recycle a freed block?



One possibility: Separate list / table

- We could have a separate list or table that holds the free and in-use information.
 - Given an address, how do we look up the information?
 - How do we update the list or table to service `mallocs` and `frees`?
 - How much overhead is there per block?
- The separate list approach could be a reasonable approach (we would have to answer all of the above questions...), but it is not often used in practice, although there are some exceptions:
 - There are some special-case allocators that use this
 - Valgrind uses this, because it needs to keep track of lots more information than just the used / free blocks.




Another Possibility

- A second possibility, and the one that is actually common and used in practice, uses what is called a **block header** to hold the information.
- The block header is actually *stored in the same memory area as the payload, and it generally precedes the payload.*




Another Possibility

- A second possibility, and the one that is actually common and used in practice, uses what is called a **block header** to hold the information.
- The block header is actually *stored in the same memory area as the payload, and it generally precedes the payload.*



Another Possibility

- A second possibility, and the one that is actually common and used in practice, uses what is called a **block header** to hold the information.
- The block header is actually *stored in the same memory area as the payload, and it generally precedes the payload.*




- This is where things start to get a bit tricky. The heap allocator has 96 bytes, and it needs to keep the free block information *in those 96 bytes* (I N C E P T I O N)
- In other words, the heap allocator is using part of the 96 bytes as housekeeping.
- In this case, 8 bytes are taken up with the information that there are 88 Free (F) bytes ahead in the block.



Another Possibility

```
a = malloc(16);
```

	Address	Value
e	0xffffe820	
d	0xffffe818	
c	0xffffe810	
b	0xffffe808	
a	0xffffe800	0x108




- This is where things start to get a bit tricky. The heap allocator has 96 bytes, and it needs to keep the free block information *in those 96 bytes* (I N C E P T I O N)
- In other words, the heap allocator is using part of the 96 bytes as housekeeping.
- Note here that there are now 16 bytes of overhead, because there are two *header blocks*.
- Here, the first 8-byte header block denotes 16 Used bytes, then there is a 16 byte payload, and then there is another 8-byte header to denote the 64 free bytes after.



Another Possibility

```
a = malloc(16);  
b = malloc(8);
```

	Address	Value
e	0xfffffe820	
d	0xfffffe818	
c	0xfffffe810	
b	0xfffffe808	0x120
a	0xfffffe800	0x108




- We changed the header to reflect the fact that 8 bytes are going to to **b**, and we added a header for the remaining 48 bytes.
- Also, note that the pointer returned for **a** is 0x108, and the pointer returned for **b** is 0x120.



Another Possibility

```
a = malloc(16);  
b = malloc(8);  
c = malloc(24);
```

	Address	Value
e	0xfffffe820	
d	0xfffffe818	
c	0xfffffe810	0x130
b	0xfffffe808	0x120
a	0xfffffe800	0x108




•Now we only have 16 bytes left for payloads...let's free some memory.



Another Possibility

```
a = malloc(16);  
b = malloc(8);  
c = malloc(24);  
free(a);
```

	Address	Value
e	0xfffffe820	
d	0xfffffe818	
c	0xfffffe810	0x130
b	0xfffffe808	0x120
a	0xfffffe800	0x108




- Notice that 0x108 will be passed to free. How do we know how much to free?
 - We have to do some pointer arithmetic, so we can grab the 16 from address 0x100 (this diagram does not reflect the `free` yet).
- As you'll find out when writing your heap allocator: the arithmetic is super important.



Another Possibility

```
a = malloc(16);  
b = malloc(8);  
c = malloc(24);  
free(a);
```

	Address	Value
e	0xfffffe820	
d	0xfffffe818	
c	0xfffffe810	0x130
b	0xfffffe808	0x120
a	0xfffffe800	0x108




- The diagram now reflects the free.
- The change to the diagram was subtle — the *only* thing that changed was that the block header now says "F" (free) instead of "U" (used). This is because the data remains, but it can be written over any time after we reassign that block — this can cause bugs! For clarity sake, on the next page, we'll remove the `aaaaaaaa`, but know that the heap allocator doesn't wipe it clean (this another reason that **free** can be fast!)



Another Possibility

```
a = malloc(16);  
b = malloc(8);  
c = malloc(24);  
free(a);  
free(c);
```

	Address	Value
e	0xfffffe820	
d	0xfffffe818	
c	0xfffffe810	0x130
b	0xfffffe808	0x120
a	0xfffffe800	0x108




- Again, 0x130 is passed in to this free, so we need to figure out that we need to look at address 0x128 for the amount of bytes to free.
- On the next slide, we will remove the `cccccccccccc`, but again: it is *not* cleared out, and we're just doing this for the sake of clarity on the diagram.



Another Possibility

```
a = malloc(16);  
b = malloc(8);  
c = malloc(24);  
free(a);  
free(c);
```

	Address	Value
e	0xfffffe820	
d	0xfffffe818	
c	0xfffffe810	0x130
b	0xfffffe808	0x120
a	0xfffffe800	0x108




- This diagram shows one possible result of the `free`. Note that we have actually fragmented our free space! It looks like we only have a block of 24 bytes and then a block of 16 bytes to allocate, yet we should have a block of 48 bytes (we can save a header, too!)



Another Possibility

```
a = malloc(16);  
b = malloc(8);  
c = malloc(24);  
free(a);  
free(c);
```


	Address	Value
e	0xfffffe820	
d	0xfffffe818	
c	0xfffffe810	0x130
b	0xfffffe808	0x120
a	0xfffffe800	0x108



- When we combine free blocks, this is called *coalescing*, and it is an important tool that the heap allocator uses to keep memory as unfragmented as possible.
- We can't coalesce any more because **b** is in the middle, and we absolutely cannot move that block until the program we gave it to frees it.




Implicit Free List



- The method just demonstrated is called an "*implicit free list*," meaning that we have a list of free blocks that we can traverse to find an appropriate fit. The header holds the size of the block and whether it is free (F) or used (U) (note: the free and used information can be stored in 1 bit). To find the next available free block, we must look from the beginning and traverse the list in order.
- As blocks fill up, implicit free lists can cause `malloc` to be slow as the heap fills up — the linear search isn't a terrific method. (We will see another type next lecture!)



Implicit Free List



- Let's answer the questions we posed before:
 - How do we track the information in a block?
 - The header block that holds the bytes in the block and the state (free or used)
 - How do we organize/find free blocks?
 - Linear search, starting from the first block.
 - How do we pick which free block from available options?
 - If the block is free and has enough space we can choose it, though there are other options (covered in the next few slides).
 - What do we do with excess space when allocating a block?
 - If we can fit another header and still have at least a block's worth of space, we can do that. If we can't, it should just become part of the block we are allocating.
 - How do we recycle a freed block?
 - Mark it free, and coalesce if we can.



Placement: first-fit, next-fit, best-fit

The method we have described simply finds the first available block that is free and fits the request, and then starts from the beginning again on a future allocation. This is called a **first-fit** placement policy. One drawback is that you always have to start from the beginning of the heap, and it can be slow. Another drawback is that it can leave "splinters" (small free blocks) towards the beginning of the list. One advantage is that it leaves large blocks towards the end of the list, which allows for larger allocations if necessary.

A second method is called **next-fit**, and was first proposed by Donald Knuth. With next-fit, you start looking for follow-on blocks after the location of the last allocation. If you found a suitable block before, you have a good chance to find another one in the same location. It is still not clear whether next-fit leads to better (or comparable) memory utilization.

The final method is called **best-fit**, and relies on searching the entire heap to find a block that matches the requested allocation the best. The obvious drawback of best-fit is that it requires an exhaustive search of the list.



Splitting and Coalescing

We have already described both splitting and coalescing as used in the implicit free list implementation.


Splitting the memory block is necessary when you have one large block to work with (which is what you will have for the heap allocator assignment). However, the heap allocator can request an increase in the size of the block of memory (using the `sbrk` *system call*), meaning that you could have a policy to use the entire block and just request more. But, we aren't going to cover that low level in this course.

Coalescing does not have to happen when you `free` — you can postpone coalescing until future `mallocs` or `reallocs`, and while it makes `malloc` a bit slower, `free`s are lightning fast.



More on Coalescing: coalescing backwards

Coalescing forwards is straightforward:



If we just freed the 24-byte block, we know exactly where the next block is in order to see if it (and subsequent blocks) are free.


However, what if we had just freed the 8 byte block? How could we coalesce the two blocks?

One way would be to look through the whole list from the beginning, keeping track of where the just-freed block is. But...this is slow.

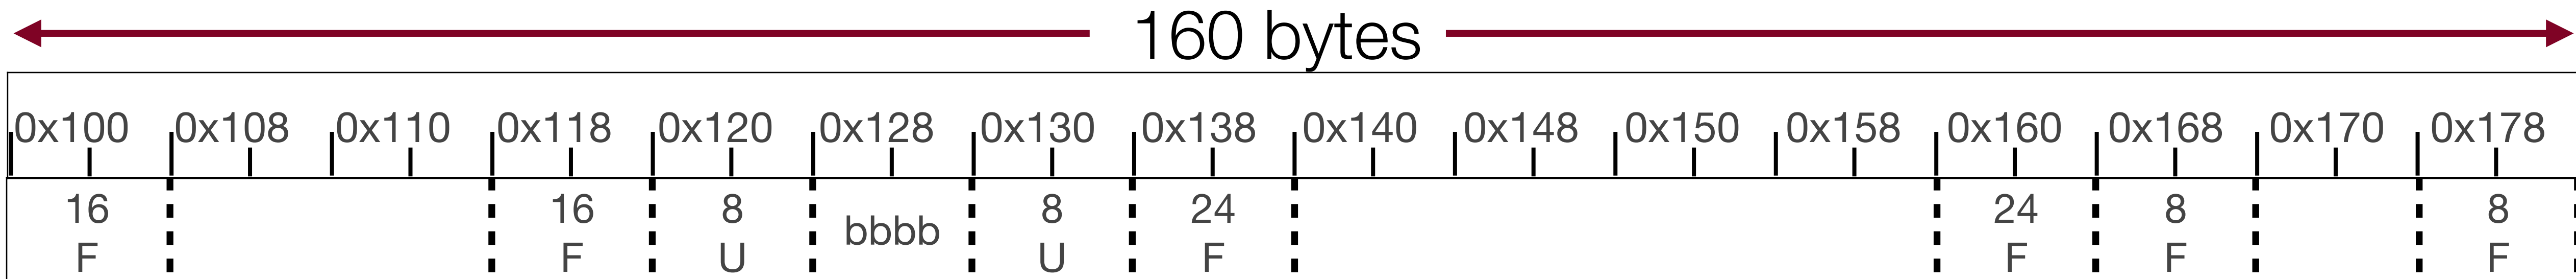


More on Coalescing: coalescing backwards

Coalescing forwards is straightforward:




Another method (described by Knuth) is to keep a footer on each block, as well. The footer is identical to the header, but it refers to the prior bytes. The above list would look like this with headers and footers (assume we were using them the whole time, and we have to add more space because of the extra overhead):




Now, let's say we just free'd the 8 byte block at 0x168. We can look eight bytes back (to 0x160) at the footer for the 24-byte block, and we can see that it is also free, and we can coalesce.

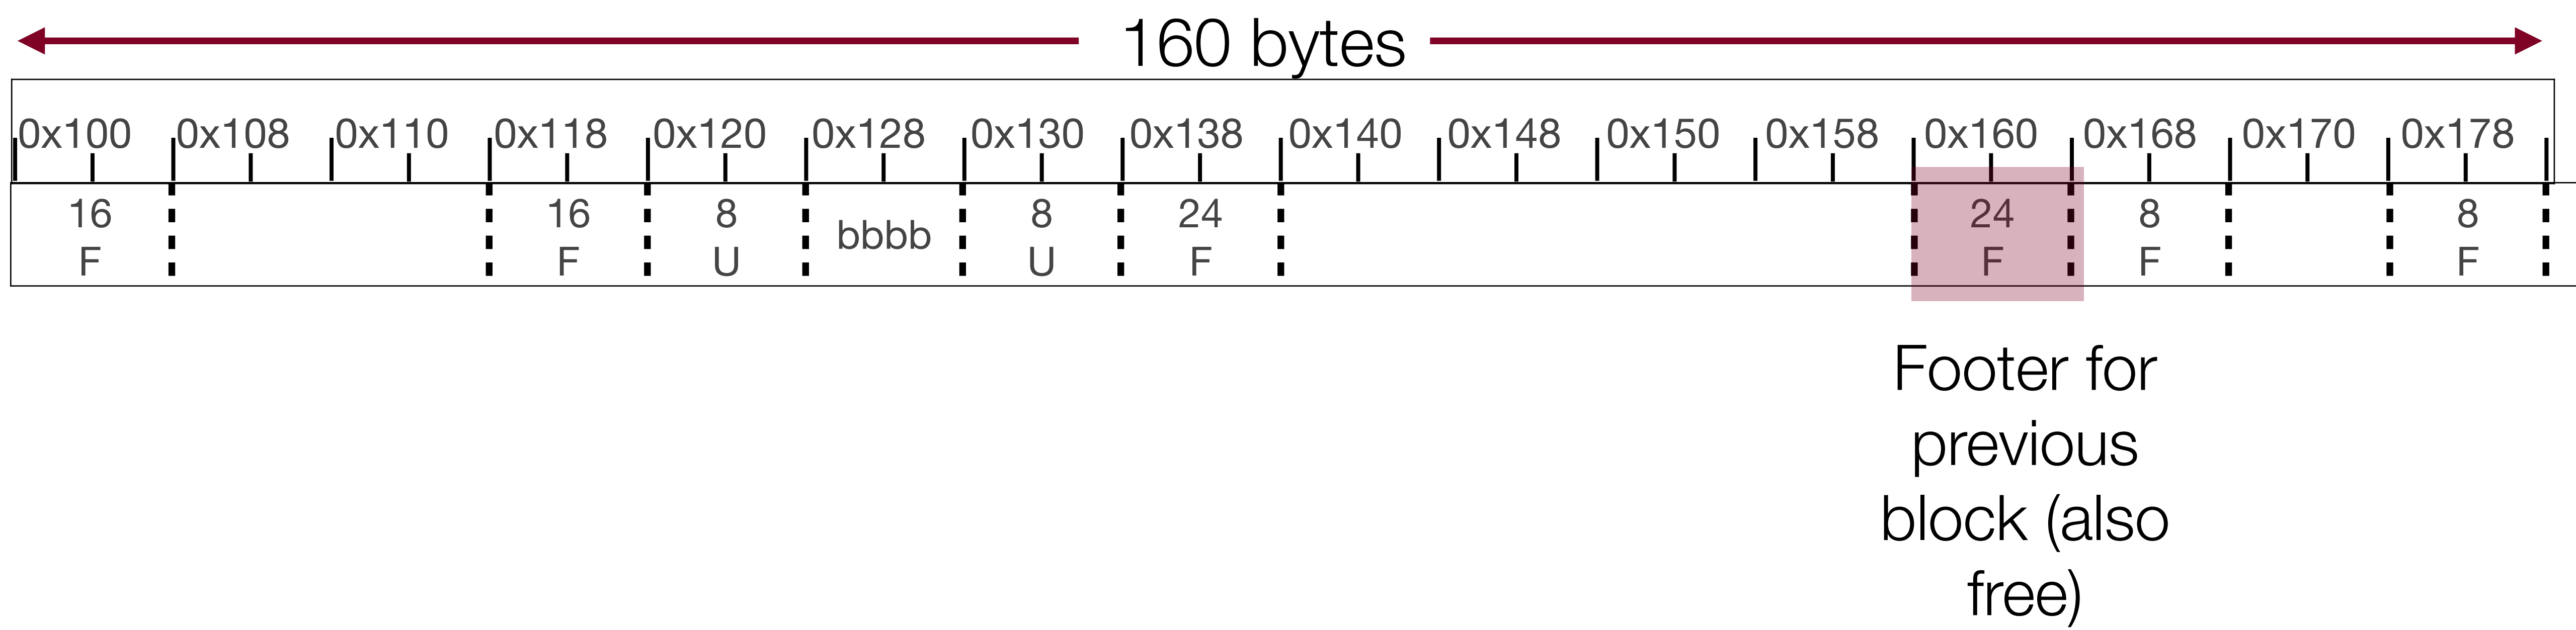
More on Coalescing: coalescing backwards




More on Coalescing: coalescing backwards




More on Coalescing: coalescing backwards



More on Coalescing: coalescing backwards



Entire free
area

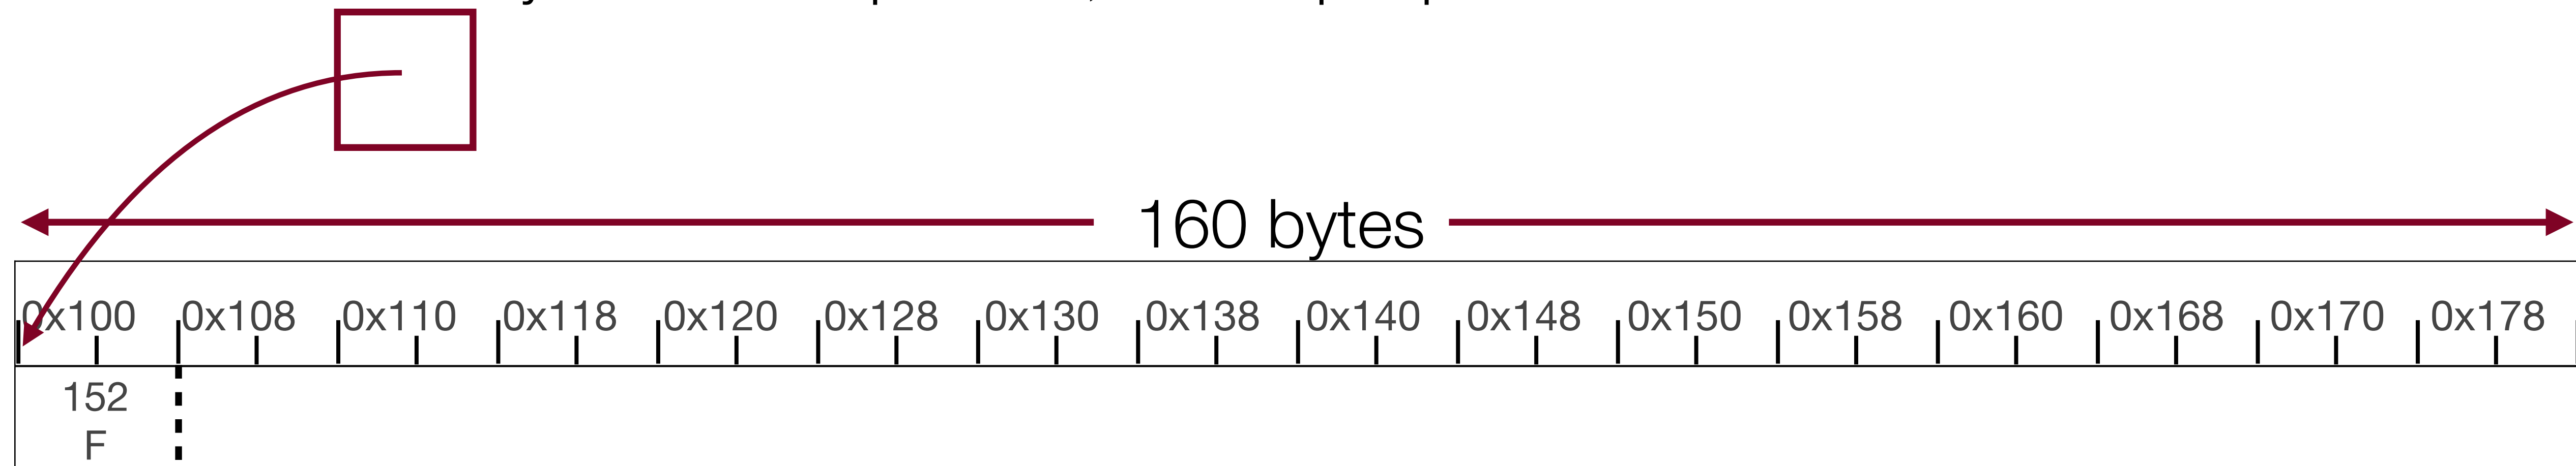


After coalescing
backwards

Explicit Free List

One critical issue with the implicit list is the problem with the linear search to find free blocks.

The *explicit* free list solves this problem by keeping a linked list of free blocks embedded in the memory. This is best shown with an example. As before, let's start with an empty block of memory. With an explicit list, we keep a pointer to the first free block.

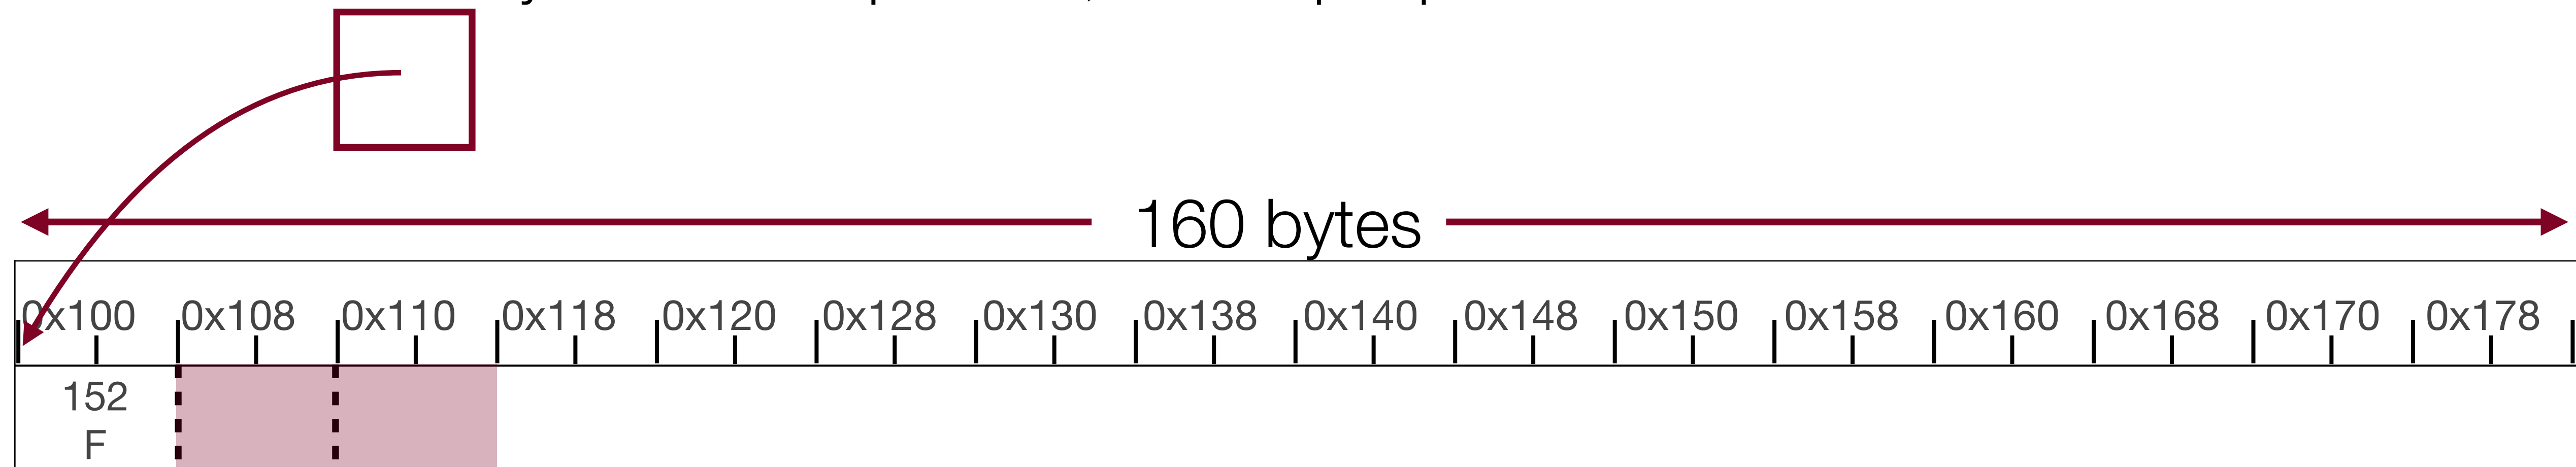


We use two blocks *in the payload* of the free block to point to the *next* and *previous* free blocks.

Explicit Free List

One critical issue with the implicit list is the problem with the linear search to find free blocks.

The *explicit* free list solves this problem by keeping a linked list of free blocks embedded in the memory. This is best shown with an example. As before, let's start with an empty block of memory. With an explicit list, we keep a pointer to the first free block.

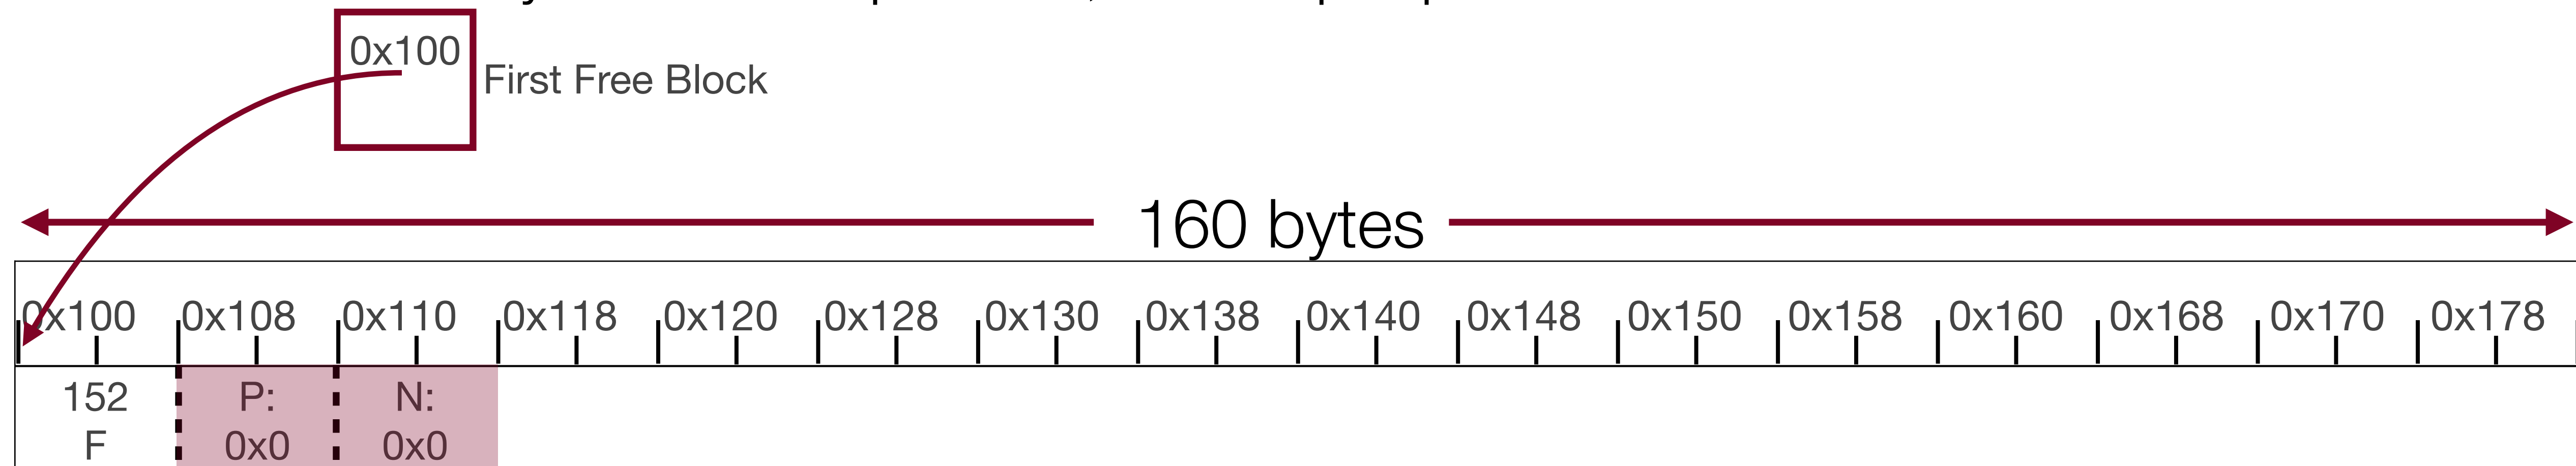


We use two blocks *in the payload* of the free block to point to the *next* and *previous* free blocks.

Explicit Free List

One critical issue with the implicit list is the problem with the linear search to find free blocks.

The *explicit* free list solves this problem by keeping a linked list of free blocks embedded in the memory. This is best shown with an example. As before, let's start with an empty block of memory. With an explicit list, we keep a pointer to the first free block.



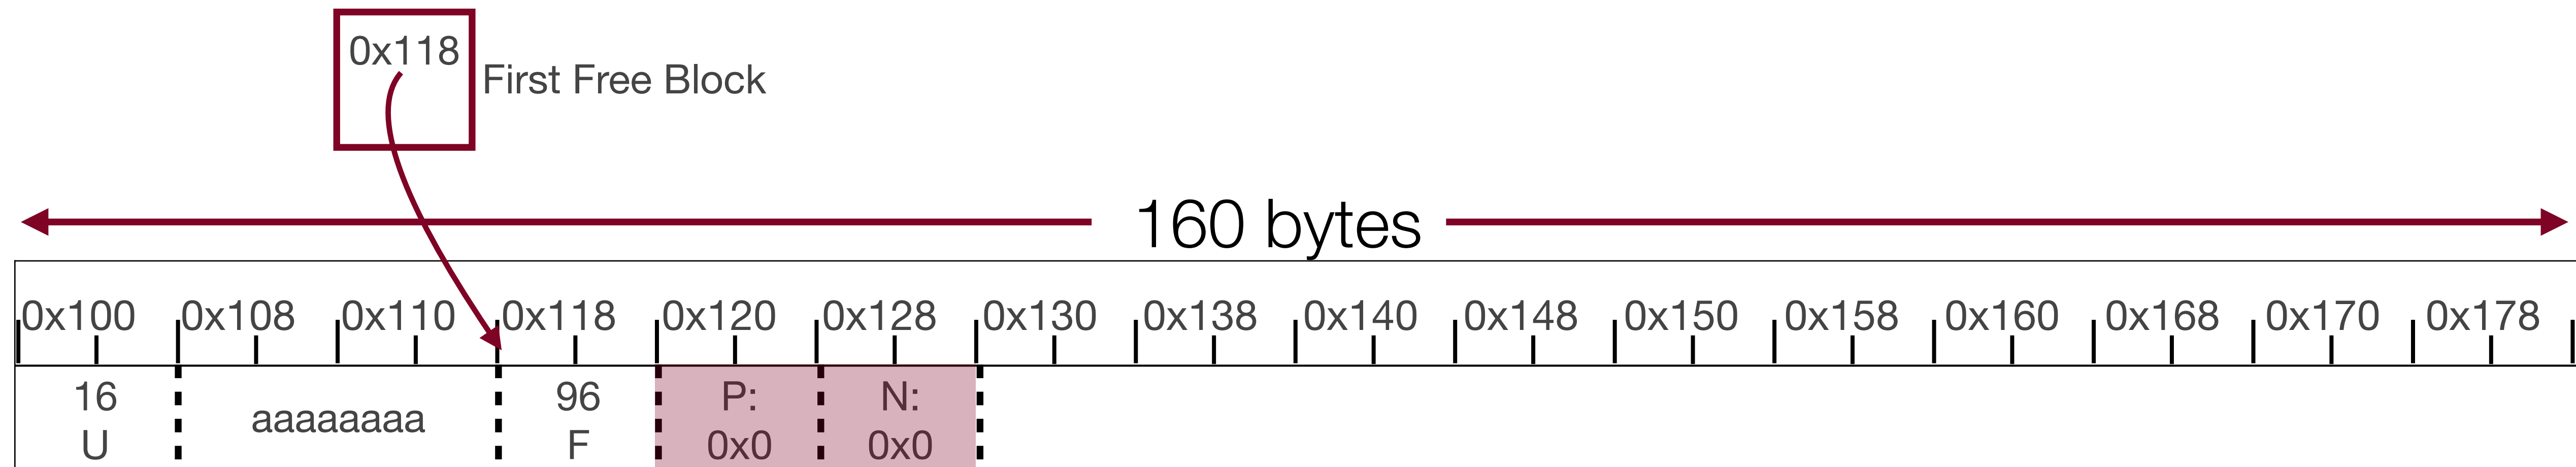
We use two blocks *in the payload* of the free block to point to the *next* and *previous* free blocks. In this case, there aren't any more free blocks, so they are **NULL** pointers.



Explicit Free List

```
a = malloc(16);
```


If we malloc 16, then we allocate as we would in the implicit list, but now we have a pointer to the next free block, and that block still has no previous or next free block.



Explicit Free List

```
a = malloc(16);  
b = malloc(8);  
c = malloc(24);
```


We continue the process. Note that we must leave at least 16 bytes in a block to save room for pointers if we eventually free (e.g., **b** has more space than it requested).



Explicit Free List


```
a = malloc(16);  
b = malloc(8);  
c = malloc(24);  
free(b);
```

Now when we free `b`, we point to the newly free'd memory, and update the pointers



Explicit Free List

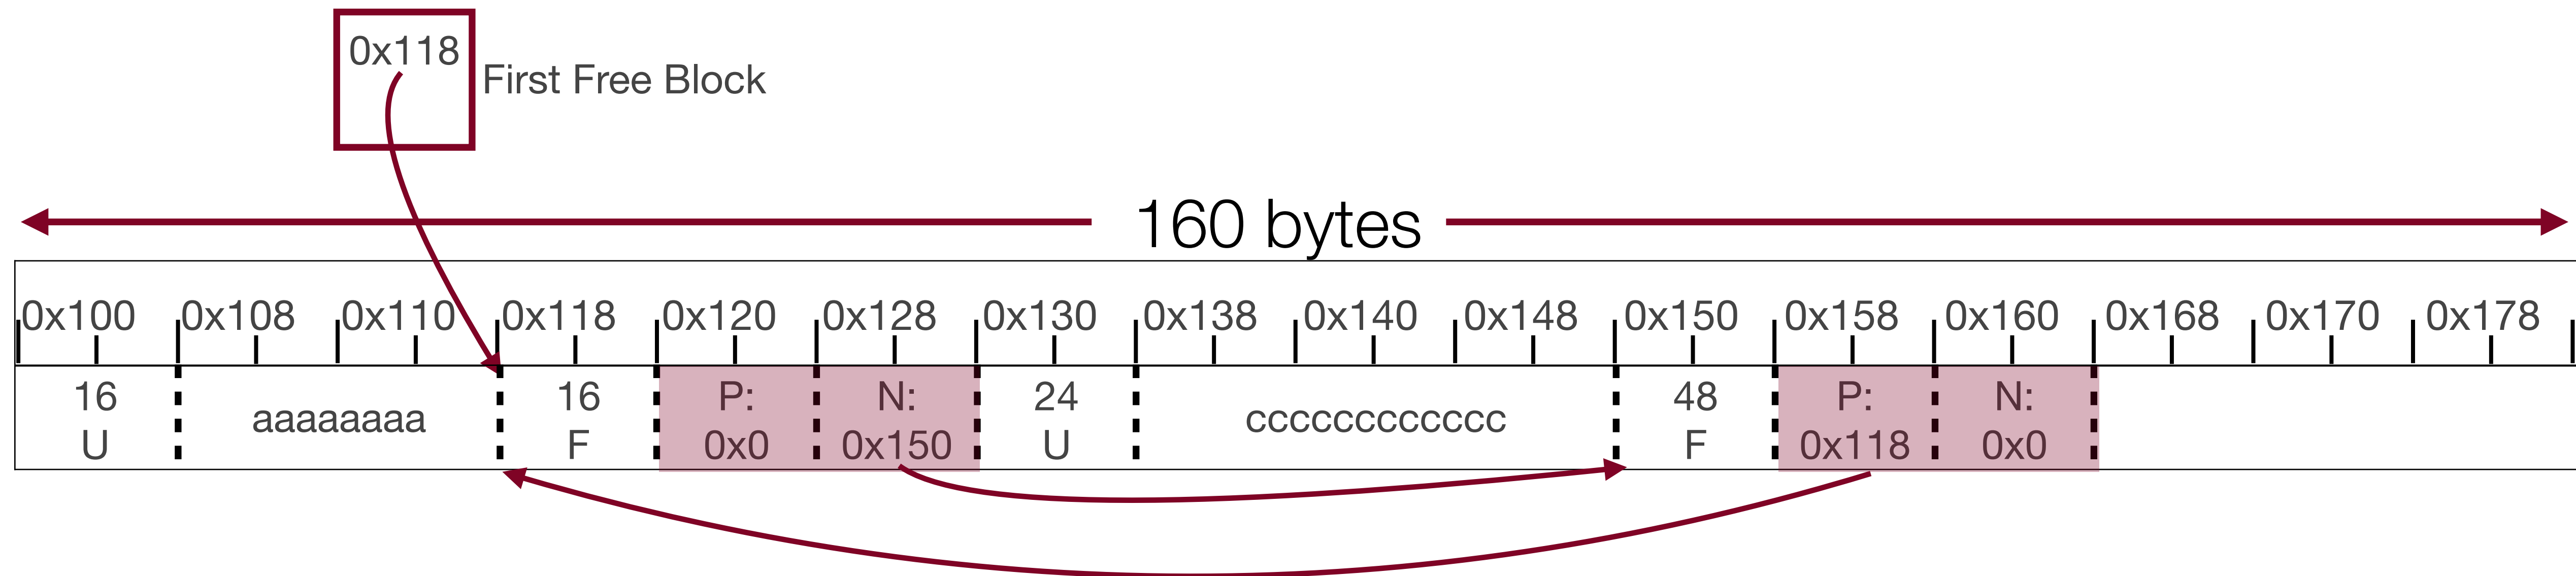
Why is this better than the implicit free list?



Explicit Free List

Why is this better than the implicit free list?

- We can now traverse only the free blocks!
- This is much faster than traversing the whole list.
- For instance, if we now tried to malloc 24 bytes, we would only need to look through two blocks (0x118 and then 0x150) to find enough space.



- More on explicit free lists next lecture!



References and Advanced Reading

References:

- The textbook is the best reference for this material.
- Here are more slides from a similar course: https://courses.engr.illinois.edu/cs241/sp2014/lecture/06-HeapMemory_sol.pdf

Advanced Reading:

- Implementation tactics for a heap allocator: <https://stackoverflow.com/questions/2946604/c-implementation-tactics-for-heap-allocators>

