

CS 153

Design of Operating Systems

Fall 20

Lecture 23: Dynamic Memory

Instructor: Chengyu Song



Some slides modified from originals by Dave O'hallaron

Dynamic Memory Allocation

- Where is it used?
 - ◆ Userspace heap (`malloc`)
 - ◆ Kernel heap (`kmalloc`)
 - ◆ Physical memory allocator
 - ◆ Problems are similar, but specific sometimes force different solutions

Dynamic Memory Allocation

- Programmers use *dynamic memory allocators* (such as `malloc`) to acquire VM at run time.
 - ◆ For data structures whose size is only known at runtime.
- Dynamic memory allocators manage an area of process virtual memory known as the *heap*.



Dynamic Memory Allocation

- Allocator maintains heap as collection of variable sized *blocks*, which are either *allocated* or *free*
- Types of allocators
 - ◆ ***Explicit allocator***: application allocates and frees space
 - » E.g., malloc and free in C
 - ◆ ***Implicit allocator***: application allocates, but does not free space
 - » E.g. garbage collection in Java, ML, and Lisp
- Will discuss explicit memory allocation

The malloc Package

```
#include <stdlib.h>
```

```
void *malloc(size_t size)
```

- ◆ Successful:
 - » Returns a pointer to a memory block of at least size bytes (typically) aligned to 8-byte boundary
 - » If size == 0, returns NULL
- ◆ Unsuccessful: returns NULL (0) and sets errno

```
void free(void *p)
```

- ◆ Returns the block pointed at by p to pool of available memory
- ◆ p must come from a previous call to malloc or realloc
- Other functions
 - ◆ calloc: version of malloc that initializes allocated block to 0.
 - ◆ realloc: Changes the size of a previously allocated block.
 - ◆ sbrk: used internally by allocators to grow or shrink the heap.

Allocation Example

`p1 = malloc(4)`



`p2 = malloc(5)`



`p3 = malloc(6)`



`free(p2)`



`p4 = malloc(2)`



Constraints

- Applications
 - ◆ Can issue arbitrary sequence of `malloc` and `free` requests
 - ◆ `free` request must be to a `malloc`'d block
- Allocators
 - ◆ Can't control number or size of allocated blocks
 - ◆ Must respond immediately to `malloc` requests
 - » *i.e.*, can't reorder or buffer requests
 - ◆ Must allocate blocks from free memory
 - » *i.e.*, can only place allocated blocks in free memory
 - ◆ Must align blocks so they satisfy all alignment requirements
 - » 8 byte alignment for GNU `malloc` (`libc malloc`) on Linux boxes
 - ◆ Can manipulate and modify only free memory
 - ◆ Can't move the allocated blocks once they are `malloc`'d
 - » *i.e.*, compaction is not allowed

Goals


- Given some sequence of `malloc` and `free` requests:
 - ◆ $R_0, R_1, \dots, R_k, \dots, R_{n-1}$
- Goals: maximize throughput and peak memory utilization
 - ◆ These goals are often conflicting
- Throughput:
 - ◆ Number of completed requests per unit time
- Utilization:
 - ◆ Percentage of the heap that is utilized
 - ◆ Poor memory utilization caused by fragmentation, or poor allocation policies

Implementation Issues

- How do we know how much memory to free given just a pointer?
- How do we keep track of the free blocks?
- What do we do with the extra space when allocating a structure that is smaller than the free block it is placed in?
- How do we pick a block to use for allocation -- many might fit?
- How do we reinsert freed block?


Knowing How Much to Free

- Standard method
 - ◆ Keep the length of a block in the word preceding the block.
 - » This word is often called the *header field* or *header*
 - ◆ Requires an extra word for every allocated block




Keeping Track of Free Blocks

- Method 1: *Implicit list* using length—links all blocks



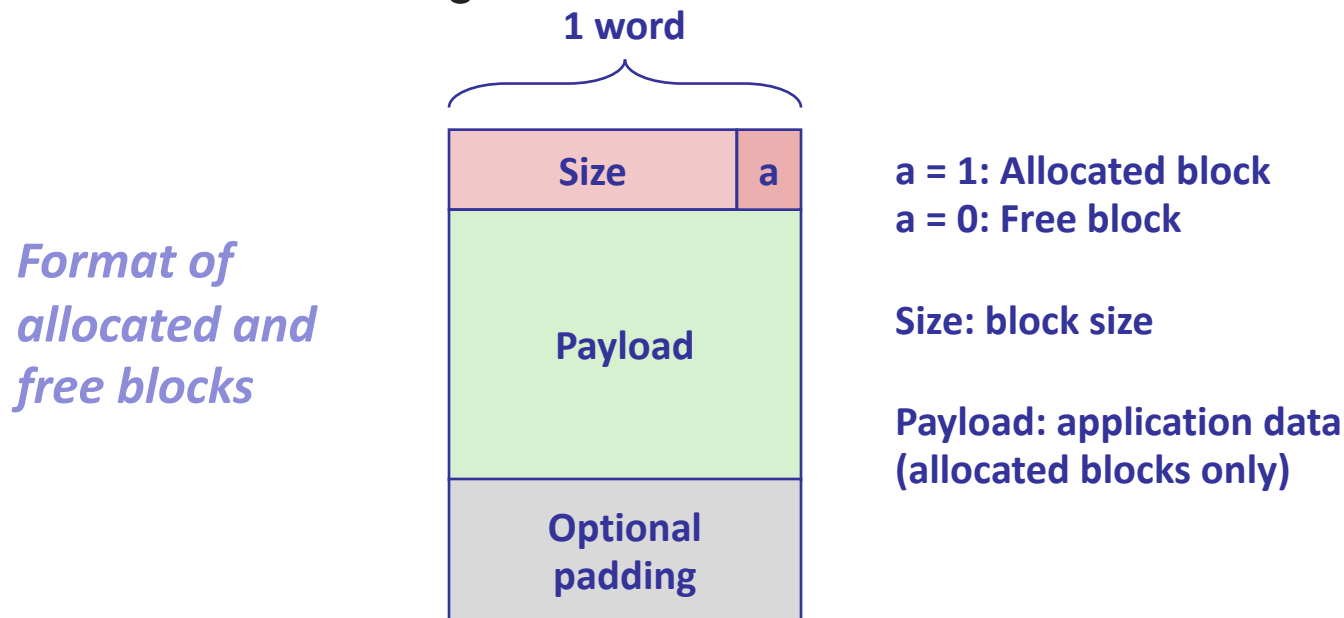
- Method 2: *Explicit list* among the free blocks using pointers




- Method 3: *Segregated free list*
 - ◆ Different free lists for different size classes
- Method 4: *Blocks sorted by size*
 - ◆ Can use a balanced tree (e.g. Red-Black tree) with pointers within each free block, and the length used as a key

Method 1: Implicit List

- For each block we need both size and allocation status
 - ◆ Could store this information in two words: wasteful!
- Standard trick
 - ◆ If blocks are aligned, some low-order address bits are always 0
 - ◆ Instead of storing an always-0 bit, use it as a allocated/free flag
 - ◆ When reading size word, must mask out this bit



Implicit Free List Example



Double-word
aligned

Allocated blocks: shaded

Free blocks: unshaded

Headers: labeled with size in bytes/allocated bit

Implicit List: Finding a Free Block

- *First fit:*
 - ◆ Search list from beginning, choose *first* free block that fits:

```
p = start;
while ((p < end) &&          \\ not passed end
      ((*p & 1) ||          \\ already allocated
      (*p <= len)))        \\ too small
  p = p + (*p & -2);        \\ goto next block (word addressed)
```


- ◆ Can take linear time in total number of blocks (allocated and free)
- ◆ In practice it can cause “splinters” at beginning of list

Implicit List: Finding a Free Block


- *Next fit:*
 - ◆ Like first fit, but search list starting where previous search finished
 - ◆ Should often be faster than first fit: avoids re-scanning unhelpful blocks
 - ◆ Some research suggests that fragmentation is worse
- *Best fit:*
 - ◆ Search the list, choose the **best** free block: fits, with fewest bytes left over
 - ◆ Keeps fragments small—usually helps fragmentation
 - ◆ Will typically run slower than first fit

Implicit List: Allocating in Free Block

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



addblock(p, 4)




```
void addblock(ptr p, int len) {
    int newsize = ((len + 1) >> 1) << 1; // round up to even
    int oldsize = *p & -2; // mask out low bit
    *p = newsize | 1; // set new length
    if (newsize < oldsize)
        *(p+newsize) = oldsize - newsize; // set length in remaining
                                            // part of block
}
```


Implicit List: Freeing a Block

- Simplest implementation:
 - ◆ Need only clear the “allocated” flag

```
void free_block(ptr p) { *p = *p & -2 }
```
 - ◆ But can lead to “false fragmentation”




malloc(5) **Oops!**

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

Implicit List: Coalescing

- Join (*coalesce*) with next/previous blocks, if they are free
 - ◆ Coalescing with next block




```
void free_block(ptr p) {  
    *p = *p & -2;           // clear allocated flag  
    next = p + *p;         // find next block  
    if ((*next & 1) == 0)  
        *p = *p + *next;   // add to this block if  
                            // not allocated  
}
```


- ◆ But how do we coalesce with *previous* block?

Implicit List: Bidirectional Coalescing

- *Boundary tags* [Knuth73]
 - ◆ Replicate size/allocated word at “bottom” (end) of free blocks
 - ◆ Allows us to traverse the “list” backwards, but requires extra space
 - ◆ Important and general technique!



*Format of
allocated and
free blocks*

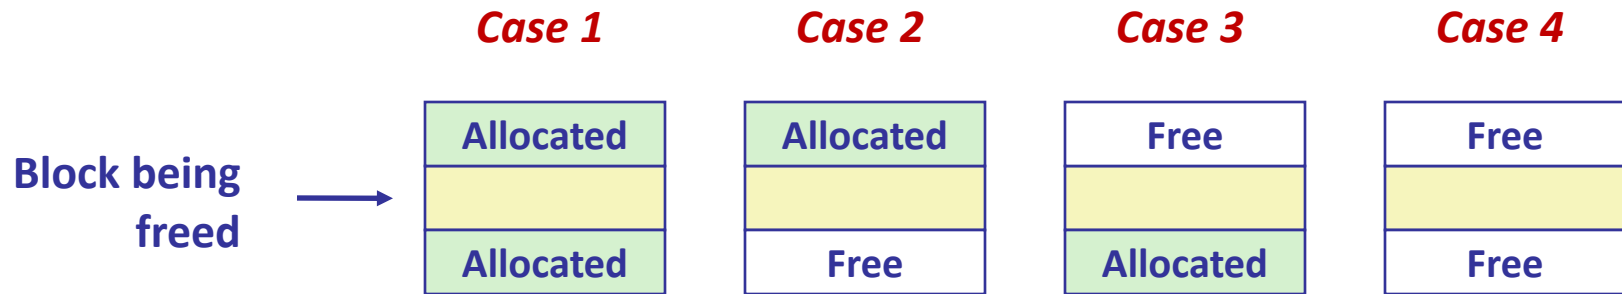


a = 1: Allocated block
a = 0: Free block


Size: Total block size

Payload: Application data
(allocated blocks only)

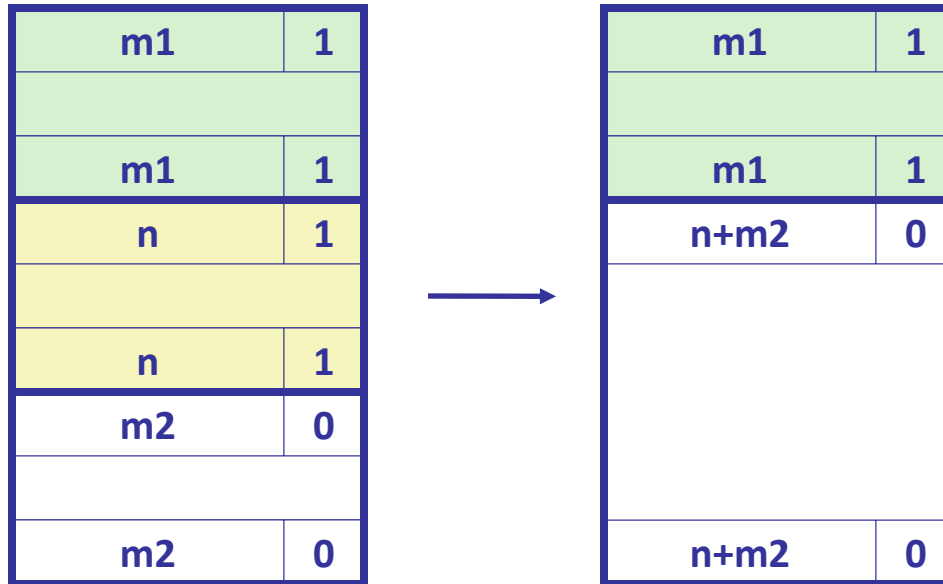
Constant Time Coalescing



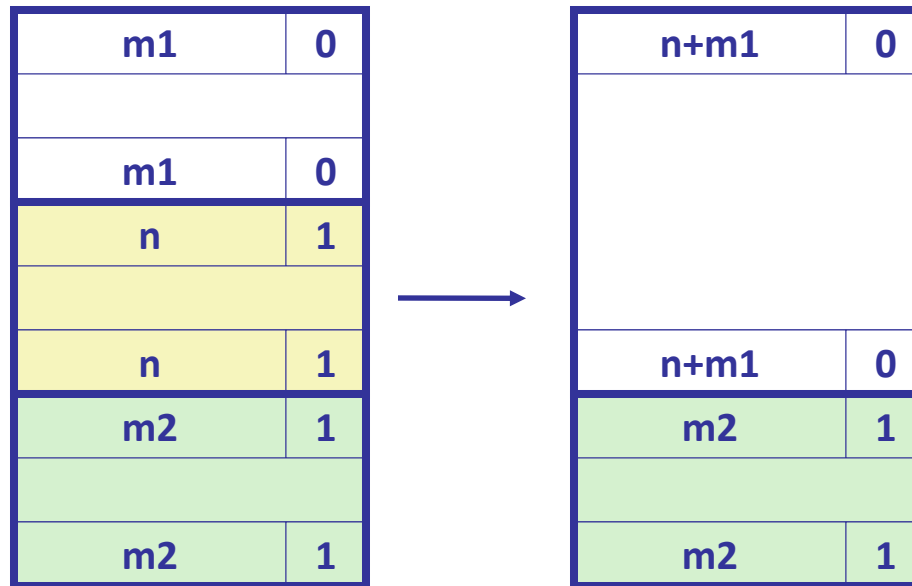
Constant Time Coalescing (Case 1)




Constant Time Coalescing (Case 2)



Constant Time Coalescing (Case 3)




Constant Time Coalescing (Case 4)




Explicit Free Lists

Allocated (as before)



Free




- Maintain list(s) of *free* blocks, not *all* blocks
 - ◆ The “next” free block could be anywhere
 - » So we need to store forward/back pointers, not just sizes
 - ◆ Still need boundary tags for coalescing
 - ◆ Luckily we track only free blocks, so we can use payload area

Explicit Free Lists

- Logically:



- Physically: blocks can be in any order



Allocating From Explicit Free Lists


conceptual graphic

Before



After


(with splitting)



Freeing With Explicit Free Lists

- **Insertion policy:** Where in the free list do you put a newly freed block?
 - ◆ **LIFO (last-in-first-out) policy**
 - » Insert freed block at the beginning of the free list
 - » **Pro:** simple and constant time
 - » **Con:** studies suggest fragmentation is worse than address ordered
 - ◆ **Address-ordered policy**
 - » Insert freed blocks so that free list blocks are always in address order:
$$addr(prev) < addr(curr) < addr(next)$$
 - » **Con:** requires search
 - » **Pro:** studies suggest fragmentation is lower than LIFO


Constant Time Coalescing



Freeing With a LIFO Policy (Case 1)


conceptual graphic

Before




- Insert the freed block at the root of the list

After




Freeing With a LIFO Policy (Case 2)

conceptual graphic




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




Freeing With a LIFO Policy (Case 3)

conceptual graphic




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




Freeing With a LIFO Policy (Case 4)

conceptual graphic



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




Explicit List Summary


- Comparison to implicit list:
 - ◆ Allocate 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 needs to splice blocks in and out of the list
 - ◆ Some extra space for the links (2 extra words needed for each block)
 - » Does this increase internal fragmentation?
- Most common use of linked 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

- Method 1: *Implicit list* using length—links all blocks




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



- Method 3: *Segregated free list*
 - ◆ Different free lists for different size classes
- Method 4: *Blocks sorted by size*
 - ◆ Can use a balanced 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



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

Seglist Allocator

- Given an array of free lists, each one for some size class
- To allocate a block of size n :
 - ◆ Search appropriate free list for block of size $m > n$
 - ◆ If an appropriate block is found:
 - » Split block and place fragment on appropriate list (optional)
 - ◆ If no block is found, try next larger class
 - ◆ Repeat until block is found
- If no block is found:
 - ◆ Request additional heap memory from OS (using **sbrk()**)
 - ◆ Allocate block of n bytes from this new memory
 - ◆ Place remainder as a single free block in largest size class.

Seglist Allocator (cont.)

- To free a block:
 - ◆ Coalesce and place on appropriate list (optional)
- Advantages of seglist allocators
 - ◆ Higher throughput
 - » log time for power-of-two size classes
 - ◆ Better memory utilization
 - » First-fit search of segregated free list approximates a best-fit search of entire heap.
 - » Extreme case: Giving each block its own size class is equivalent to best-fit.

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)

Implicit Memory Management: Garbage Collection

- *Garbage collection*: automatic reclamation of heap-allocated storage—application never has to free

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

- Common in functional languages, scripting languages, and modern object oriented languages:
 - ◆ Lisp, ML, Java, Perl, Python
- Variants (“conservative” garbage collectors) exist for C and C++
 - ◆ However, cannot necessarily collect all garbage

Garbage Collection


- How does the memory manager 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 there are no pointers to them
- Must make certain assumptions about pointers
 - ◆ Memory manager can distinguish pointers from non-pointers
 - ◆ All pointers point to the start of a block
 - ◆ 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)
 - ◆ Collection based on lifetimes
 - » Most allocations become garbage very soon
 - » So focus reclamation work on zones of memory recently allocated
- For more information
 - ◆ Jones and Lin, “*Garbage Collection: Algorithms for Automatic Dynamic Memory*”, John Wiley & Sons, 1996.

Memory as a Graph

- We view memory as a directed graph
 - ◆ Each 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, locations on the stack, 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)

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 head 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



Assumptions For a Simple Implementation

- Application
 - ◆ `new(n)`: returns pointer to new block with all locations cleared
 - ◆ `read(b, i)`: read location `i` of block `b` into register
 - ◆ `write(b, i, v)`: write `v` into location `i` of block `b`
- Each block will have a header word
 - ◆ addressed as `b[-1]`, for a block `b`
 - ◆ Used for different purposes in different collectors
- Instructions used by the Garbage Collector
 - ◆ `is_ptr(p)`: determines whether `p` is a pointer
 - ◆ `length(b)`: returns the length of block `b`, not including the header
 - ◆ `get_roots()`: returns all the roots

Mark and Sweep (cont.)

Mark using depth-first traversal of the memory graph


```
ptr mark(ptr p) {  
    if (!is_ptr(p)) return;           // do nothing if not pointer  
    if (markBitSet(p)) return;       // check if already marked  
    setMarkBit(p);                   // set the mark bit  
    for (i=0; i < length(p); i++)   // call mark on all words  
        mark(p[i]);                 // in the block  
    return;  
}
```

Sweep using lengths to find next block


```
ptr sweep(ptr p, ptr end) {  
    while (p < end) {  
        if markBitSet(p)  
            clearMarkBit();  
        else if (allocateBitSet(p))  
            free(p);  
        p += length(p);  
    }  
}
```

Conservative Mark & Sweep in C

- A “conservative garbage collector” for C programs
 - ◆ `is_ptr()` determines if a word is a pointer by checking if it points to an allocated block of memory
 - ◆ But, in C pointers can point to the middle of a block



- So how to find the beginning of the block?
 - ◆ Can use a balanced binary tree to keep track of all allocated blocks (key is start-of-block)
 - ◆ Balanced-tree pointers can be stored in header (use two additional words)



Left: smaller addresses
Right: larger addresses