

CSc 553

Principles of Compilation

6 : Malloc Algorithms

Department of Computer Science
University of Arizona

collberg@gmail.com

Copyright © 2011 Christian Collberg

- Suppose we want to write a program that keeps track of a company's employees. The information for each employee will be kept in a structure:

- a structure is created when an employee is hired, and
- deleted when the employee is fired or quits.

How should the structures be allocated and freed?

Static allocation

- In assembly code, use the `.word` directive. Memory is allocated before the program runs.
 - Requires allocating the maximum number of employee structures.
 - Not all may be used, wasting space.
 - Increasing the maximum requires reassembling the program.
- This is not what we want.

Stack allocation

- Use a stack, similar to stack frames for subroutines.
 - A stack supports push and pop.
 - A stack is a First-In-First-Out data structure.
 - Employees must quit or be fired in the reverse order in which they were hired.
- This is not what we want.

- We need a fully dynamic memory allocation strategy. Employees can quit or be fired in any order relative to their hiring.
- The data structure used to provide general memory allocation such as this is called a heap. It's not the same type of heap you learned about (or will learn about) in a data structures course (those heaps were used to implement a priority queue). Our heap is simply a data structure for keeping track of memory.

- The memory of a running program is divided into three segments:

text: the program instructions

stack: stack frames for subroutines

heap: program data, both static and dynamic.



- The stack grows downward from the highest address.

- The heap grows upward from the top of the text. Initially the heap contains the static data from the program, and one big free region to hold dynamic data. Once this space is used up, the operating system is invoked to increase the size of the heap by moving the boundary up towards the stack.
- Each program has its own heap and stack.
- When the stack and heap collide your program runs out of memory.

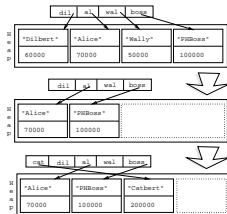
- The heap starts with one big free region, but after the program has been running for a while it will probably be divided into some regions that are being used by the program, and some that are not (employees have been hired and fired). In the following, shaded regions of memory are in-use (they contain data that are being used by the program). Blank regions are not in-use (free).



```

class Employee {String name; int salary;}
class Database {
    public static void main () {
        Employee dil = new Employee("Dilbert",60000);
        Employee al = new Employee("Alice",70000);
        Employee wal = new Employee("Wally",50000);
        Employee boss = new Employee("PHBoss",100000);
        wal = null; dil = null;
        Employee cat = new Employee("Catbert",200000);
    } }

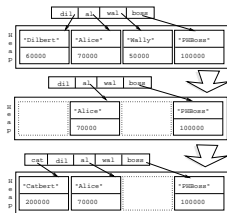
```

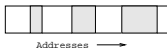


```

program P;
    record Employee String name; int salary; end
begin
    Employee dil = new Employee("Dilbert",60000);
    Employee al = new Employee("Alice",70000);
    Employee wal = new Employee("Wally",50000);
    Employee boss = new Employee("PHBoss",100000);
    dispose(wal); dispose(dil);
    Employee cat = new Employee("Catbert",200000);
end.

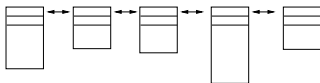
```





- When the OS increases the size of the heap, more free memory is added to the right of this picture.
- When the program wants memory from the heap, we must allocate some from one of the free regions. For each free region we need to know its starting address and its size.

- Where should this information be kept? In an array? This has the same problems as allocating the employee structures themselves in an array. A linked-list called the free list is a better choice. Since the regions we are linking are free space, we can put the linked-list structures in the free regions.

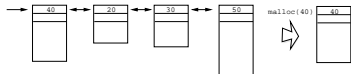


- Think of each free region as a structure. Different regions are different sizes, so the structures have different sizes, but all begin with a standard header. The header contains the links for making the linked-list of free regions, and the size of the region.
- The previous picture of memory looks like the following with the free-list added.

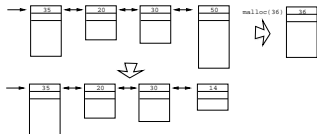


- Memory is allocated from the heap via `malloc(int size)` where `size` is the number of bytes needed. `malloc` returns the address of (a pointer to) a region of free memory of at least `size` bytes.
- `malloc` returns 0 (NULL) if there isn't a big enough free region to satisfy the request.

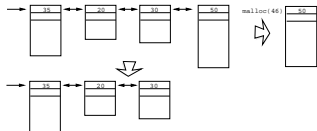
- malloc searches the free list for a free region that's big enough, removes it from the free list, and returns its address.



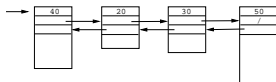
- If the region is bigger than requested, malloc takes what is needed from the end of the region and leaves the rest (the beginning) on the free list.



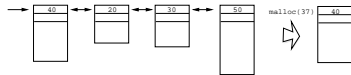
- If what is left is too small to hold a header malloc simply returns the entire region. The program that called malloc will get more memory than it asked for, but that doesn't matter.



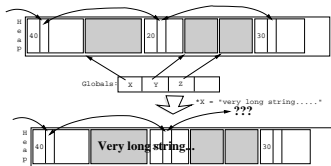
- A doubly-linked-list is often used to make insertion and deletion easier.



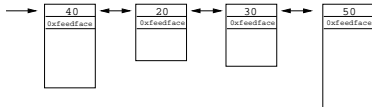
- malloc should return a word-aligned region that can store integers. The easiest way to do this is to start with the heap word-aligned, and always allocate memory in multiples of the word size.



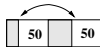
- What happens if the program asks for 50 bytes, but then writes 60 bytes to the region? The last 10 bytes overwrite the first 10 bytes of the next region. This will corrupt the free list if the next region is free (and probably crash the program if it is not).



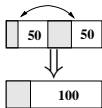
- For this reason it's a good idea to have a magic number in the free list node headers. This is a distinctive value that malloc checks when traversing the free list, and complains if the value changes (which indicates the list is corrupted). For example, put a field in the header whose value is always 0xfeedface.



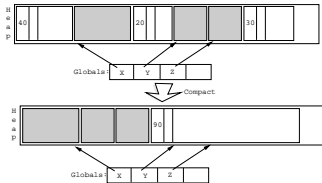
- There may be many regions big enough to satisfy a request. Which one should be used? First, why does it make a difference? Consider a heap that has two free regions, separated by allocated memory. Each region contains 50 bytes. malloc can't satisfy a request for 100 bytes, even though there is enough total free memory.



- Why not move the shaded region out of the way, by copying its contents to the beginning of the heap?

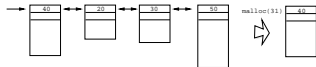


- This is called memory compaction, and doesn't work in general because it changes the address of memory previously returned by `malloc`. Most programs cannot handle this.

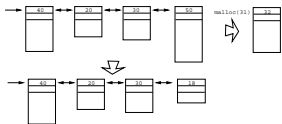


- Memory fragmentation is what occurs when the free space is in regions too small to be useful. We want our allocation scheme to avoid fragmentation, if possible.
- Note that there is no fragmentation with stack allocation because the free space is always in one big piece.
- Without compaction it is impossible to avoid fragmentation entirely.

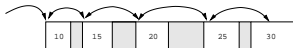
- `malloc` returns the smallest region that's big enough to satisfy the request.
- This favors small regions over large regions, leaving large regions for large requests.
- Unfortunately, it leaves lots of very tiny, useless regions because it always chooses the region whose size is closest to the request.



- `malloc` returns the biggest region that satisfies the request. This avoids the “lots of tiny regions” problem. The idea is that the amount “left over” is likely to be big enough to satisfy a subsequent request. This tends to produce free regions of about the same size.

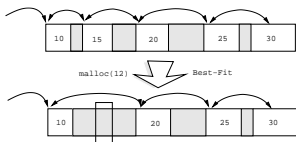


- Example: suppose the free list has regions of size 10, 15, 20, 25, and 30 bytes, and `malloc` is called to allocate a region of size 12.

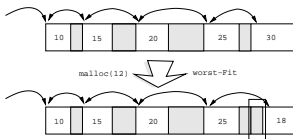


- Both of these algorithms require searching the entire free list, which can be expensive. They also cannot prevent fragmentation.

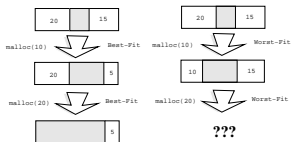
- Best-fit will allocate from the 15-byte region, leaving a region of 3 bytes. Probably will return all 15 bytes as the 3-byte region is too small to hold a node structure.



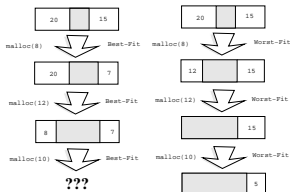
- Worst-fit will allocate from the 30-byte region, leaving a region of 18 bytes.



- Suppose the free list has two regions of size 20 and 15 and the allocations are 10 then 20: which algorithm wins?

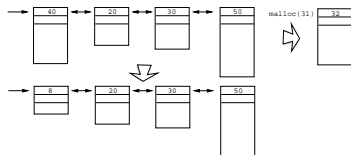


- Who wins if the allocations are 8, 12, then 10?



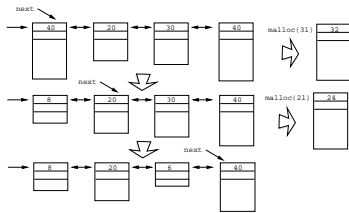
Memory Allocation Schemes—First-fit

- To avoid searching the entire list First-fit allocates space from the first region on the free list that is big enough. It has the side-effect of leaving lots of small regions at the front of the list, while the larger are at the end.

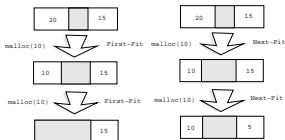


Memory Allocation Schemes—Next-fit

- Next-fit works like First-fit except that it starts searching at the region following the last region allocated. It tends to leave average-sized regions.



- Suppose the free list has two regions of size 20 and 15, and the allocations are 10 then 10. Who wins?



- The routine

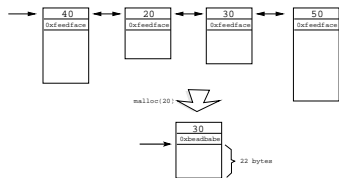
```
free(void *address)
```

is used to release memory when it is no longer needed (e.g. an employee quits or is fired).

- The address parameter is a pointer to the region to be freed, and it must have previously been returned by `malloc`.
- In C the type "void *" means a pointer to an unspecified type (it can be a pointer to anything, much like a reference to an object of type "Object" in Java).

- `free` inserts the region into the free list. It must fill in the header to do so, but this introduces a problem. How does `free` determine the size of the region? It is only given its address, not its size. One solution is to store the size elsewhere, such as in an array.
- A simple solution is to put a hidden header on allocated regions. This header contains the size of the region, and is stored in memory just before the allocated region.
- I call it "hidden" because it is not seen by the program that called `malloc`; `malloc` returns the address of the first byte after the hidden header.

- Free subtracts the size of the hidden header from the address it is given to get the address of the hidden header.
 - `malloc` must account for the hidden header by adding its size to the amount of memory requested.
 - It's a good idea to have a (different) magic number in the hidden header.
 - It's probably easiest to use the same header as the linked-list, although not the most space efficient.



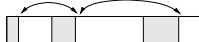
- There is a significant problem with this simple implementation of free. Consider the following memory status:



An employee quits, making an allocated region free:



Note that there are now three consecutive free regions (A,B,C), without any allocated regions between them. They should be coalesced into one large free region:



- Free needs to find the two regions on either side (A,C) of the region being freed (B), and if either is also free, merge it with the current region into a single region.
 - How does free find adjacent regions? One option is to keep the free list sorted by address. When a region is freed insert it into the correct place in the list, then check to see if the free regions before and after it on the list are adjacent.

- How are two regions coalesced? The one with the higher address is merged into the one with the lower address.
- Increment the size of the lower region by the size of the higher region (including the size of the higher region's header).
- Remove the higher region from the free list.

- Each free region of memory contains a header with the following fields:
 - Links for the doubly-linked free list
 - Size of region (including header)
- Each allocated region of memory contains a hidden header:
 - Size of region (including hidden header)

- Add size of hidden header to size.
- Round size up to a whole number of words.
- Use Next-fit to search free list for first free region of at least size bytes.
- Start search using pointer kept in global variable.
- Check magic numbers as you go.
- If there isn't a free region big enough, return 0 (NULL).
- If free region isn't exactly size bytes, compute how much is left-over (excess).

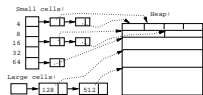
- If excess is at least the size of a header plus one word:
 - Address of allocated region = address of free region + excess
 - else
 - Remove free region from free list
 - Address of allocated region = address of free region
- Put size of allocated region into hidden header at the start of the region.
- Return address of first byte after hidden header.
- Update starting search pointer to next region on free list.

- Subtract size of hidden header from address to get hidden header's address.
- Check magic number in hidden header.
- Use size in hidden header to fill in free-list node header at start of region.
- Insert free region into sorted free list. The pointer to the start of the free list is stored in a global variable.

- ④ If previous region on free list is adjacent to new free region:
 - ① Coalesce by adding size of higher region (including header) to size of lower region.
 - ② Remove higher region from free list.
 - ③ Also coalesce subsequent region on free list if it is adjacent.

- There are many different memory allocation algorithms. Many optimize the amount of space used for the header.

Binning: keeps separate lists for regions of the same size. Good if the program allocates and frees lots of regions of the same size (e.g. employee structures).



Bit map: if regions are of the same size (blocks), keep a bit map instead of a free list. One bit per block: 1 means allocated, 0 means free. One word of the bit map represents 32 blocks. Index of bit in bit map is the index of the region in the "array" of regions.

Garbage collection: if a program does not have a pointer to a particular region of memory, that memory cannot be accessed by the program and can be automatically added to the free list. The programmer allocates memory, but never calls free. Examples: Lisp, Java

- [Read Scott, pp. 369–380.](#)