

15-213, Fall 2006
Malloc Lab: Writing a Dynamic Storage Allocator
Assigned: Fri, Oct. 27, Due: Fri, Nov. 10, 11:59PM
Last Possible Time to Turn in: Sun, Nov. 12, 11:59PM

Donnie H Kim (dhjkim@cs.cmu.edu) is the lead person for this assignment.

1 Introduction

In this lab you will be writing a dynamic storage allocator for C programs, that is, your own version of the `malloc`, `free`, `realloc`, and `calloc` functions. You are encouraged to explore the design space creatively and implement an allocator that is correct, efficient and fast.

2 Logistics

This is an individual project. You should do this lab on one of the class saltwater fish machines. As always, clarifications and corrections will be posted to the Autolab message board.

3 Hand Out Instructions

Start by downloading `mallocclab-handout.tar` from Autolab to a protected directory in which you plan to do your work. Then give the command `tar xvf mallocclab-handout.tar`. This will cause a number of files to be unpacked into the directory.

The only file you will be modifying and turning in is `mm.c`, which contains your solution. The `mdriver.c` program is a driver program that allows you to evaluate the performance of your solution. Use the command `make` to generate the driver code and run it with the command `./mdriver`.

4 How to Work on the Lab

Your dynamic storage allocator will consist of the following functions, which are declared in `mm.h` and defined in `mm.c`.

```
int    mm_init(void);
void *malloc(size_t size);
void free(void *ptr);
void *realloc(void *ptr, size_t size);
void *calloc (size_t nmemb, size_t size);
void mm_heapcheck(void);
```

The `mm.c` file we have given you implements nothing. However, we have also provided you with a program called `mm-naive.c`, which implements everything correctly, but naively. We have also provided you with a working version of the implicit list allocator described in your textbook, called `mm-implicit.c`.

You may use either of these examples as starting points for your own `mm.c` file. Implement the functions (and possibly define other private `static` functions), so that they obey the following semantics:

- `mm_init`: Performs any necessary initializations, such as allocating the initial heap area. The return value should be -1 if there was a problem in performing the initialization, 0 otherwise.

Every time the driver executes a new trace, it resets your heap to the empty heap by calling your `mm_init` function.

- `malloc`: The `malloc` routine returns a pointer to an allocated block payload of at least `size` bytes. The entire allocated block should lie within the heap region and should not overlap with any other allocated chunk.

Since the standard C library (`libc`) `malloc` always returns payload pointers that are aligned to 8 bytes, your `malloc` implementation should do likewise and always return 8-byte aligned pointers.

- `free`: The `free` routine frees the block pointed to by `ptr`. It returns nothing. This routine is only guaranteed to work when the passed pointer (`ptr`) was returned by an earlier call to `malloc`, `calloc`, or `realloc` and has not yet been freed. `free(NULL)` has no effect.
- `realloc`: The `realloc` routine returns a pointer to an allocated region of at least `size` bytes with the following constraints.

- if `ptr` is `NULL`, the call is equivalent to `malloc(size)`;
- if `size` is equal to zero, the call is equivalent to `free(ptr)`, and should return `NULL`;
- if `ptr` is not `NULL`, it must have been returned by an earlier call to `malloc` or `realloc`, and not yet have been freed. The call to `realloc` changes the size of the memory block pointed to by `ptr` (the *old block*) to `size` bytes and returns the address of the new block. Notice that the address of the new block might be the same as the old block, or it might be different, depending

on your implementation, the amount of internal fragmentation in the old block, and the size of the `realloc` request.

The contents of the new block are the same as those of the old `ptr` block, up to the minimum of the old and new sizes. Everything else is uninitialized. For example, if the old block is 8 bytes and the new block is 12 bytes, then the first 8 bytes of the new block are identical to the first 8 bytes of the old block and the last 4 bytes are uninitialized. Similarly, if the old block is 8 bytes and the new block is 4 bytes, then the contents of the new block are identical to the first 4 bytes of the old block.

- `calloc`: Allocates memory for an array of `nmemb` elements of `size` bytes each and returns a pointer to the allocated memory. The memory is set to zero before returning.

Note: Your `calloc` will not be graded on throughput or performance. Therefore a correct simple implementation will suffice.

- `mm_checkheap`: The `mm_checkheap` function scans the heap and checks it for consistency. This function will be very useful in debugging your `malloc` implementation. Some `malloc` bugs are very hard to debug using conventional `gdb` techniques. The only effective technique for some of these bugs is to use a heap consistency checker. When you encounter a bug, you can isolate it with repeated calls to the consistency checker until you find the instruction that corrupted your heap. Because of the importance of the consistency checker, it will be graded.

These semantics match the semantics of the corresponding `libc` routines. Type `man malloc` to the shell for complete documentation.

5 Support Routines

The `memlib.c` package simulates the memory system for your dynamic memory allocator. You can invoke the following functions in `memlib.c`:

- `void *mem_sbrk(int incr)`: Expands the heap by `incr` bytes, where `incr` is a positive non-zero integer and returns a generic pointer to the first byte of the newly allocated heap area. The semantics are identical to the Unix `sbrk` function, except that `mem_sbrk` accepts only a positive non-zero integer argument.
- `void *mem_heap_lo(void)`: Returns a generic pointer to the first byte in the heap.
- `void *mem_heap_hi(void)`: Returns a generic pointer to the last byte in the heap.
- `size_t mem_heapsize(void)`: Returns the current size of the heap in bytes.
- `size_t mem_pagesize(void)`: Returns the system's page size in bytes (4K on Linux systems).

6 The Trace-driven Driver Program

The driver program `mdriver.c` in the `malloclab-handout.tar` distribution tests your `mm.c` package for correctness, space utilization, and throughput. The driver program is controlled by a set of *trace files* that are included in the `malloclab-handout.tar` distribution. Each trace file contains a sequence of allocate and free directions that instruct the driver to call your `malloc` and `free` routines in some sequence. The driver and the trace files are the same ones we will use when we grade your `handin mm.c` file.

When the driver program is run, it will run each trace file 12 times: once to make sure your implementation is correct, once to determine the space utilization, and 10 times to determine the performance.

The driver `mdriver.c` accepts the following command line arguments. The normal operation is to run it with no arguments, but you may find it useful to use the arguments during development.

- `-t <tracedir>`: Look for the default trace files in directory `tracedir` instead of the default directory defined in `config.h`.
- `-f <tracefile>`: Use one particular `tracefile` instead of the default set of tracefiles for testing correctness and performance.
- `-c <tracefile>`: Run a particular `tracefile` exactly once, testing only for correctness. This option is extremely useful if you want to print out debugging messages.
- `-h`: Print a summary of the command line arguments.
- `-l`: Run and measure `libc malloc` in addition to the student's `malloc` package. This is interesting mainly to see how slow the `libc malloc` package is.
- `-V`: Verbose output. Prints additional diagnostic information as each trace file is processed. Useful during debugging for determining which trace file is causing your `malloc` package to fail.
- `-v <verbose level>`: This optional feature lets you set your verbose level manually to a particular integer.
- `-d <i>`: At debug level 0, very little validity checking is done. This is useful if you're mostly done but just tweaking performance.

At debug level 1, every array the driver allocates is filled with random bits. When the array is freed or reallocated, we check to make sure the bits haven't been changed. This is the default.

At debug level 2, every time any operation is done, all arrays are checked. This is very slow, but useful to discover problems very quickly.

- `-D`: Equivalent to `-d2`.
- `-s <s>`: Timeout after `s` seconds. The default is to never timeout.

7 Programming Rules

- You should not change any of the interfaces in `mm.h`. However, we strongly encourage you to use static functions in `mm.c` to break up your code into small, easy-to-understand segments.
- You should not invoke any external memory-management related library calls or system calls. This excludes the use of the `libc malloc, calloc, free, realloc, sbrk, brk` or any other memory management packages in your code.
- You are not allowed to define any global data structures such as arrays, structs, trees, or lists in your `mm.c` program. However, you *are* allowed to declare global scalar variables such as integers, floats, and pointers in `mm.c`.

The reason for this restriction is that the driver can't account for such global variables in its memory utilization measure. If you need space for large data structures, you can put them at the beginning of the heap.

- You are not allowed to simply hand in the code for the allocators from the CS:APP or K&R books. If you do so you will receive no credit.

However, we encourage you to study these codes and to use them as starting points. For example, you might modify the CS:APP code to use an explicit list with constant time coalescing. Or you might modify the K&R code to use constant time coalescing. Or you might use either codes as the basis for a segregated list allocator.

- It is OK to look at any descriptions of algorithms found in the textbook or elsewhere, but it is **not** acceptable to copy any code of malloc implementations found online or in other sources, except for the implicit list allocator described in your book.
- We encourage you to study the trace files and optimize for them, but your code must be correct on any trace. The score you get is averaged over all traces marked '*'. The utilization score weights all traces equally, whereas the performance score weights by the number of operations. In other words, if you are worried about speed, optimize for the largest traces.
- For consistency with the `libc malloc` package, which returns blocks aligned on 8-byte boundaries, your allocator must always return pointers that are aligned to 8-byte boundaries. The driver will check this requirement.

8 Evaluation

There are a total of 120 points. You will receive **zero points** if you break any of the rules or your code is buggy and crashes the driver. Otherwise, your grade will be calculated as follows:

- *Performance (100 points)*. Two metrics will be used to evaluate your solution:

- *Space utilization*: The peak ratio between the aggregate amount of memory used by the driver (i.e., allocated via `malloc` but not yet freed via `free`) and the size of the heap used by your allocator. The optimal ratio equals to 1. You should find good policies to minimize fragmentation in order to make this ratio as close as possible to the optimal.
- *Throughput*: The average number of operations completed per second.

The driver program summarizes the performance of your allocator by computing a *performance index*, $0 \leq P \leq 100$, which is a weighted sum of the space utilization and throughput

$$P = 100 * \left(w \min \left(1, \frac{U}{U_{thresh}} \right) + (1 - w) \min \left(1, \frac{T}{T_{thresh}} \right) \right)$$

where U is your space utilization, T is your throughput, and U_{thresh} and T_{thresh} are the estimated space utilization and throughput of an optimized malloc package.¹ The performance index favors space utilization over throughput: $w = 0.6$.

Observing that both memory and CPU cycles are expensive system resources, we adopt this formula to encourage balanced optimization of both memory utilization and throughput. Since each metric will contribute at most w and $1 - w$ to the performance index, respectively, you should not go to extremes to optimize either the memory utilization or the throughput only. To receive a good score, you must achieve a balance between utilization and throughput.

The 100 performance points (`$perfpoints`) will be allocated as a function of the performance index (`$perfindex`):

```
if ($perfindex < 60) {
    $perfpoints = 0;
}
elseif ($perfindex < 98) {
    $perfpoints = (40 + ((3 * $perfindex)/5));
}
else {
    $perfpoints = 100;
}
```

We chose this function so that, when run on the saltwater fish machines, the CS:APP implicit list allocator receives 0/100 points, a good explicit list allocator receives around 85/100 points, a good segregated list allocator gets around 96/100 points, and a highly tuned seglist allocator can get 100/100 points.

Partial credit. Partial credit will be available for those students who are unable to finish a working allocator. If you can't get your allocator to work by the due date, document and handin what you have and your instructors will award partial credit manually.

¹The values for U_{thresh} and T_{thresh} are constants in the driver (0.93 and 15,000 Kops/s) that your instructor established when they configured the program. This means that once you beat 93% utilization and 15,000 Kops/s, your performance index is perfect.

- *Heap Consistency Checker (10 points)*. Ten points will be awarded based on the quality of your implementation of `mm_checkheap`. It is up to your discretion how thorough you want your heap checker to be. The more the checker tests, the more valuable it will be as a debugging tool.

However, to receive full credit for this part, we require that the header comments for your heap checker list **all** of the invariants of your data structure. For each such invariant, you should state whether or not your heap checker verifies that it is satisfied. (It is not OK to list all the invariants and not check any of them - you should at least verify the critical portions). Some examples of what your heap checker should check are provided below.

- Checking the heap (implicit list, explicit list, segregated list):
 - * Check epilogue and prologue blocks.
 - * Check block's address alignment.
 - * Check heap boundaries.
 - * Check each block's header and footer: size (minimum size, alignment), prev/next allocate/free bit consistency, header and footer matching each other.
 - * Check coalescing: no two consecutive free blocks in the heap.
- Checking the free list (explicit list, segregated list):
 - * All next/prev pointers are consistent (if A's next pointer points to B, B's prev pointer should point to A).
 - * All free list pointers points between `mem_heap_lo()` and `mem_heap_high()`.
 - * Count free blocks by iterating every block, and traversing free list by pointers and see if they match.
 - * All blocks in each list bucket fall within bucket size range (segregated list).
- *Style (10 points)*.
 - Your code should be decomposed into functions and use as few global variables as possible. You should use macros or inline functions to isolate the pointer arithmetic to as few places as possible.
 - Your code must begin with a header comment that gives an overview of the structure of your free and allocated blocks, the organization of the free list, and how your allocator manipulates the free list.
 - In addition to this overview header comment, each function should be preceded by a header comment that describes what the function does.

9 Handin Instructions

Make sure you have included your name and Andrew ID in the header comment of `mm.c`.

Hand in your `mm.c` file by uploading it to Autolab. You may submit your solution as many times as you wish up until the due date.

Only the last version you submit will be graded.

For this lab, you must upload your code for the results to appear on the class status page.

10 Hints

- *Use the `mdriver -c` option or `-f` option.* During initial development, using tiny trace files will simplify debugging and testing. The first several traces that `mdriver` runs are such small trace files.
- *Use the `mdriver -V` options.* The `-V` option will also indicate when each trace file is processed, which will help you isolate errors.
- *Use the `mdriver -D` option.* This does a lot of checking to quickly find errors.
- *Compile with `gcc -g` and use a debugger.* A debugger will help you isolate and identify out of bounds memory references. You may want to modify the Makefile and remove the `-O2` option during initial testing.
- *Use `gdb`'s `watch` command* to find out what changed some value you didn't expect to have changed.
- *Understand every line of the implicit list malloc implementation in the textbook.* A working version of this allocator is included in your handout directory in `mm-implicit.c`.
- *Encapsulate your pointer arithmetic in C preprocessor macros or inline functions.* Pointer arithmetic in memory managers is confusing and error-prone because of all the casting that is necessary. You can reduce the complexity significantly by writing macros for your pointer operations. See the text for examples.
- *Remember we are working with 64-bit fish machines.* Pointers take up 8 bytes of space, so you should understand the macros in the book and port them to 64-bit machines. Notably, `sizeof(size_t) == 8` on 64-bit machines.
- *Use your heap consistency checker.* We are assigning ten points to your `mm_heapcheck` function for a reason. A good heap consistency checker will save you hours and hours when debugging your malloc package. You can use your heap checker to find out where exactly things are going wrong in your implementation (hopefully not in too many places!). Make sure that your heap checker is detailed. Your heap checker should scan the heap, performing sanity checks and possibly printing out useful debugging information. Every time you change your implementation, one of the first things you should do is think about how your `mm_heapcheck` will change, what sort of tests need to be performed, etc.
- *Use a profiler.* You may find the `gprof` tool helpful for optimizing performance.
- *Versioning your implementation.* You may find it useful to manage a couple of different versions of implementation (i.e. explicit list, segregated list) during the assignment. Since `mdriver` looks for `mm.c`, creating a symbolic link between files is useful in this case. For example, you can create a symbolic link between `mm.c` and your implementation such as `mm-explicit.c` with command

line `ln -s mm-explicit mm.c`. An example of linking `mm-implicit.c` to `mm.c` is provide in the handout `Makefile` (`make naive` will create a link between `mm.c` and `mm-implicit.c`). You can easily extend it to suit your implementation.

- *Start early!* It is possible to write an efficient malloc package with a few pages of code. However, we can guarantee that it will be some of the most difficult and sophisticated code you have written so far in your career. So start early, and good luck!

11 More Hints

Basically, you want to design an algorithm and data structure for managing free blocks that achieves the right balance of space utilization and speed. Note that this involves a tradeoff. For space, you want to keep our internal data structures small. Also, while allocating a free block, you want to do a thorough (and hence slow) scan of the free blocks, to extract a block that best fits our needs. For speed, you want fast (and hence complicated) data structures that consume more space. Here are some of the design options available to you:

- Data structures to organize free blocks:
 - Implicit free list
 - Explicit free list
 - Segregated free lists
- Algorithms to scan free blocks:
 - First fit/Next fit
 - Blocks sorted by address with first fit
 - Best fit

You can pick (almost) any combination from the two. For example, you can implement an explicit free list with next fit, a segregated list with best fit, and so on. Also, you can build on a working implementation of a simple data structure to a more complicated one.

In general, we suggest that you start with an implicit free list, then change this to an explicit list, and then use the explicit list as the basis for a final version based on segregated lists.