

CSE 351, Autumn 2010

Lab 7: Writing a Dynamic Storage Allocator

Due: Thursday Dec. 9, 11 PM

1 Instructions

In this lab you will be writing a dynamic storage allocator for C programs, i.e., your own version of the `malloc` and `free` routines.

Start by extracting `/projects/instr/10au/cse351/labs/lab7.tar` to a (protected) directory in which you plan to do your work:

```
tar xvf /projects/instr/10au/cse351/labs/lab7.tar
```

(You can also obtain the file from the direct link on the web page.)

This will cause a number of files to be unpacked into the directory. The only file you will be modifying and handing in is `mm.c`. 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 -v`. (The `-v` flag displays helpful summary information.)

1.1 Submitting Your Work

When you have completed the lab, you will hand in only one file (`mm.c`), which contains your solution. To turn in `mm.c`, run `make submit` and turn in the resulting archive file using the regular Catalyst dropbox, or just turn in the single `mm.c` file. For this lab, you can work anywhere there's a C compiler and `make`, but make sure your allocator works on `attu`, where we'll be testing it.

2 How to Work on the Lab

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

```
int    mm_init(void);
void*  mm_malloc(size_t size);
void   mm_free(void* ptr);
```

The `mm.c` file we have given you partially implements an allocator using an explicit free list. Your job is to complete this implementation by filling out `mm_malloc` and `mm_free`. The three main memory management functions should work as follows:

- `mm_init` (provided): Before calling `mm_malloc` or `mm_free`, the application program (i.e., the trace-driven driver program that you will use to evaluate your implementation) calls `mm_init` to perform any necessary initializations, such as allocating the initial heap area. The return value is `-1` if there was a problem in performing the initialization, `0` otherwise.
- `mm_malloc`: The `mm_malloc` routine returns a pointer to an allocated block payload of at least `size` bytes. (`size_t` is a type for describing sizes; it's basically an unsigned integer.) The entire allocated block should lie within the heap region and should not overlap with any other allocated chunk.

We will compare your implementation to the version of `malloc` supplied in the standard C library (`libc`). Since the `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.

- `mm_free`: The `mm_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 `mm_malloc` and has not yet been freed.

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

3 Supporting Code

We define a `BlockInfo` struct designed to be used as a node in a doubly-linked explicit free list, and the following functions for manipulating free lists.

- `BlockInfo* searchFreeList(int reqSize)` returns a block of at least the requested size if one exists (and `NULL` otherwise).
- `void insertFreeBlock(BlockInfo* blockInfo)` inserts the given block in the free list in LIFO manner.
- `void removeFreeBlock(BlockInfo* blockInfo)` removes the given block from the free list.
- Helper functions for implementing list functions:
 - `BlockInfo* getFreeListHead()` returns a pointer to the first block in the free list.
 - `void getFreeListHead(BlockInfo* newHead)` takes a pointer to a block and sets the head of the free list to point to the block referenced by the pointer.

In addition, we implement `mm_init` and provide two helper functions implementing important parts of the allocator:

- `void requestMoreSpace(int incr)` enlarges the heap by `incr` bytes (if enough memory is available on the machine to do so).
- `void coalesceFreeBlock(BlockInfo* oldBlock)` coalesces any other free blocks adjacent in memory to `oldBlock` into a single new large block and updates the free list accordingly.

Finally, we use a number of C Preprocessor macros to extract common pieces of code (constants, annoying casts/pointer manipulation) that might be prone to error. Each is documented in the code. You are welcome to use macros as well, though the ones already included in `mm.c` are the only ones we used in our sample solution, so it's possible without more. For more info on macros, check the GCC manual at <http://gcc.gnu.org/onlinedocs/cpp/Macros.html>.

3.1 Memory System

The `memlib.c` package simulates the memory system for your dynamic memory allocator. In your allocator, you can call the following functions (if you use the provided code for an explicit free list, most uses of the memory system calls are already covered).

- `void* mem_sbrk(int incr)`: Expands the heap by `incr` bytes, where `incr` is a positive non-zero integer and returns a 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. (Run `man sbrk` if you want to learn more about what this does in Unix.)
- `void* mem_heap_lo()`: Returns a pointer to the first byte in the heap.
- `void* mem_heap_hi()`: Returns a pointer to the last byte in the heap.
- `size_t mem_heapsize()`: Returns the current size of the heap in bytes.
- `size_t mem_pagesize()`: Returns the system's page size in bytes (4K on Linux systems).

4 The Trace-driven Driver Program

The driver program `mdriver.c` in the `lab7.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 `lab7.tar` distribution. Each trace file contains a sequence of `allocate` and `free` directions that instruct the driver to call your `mm_malloc` and `mm_free` routines in some sequence. The driver and the trace files are the same ones we will use when we grade your submitted `mm.c` file.

The driver `mdriver` accepts the following command line arguments:

- `-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` for testing instead of the default set of tracefiles.
- `-h`: Print a summary of the command line arguments.
- `-l`: Run and measure `libc` `malloc` in addition to the student's `malloc` package.
- `-v`: Verbose output. Print a performance breakdown for each tracefile in a compact table.
- `-V`: More 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.

5 Programming Rules

- You should not change any of the interfaces in `mm.c`.
- You should not invoke any memory-management related library calls or system calls. This excludes the use of `malloc`, `calloc`, `free`, `realloc`, `sbrk`, `brk` or any variants of these calls in your code. (You may use all the functions in `memlib.c`, of course.)
- You are not allowed to define any global or `static` compound data structures such as arrays, structs, trees, or lists in your `mm.c` program. You *are* allowed to declare global scalar variables such as integers, floats, and pointers in `mm.c`, but try to keep these to a minimum. (It is possible to complete the implementation of the explicit free list without adding any global variables.)
- For consistency with the `malloc` implementation in `libc`, which returns blocks aligned on 8-byte boundaries, your allocator must always return pointers that are aligned to 8-byte boundaries. The driver will enforce this requirement for you.

6 Evaluation

Your grade will be calculated out of a total of 60 points as follows:

- *Correctness (45 points)*. You will receive 5 points for each test performed by the driver program that your solution passes. (9 tests)
- *Style (10 points)*.
 - Your code should use as few global variables as possible (ideally none!).
 - Your code should be as clear and concise as possible.
 - Since some of the unstructured pointer manipulation inherent to allocators can be confusing, short inline comments on steps of the allocation algorithms are also recommended. (These will also help us give you partial credit if you have a partially working implementation.)
 - Each function should have a header comment that describes what it does and how it does it.

- *Performance (5 points)*. Performance represents a small portion of your grade. We are most concerned about the correctness of your implementation. For the most part a correct implementation will yield reasonable performance. Two performance 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 `mm_malloc` but not yet freed via `mm_free`) and the size of the heap used by your allocator. The optimal ratio is 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*, P , which is a weighted sum of the space utilization and throughput

$$P = 0.6U + 0.4 \min\left(1, \frac{T}{T_{libc}}\right)$$

where U is your space utilization, T is your throughput, and T_{libc} is the estimated throughput of `libc malloc` on your system on the default traces.¹ The performance index favors space utilization over throughput. You will receive $5(P + 0.1)$ points, rounded *up* to the closest whole point. For example, a solution with a performance index of 0.63 or 63% will receive 4 performance points. Our complete version of the explicit free list allocator has a performance index between 0.7 and 0.8; it would receive 5 points.

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. Ideally, the performance index will reach $P = 1$ or 100%. To receive a good performance score, you must achieve a balance between utilization and throughput.

7 Hints

- *Use the `mdriver -f` option*. During initial development, using tiny trace files will simplify debugging and testing. We have included two such trace files (`short1-bal.rep` and `short2-bal.rep`) that you can use for initial debugging.
- *Use the `mdriver -v` and `-V` options*. The `-v` option will give you a detailed summary for each trace file. The `-V` will also indicate when each trace file is read, which will help you isolate errors.
- *Compile with `gcc -g` and use `gdb`*. The `-g` flag tells `gcc` to include debugging symbols, so `gdb` can follow the source code as it steps through the executable. To use the `-g` flag with the `Makefile`, edit the `CFLAGS` variable in the `Makefile` OR just run `make` like this:

```
make CFLAGS="-Wall -g"
```

¹The value for T_{libc} is a constant in the driver (1800 Kops/s) that is close to the average throughput of the `libc` allocator on the same traces, measured on `attu`. The performance index will vary from system to system based on the local `libc` allocator throughput, so run on `attu` for a good idea of where you stand.

This has the effect of replacing the value of `CFLAGS` as defined in the `Makefile` with whatever you type on the command line instead. (Just be sure to quote it if it contains spaces, and do not put spaces around the `=` sign.) In general, when debugging, you want to turn off the `-O2` (that's a capital *o*, not a zero), since it tells the compiler to perform optimizations that can occasionally make following in the debugger confusing. A debugger will help you isolate and identify out of bounds memory references.

- *Understand every line of the malloc implementation in the textbook.* The textbook has a detailed example of a simple allocator based on an implicit free list. Use this as a point of departure. Don't start working on your allocator until you understand everything about the simple implicit list allocator.
- *Encapsulate your pointer arithmetic in C preprocessor macros.* 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 textbook for examples.
- *Use a profiler.* You may find the `gprof` tool helpful for optimizing performance. (`man gprof` or searching online for `gprof` documentation will get you the basics.) If you use `gprof`, see the hint about debugging above for how to pass extra arguments to GCC in the `Makefile`.
- *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!

7.1 Heap Consistency Checker

Dynamic memory allocators are notoriously tricky beasts to program correctly and efficiently. They are difficult to program correctly because they involve a lot of untyped pointer manipulation. In addition to the usual debugging techniques, you may find it helpful to write a heap checker that scans the heap and checks it for consistency.

Some examples of what a heap checker might check are:

- Is every block in the free list marked as free?
- Are there any contiguous free blocks that somehow escaped coalescing?
- Is every free block actually in the free list?
- Do the pointers in the free list point to valid free blocks?
- Do any allocated blocks overlap?
- Do the pointers in a heap block point to valid heap addresses?

Your heap checker will consist of the function `int mm_check(void)` in `mm.c`. Feel free to rename it, break it into several functions, and call it wherever you want. It should check any invariants or consistency conditions you consider prudent. It returns a nonzero value if and only if your heap is consistent. This is not required, but may prove useful. When you submit `mm.c`, make sure to remove any calls to `mm_check` as they will slow down your throughput.