

Project 2: Semispace Copying Garbage Collector

Goal

Implement a semispace copying garbage collector.

Time

This project is assigned as of 9AM Wednesday, October 7th and is due by 12PM (**noon**, not midnight), Monday, October 26th.

Requirements

- The GC must be compatible with the public GGGGC API, accessible through `ggggc/gc.h`, plus the function `void ggggc_collect(void)`. Simply using the GGGGC template provided is sufficient so long as you don't reduce the API.
- The GC must include a `Makefile`, and `make` with no arguments should build `libggggc.a`. `libggggc.a` should have no further library requirements.
- The GC must compile on `linux.student.cs.uwaterloo.ca` in 64-bit mode.
- The GC must include a file named `project2.txt` which briefly details the techniques used in the implementation of the GC. As there are few choices, I expect this file to be *very* brief.
- A single-file test program must be included named `project1-test.c` if implemented in C, or `project1-test.cc` if implemented in C++, expected to be the same test program as in project 1 (hence the name). It may be different if you prefer, but must follow the same requirements as project 1's test program.
- The GC must be implemented in C or C++. If implemented in C++, the `Makefile` must be updated to call `g++` or `c++` as appropriate, and its API must be in C. The template provided will expose a C API even if compiled as C++.
- The GC must use a semispace copying approach, with a `fromspace` and `tospace` of equal size, and collection involving swapping spaces then copying objects from the `fromspace` to the `tospace`.
- After an explicit call to `ggggc_collect`, all references must be updated correctly, all living objects copied only once, and no dead objects copied at all.
- `ggggc_collect` must eliminate all external fragmentation.
- Objects returned by `ggggc_malloc` must be correctly allocated, include a pointer to their `struct GGGGC_Descriptor` in their object header (i.e., `obj->header.descriptor_ptr`), and otherwise be zeroed.
- Objects returned by `ggggc_malloc` must be unique amongst reachable objects. i.e., no two calls to `ggggc_malloc` without an intervening collection should return overlapping space.
- Objects must be allocated using a simple bump-pointer strategy and no free-list.

Options

- The included test program, `project1-test.c`, may perform any task you please, so long as it allocates enough memory to require GC.
- When `ggggc_collect` is not explicitly called, you may use any heuristic you deem fit decide when to collect, so long as it is transparent to the user and does not substantially burden performance.
- Space required for collection itself may be allocated from the C heap to ease collection.

Marks

Your submission will be graded partially on “black-box” tests, which test its correctness with no examination of the code, and partially on “white-box” tests, which involve direct inspection:

- 33.34%: All tests included in the GGGGC template, plus your test, plus private grading tests, compile correctly, give correct output and run without consuming excessive amounts of memory. If the submitted GC does not compile, none of these points are available. Although no performance requirements are set in this project, they are graded by a human (me), so execution times which test human patience will be given no points.
 - 11.12%: All tests compile against the submitted GC.
 - 11.11%: All included and grading tests give correct output.
 - 11.11%: Memory consumption on tests is not excessive¹.
- 15%: Allocation is implemented correctly.
 - 5%: Allocation from free space is implemented correctly, in the tospace of the previous collection, without external fragmentation.
 - 10%: Allocation of new pools is implemented correctly, and fromspace and tospace are always of equal size.
- 51.66%: Collection is implemented correctly.
 - 10%: All roots are correctly identified.
 - 20%: All reachable objects are copied to tospace.
 - 10%: Objects are never copied twice.
 - 21.66%: All references are correctly updated.

Notes

- GGGGC is released under the ISC license. You are free to look at the original source, but since it implements a very different GC, it’s unlikely to help much. Your changes are your own.
- Typically, object headers start *before* the object, i.e. at `&((struct ObjectHeader *) ptr)[-1]`. In GGGGC, object headers start *with* the object, i.e. at `((struct GGGC.Header *) ptr)`. The header is contained in the size of the object. Expanding the header will not break this invariant. If you expand the header, you will have to change `GGGGC_DESCRIPTOR_DESCRIPTION`, which is the pointer bitmap for `struct GGGC.Descriptor`.
- Descriptors in GGGGC contain the relevant type information in terms of a pointer bitmap. The pointer bitmap is in a partially-reversed order: The *low* bit (last bit) of the *first* word of the pointer map corresponds to the *first* word of the object. Because GGGGC objects include the header, this is the first word of the header. The *high* bit (first bit) of the first word of the pointer map corresponds to the 64th word of the object on a 64-bit system. The *low* bit (last bit) of the second word of the pointer map corresponds to the 65th word of the object on a 64-bit system.
- Since descriptors are part of the heap, GGGGC creates “descriptor descriptors”. This is a complicated problem and you shouldn’t have to study the code that does it in detail, but there is one important point: It sometimes allocates descriptor-descriptors by calling `gggc_malloc` with a temporary descriptor that is not actually on the heap. If your GC runs at the wrong time, it could find these objects and be confused. Your GC may easily distinguish such temporary descriptors by the fact that their `descriptor_ptr` field is set to `NULL`. Simply ignore any object with a `NULL` descriptor.

¹“Excessive” is not precisely defined. Generally speaking, if your collector is incorrect, it will leak memory, and thus any minor memory problem will be amplified over time, sometimes exponentially. Such bugs often crash the program or even cause system instability. Such problems are considered excessive. For these tests, minor memory bugs which do not present such visible errors will not be considered.

- The first word of the header should always be the object descriptor reference. As a result, the first word of every object is always a reference. Taking advantage of this fact, the low bit of the pointer bitmap is used for another purpose: If the object has no references aside from the descriptor, the low bit will be set to 0. A collector can thereby avoid scanning the whole object when it contains no references.
- Typical user code never calls `gggc_collect` directly. However, test code for this project will.
- Project 3 will reuse the code created for this project as well as the code created for project 1.
- There are very few software components more difficult to debug than garbage collectors. Valgrind will not help you. `gdb` will catch bugs long after the root cause. Print statements are crucial but not sufficient. Your best bet for debugging is what are called “canaries”: Fields in objects, in object headers or between objects in the heap which have fixed values, typically `0xDEADBEEF`. By inserting canaries and then checking them frequently, it is possible to determine what code corrupted the heap.
- Do not be afraid of rewriting. Rewriting a portion of this project, or even the entire project, may take considerably fewer hours than tracking down an obscure bug.

Submission

Project code must be submitted through the `submit` command available on `linux.student.cs.uwaterloo.ca` or its web interface. The course is `cs842` and the project is `gcproject2`. e.g.: `submit cs842 gcproject2 ./mygcprojectcode`. See `man submit` for further instructions on the `submit` command.

Late submissions are not accepted without prior agreement from the instructor. Extensions will rarely be granted. Requests for extensions sent within the 24 hours prior to the due date, or sent after the due date, will be ignored, except where existing university regulations on extensions, e.g. due to illness, are relevant.

Exceptions and help

Due to the size of this course, uniform projects are required. However, if you need to deviate from this document for exceptional reasons, or if you want any clarification or help, feel free to email the instructor. My email is available at my web site, <http://the.gregor.institute/>.