

# 15-411 Compiler Design, Spring 2024

## Lab 6 - Garbage Collection

Jan and co.

Due 11:59pm, Saturday, May 4th, 2024

### 1 Introduction

The main goal of Lab 6 is to explore advanced aspects of compilation, and for the Garbage Collection option, that means modifying your compiler and its associated runtime so that unneeded blocks of memory on the heap can be garbage collected. The source language you will compile is still L4, which remains unchanged from Labs 4 and 5.

### 2 Deliverables

#### 2.1 Garbage Collecting Compiler and Runtime

You have complete freedom with regard to which kind of garbage collector to implement. A garbage collector will consist of the compiler proper and the runtime system. The interface from the compiled code to the runtime system should be part of your design. Reasonable choices are a mark-and-sweep or a copying collector, but even a conservative collector is acceptable. Incremental or generational collectors are significantly harder and should only be attempted if you already have a basic collector working.

Grading criteria includes:

1. Functional correctness is paramount. You should not mutate the heap in such a way as to result in the incorrect execution of programs. Your compiler should continue to function correctly, despite any changes to the binary interface you use to interface with the garbage collector.
2. Developing a convincing framework for *understanding* and *quantifying* the performance of your garbage collector is also a priority.
3. The absence of memory leaks comes second. A garbage collector that takes a whole lot of processor time is not of much use if it cannot effectively collect parts of the heap that are no longer referenced. However, conservative collectors may not be able to reclaim all memory.
4. Performance is a distant third – actual performance has a very minor effect on your grade. Optimizations to garbage collectors require a significant amount of time. Therefore, we recommend that you avoid premature optimizations.

Issuing the shell command

```
% make
```

should generate the appropriate files so that

```
% bin/c0c --exe <args>
```

will run your compiler in safe mode with support for garbage collection. You may stop supporting all other compiler flags except `-t`. This means you no longer need to support the `--unsafe` flag or any optimization flags for this assignment.

After running `make`, issuing the shell command

```
% make test
```

should run all your own tests and print out informative output. The command

```
% make clean
```

should remove all binaries, heaps, and other generated files.

You should also update the `README` file and insert a roadmap to your code. This will be a helpful guide for the grader. In particular, since there are likely to be many different projects undertaken, do introduce the project at the very top of the `README`.

## Runtime environment

You will need to extend the `run411.c` runtime file and/or write new runtime files to implement garbage collection. Because we don't know exactly how your runtime works, your compiler should have an additional flag `--exe`. If compiler is given a well-formed input file `foo.11` or `foo.12` as a command-line argument and is also given the `--exe` argument, it should generate a target file called `foo.11.s` or `foo.12.s` (respectively) in the same directory as the source file, and should *also* compile the garbage collector and runtime and link your generated assembly to create an executable `foo.11.exe` or `foo.12.exe` (respectively).

Your `README` should explain how your compiler implements the `--exe` option, and should also explain how to link other runtimes and how to use any other command-line options used by the compiler or by the generated programs. For example, you may want a command-line argument accepted by either the compiler or the generated executable that determines how much heap space the executable is allowed to use.

## 2.2 Tests and Measurement Tools

At minimum, your testing *must* be sufficient to provide strong evidence of two things:

1. Your garbage collector does not corrupt the heap, and does not leak memory.
2. Your garbage collector allows programs to run that would otherwise have failed due to lack of resources.

To this end, feel free to search through all of the test cases that we have accumulated through this semester for programs that are both realistic and usefully contrived to assemble a test suite. You will need to write contrived test cases designed specifically to ensure that your garbage collector goes through several collection cycles without corrupting the heap or leaking any memory. You will also need to run your compiled programs in an environment where memory is artificially constrained. You will be graded on how well you test your garbage collector.

Please submit test cases and any associated testing framework in a clearly marked directory, such as `lab6-tests/`. If they are to be run in a different way than a vanilla L4 test, you should include a `README` file explaining exactly how to use your tests, and `make test` should run your tests.

## 2.3 Term Paper

You need to describe and critically evaluate your garbage collector in a term paper of about 5-10 pages. You may use more space if you need it. Your paper should follow this outline. Your paper should follow this outline.

1. Introduction: Provide an overview of your implementation and briefly summarize the results you obtained.
2. Compilation: Describe the data structures, code, and information generated by the compiler in order to support the garbage collector. If applicable, describe the syntax and semantics of new language features.
3. Runtime System: Describe the runtime system of the garbage collector, giving details of the algorithms and also its implementation (most likely in C).
4. Testing: Describe the criteria based on which you selected and designed your tests, and explain how you use them to verify the functionality of your garbage collector.
5. Analysis: Critically evaluate your collector and sketch future improvements one might make to its basic design.

Submit your term paper in PDF form via Gradescope before the stated deadline. Early submissions are much appreciated since it lessens the grading load of the course staff near the end of the semester.

**You may not use any late days on the term paper describing your implementation of Lab 6!**

## 3 Something Extra

Beyond the basic implementation and perfunctory analysis described above, an excellent and high-scoring final project will include a little something extra. Exactly what your “something extra” is is up to you, but it should represent a different direction, rather than (just) being an incremental or generational version of the collector you implemented.

One category of “something extra” ideas involves conducting a significantly more quantitative analysis of your garbage collector’s performance. These are couple of (non-exhaustive) suggestions:

- Compare the performance of multiple garbage collection strategies. Note that this requires implementing multiple garbage collection strategies, which is rather ambitious!
- Compare the performance of your compiler to the Boehm-Demers-Weiser conservative collector used by the reference compiler, which can be used as a drop-in replacement for `malloc()`. You can do this either by compiling with the reference compiler and runtime or by obtaining the Boehm-Demers-Weiser collector (see <http://www.hboehm.info/gc/>).
- Extend L4 with manual memory management using a built-in `free()` function that accepts any pointer or array type. Compare the performance of the manual and garbage-collected memory management.
- Analyze the way in which performance is influenced by varying the total amount of memory available to your program. (You can measure performance both in terms of total running time and in terms of total number of collections.)

Another way to do something extra is to explore a language feature that is enabled by garbage collection:

- Implement *weak pointers*, references that do not keep the allocations they point to from being reclaimed. In order to preserve memory safety, it must always be possible to determine if the target of a weak pointer has been reclaimed. Demonstrate an interesting program using weak pointers. (See section 7.1 of the Wilson review.)
- Implement *finalizers*, functions that are associated with a specific pointer or a specific type of pointer and that run when the pointer is freed by garbage collection. Demonstrate an interesting program using finalizers. (See section 7.2 of the Wilson review.)
- Implement a leak detector: use your garbage collection infrastructure to explore a safe implementation of `free()` that always safely detects double-frees (signaling a memory error) and reports the number of un-freed allocations when the program exits.

## 4 Notes and Hints

- Limit optimizations. Garbage collection is easier if fewer optimizations are applied to the code, especially where memory references are concerned. In order to concentrate on the garbage collector, it is probably a good idea to stay away from optimizations altogether.
- Apply regression testing. It is very easy to get caught up in the new functionality.
- Copying vs. mark-and-sweep collector. Experience in previous years indicates that a copying collector is easier to implement for our language than a mark-and-sweep collector because the data structures are simpler.