

## Important Dates:

**Design Proposal: Thursday, November 30**

**Final Compiler and Documentation: Monday, December 12**

**Derby: Wednesday, December 13**

In this final phase of the project your goal is to reduce the execution time of your generated code while preserving the semantics of the given application. This phase is completely opened-ended; you are not required to implement any specific transformations. Your grade will be based on your overall design (with heavy emphasis on the writeup) and how well your compiler performs as compared to the submissions of the other groups. During the final class, we will hold the Compiler Derby where your compiler implementation will compete against the implementations of your fellow classmates on a hidden program (the derby program).

In the final month of class, we will cover various code improving optimizations. The AMD64 architecture is very complex and aggressive. A substantial portion of this project phase is to determine which optimizations will provide a benefit given the programs of the provided benchmark suite and the target architecture. Some optimizations that we will cover in class include:

- **Register Allocation** – Your compiler can implement a graph coloring based register allocator. See Chapter 16 of the Whale book and §9.7 of the Dragon book. Your register allocator should take advantage of the full set of general-purpose registers provided by the x86-64 ISA. It should also respect the Linux calling convention including caller-save and callee-save properties of each register.
- **List Scheduling** – Instruction scheduling minimizes pipeline stalls for long latency operations such as loads, multiplies, and divides. See Chapter 17 of Whale book.
- **Instruction Selection** – So far, we have been using a restricted subset of the x86-64 ISA. As a peephole optimization, you might replace a sequence of simple instructions with a single, complex instruction that performs the same operation in fewer cycles.
- **Data Parallelization** – So far, we have treated the `forpar` construct in the same manner as a normal `for` loop. A `forpar` loop reveals that the iterations of the loop can be executed concurrently. Using a provided parallelization library (to be presented in a future document), you can schedule different iterations of a `forpar` loop to execute on different cores of our target architecture.

In order to identify and prioritize optimizations, we have provided you with a benchmark suite of three image-processing programs. These programs are more complex than the code that has

been provided during the previous phases, so your first priority is to produce correct unoptimized assembly code for each benchmark.

After your compiler produces correct unoptimized code, you should begin to study the applications of the benchmark suite. You will use these programs (and any other programs provided during previous phases) to determine which optimizations you will implement and in what order they will be applied. You are expected to analyze the assembly code produced by your compiler to classify the effectiveness of each proposed optimization, perhaps first applying the optimization manually and empirically measuring the benefit. Your writeup should include evidence for the effectiveness of each optimization you considered.

You are not limited to the optimizations covered in class. You are free to implement any optimization that you come across in your reading or any optimization that you devise. However, your writeup must demonstrate that each optimization is effective and semantics preserving. Furthermore, you must argue that each optimization is general, meaning it is not a special-case optimization that works only for the derby program or a specific application in the benchmark suite.

You should consult AMD's documentation for details regarding our target architecture and the ISA. The documentation is linked from the class website. To benchmark your generated assembly code, provide the `-pg` option to `gcc` or `cc` while assembling in order to generate profiling information. After the code is executed, you can use `gprof` to examine the generated profile statistics. Also, you can use the Unix command `time` focusing on the "user" time. A more accurate timing mechanism may be provided in the future.

## Writeup

The writeup for this project is extremely important. Although it explicitly accounts for only 20% of the grade, it will also be used to determine your score for the Implementation aspect of the project (40%).

Your written documentation must discuss each optimization that you considered. The thoroughness of your exploration of the optimization space will be an important aspect of your grade. For each optimization that you implement, your writeup must convince the reader that the optimization was beneficial, general, and correct. For each optimization that you decide not to implement, you must convince the reader that the optimization would not be beneficial given your generated code for the given benchmarks and our target architecture.

You should include assembly code examples for each optimization. Show how your generated code is transformed by the optimization (might be hand-applied for optimizations you did not implement). Highlight the benefit of the optimization on the assembly code. Discuss exactly how the benefit was achieved given the characteristics of our target architecture. Include empirical evidence that proves your conclusion for the optimization.

Your compiler should include a "full optimizations" command-line option (see below). Your written documentation should present a detailed discussion of this option including how you determined the order in which your optimizations are performed and how many times you apply the optimizations. Finally, describe any hacks or solutions to tricky problems that you encountered.

## Design Proposal

As with previous phases, this phase includes a design proposal that must be submitted. This document will allow us to judge your progress and comment on your design. We do not expect this document to be a complete description of your implementation. It should include a detailed discussion of the optimizations that you have considered and which optimizations are planned for implementation. The proposal should include a discussion on the “full optimizations” option.

We expect your proposal to provide you with a substantiated plan of action for implementation. Of course, we are not ruling out the implementation of optimizations not discussed in your proposal.

## Derby

On the final day of class we will hold the compiler derby. Your compiler will be pitted against the submissions from the other groups in the class. The derby benchmark will be revealed one day prior to the derby. 40% of the grade will be determined by your ranking in the derby. At this time, we will not reveal the exact formula for determining this portion of the project grade.

## Compiler Submission

Your compiler must provide a command-line option for each optimization. Your project writeup should include documentation for each command-line option. For example:

- `-opt regalloc`: turns on register allocation
- `-opt instsel`: turns on instructions selection peephole optimizations
- `-opt sched`: turns on list scheduling
- `-opt par`: turns on coarse-grained data parallelization of `forpar` loops

You must provide a `-opt all` flag to turn on all optimizations and apply them in the order you have determined (“full optimizations”). This option should consider how many times each optimization is applied and the application order of the optimizations.

As before, your generated code must perform the runtime checks listed in the language specification. These may be optimized (or removed in some cases) as long as they report a runtime error whenever the unoptimized program reports an error.

Follow the instructions in Handout 3 on what to turn in. You should submit files named `leNN-optimizer.jar` and `leNN-optimizer.tar.gz`.

## Test Cases

We have provided three sample programs that perform image processing and filtering. The programs are located in `/mit/6.035/provided/optimizer/decaf`. These programs must be linked against the library provided in `/mit/6.035/provided/optimizer/lib`. You should make sure that any valid program provided during previous phases continues to run correctly. We will also test your compiler (including optimizations) on a suite of hidden programs.