

15-411 Compiler Design, Fall 2014

Lab 6 - LLVM

Instructor: Frank Pfenning
TAs: Flávio Cruz, Tae Gyun Kim, Rokhini Prabhu, Max Serrano

Compilers due 11:59pm, Thursday, December 4, 2014
Papers due 11:59pm, Tuesday, December 9, 2014

1 Introduction

The main goal of the lab is to explore advanced aspects of compilation. This writeup describes the option of retargeting the compiler to generate LLVM code; other writeups detail other options. The language L4 does not change for this lab and remains the same as in Labs 4 and 5.

2 Requirements

You are required to hand in two separate items:

1. The working compiler and runtime system, and
2. a term paper describing and critically evaluating your project.

2.1 Compilers

Your compilers should treat the language L4 as in Labs 4 and 5. You are required to support safe and unsafe memory semantics. Note that safe memory semantics is technically a valid implementation of unsafe memory semantics; therefore, if you have trouble getting the exception semantics of L4 working in a manner that corresponds directly to x86-64, use the safe semantics as a starting point, and try to remove as much of the overhead as you can.

When generating code for the LLVM, given file *name.l4*, your compiler should generate: *name.ll*, which is in the LLVM human-readable assembly language. The driver will use LLVM commands to generate from this the file *name.s*, in the x86-64 assembly language.

If you would like to apply LLVM optimizations yourself, your compiler may generate a *name.bc* file using `llvm-as` and run optimizations on it. (See the LLVM documentation.) In this case, the driver will ignore any existing *name.ll* file and use the *name.bc* file instead.

After all is said and done, your compiler must support both LLVM and x86-64 as backends.

2.2 Testing

You are not required to hand in new tests. We will use a subset of the tests from the previous labs to test your compiler.

However, if you wish to use LLVM to optimize your code you can use the cycle counting benchmarking tools developed for Lab 5. Porting and applying those tools will be your responsibility and should be described in the term paper.

2.3 Term Paper

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

1. Introduction. This should provide an overview of your implementation and briefly summarize the results you obtained.
2. Comparison. Compare compilation to LLVM, followed by native code generate with direct native code generation. How does the structure of your compiler differ? How does the generated code differ? If you are applying optimizations at the LLVM level, describes those optimizations and their rationale.
3. Analysis. Critically evaluate the results of your compiler via LLVM, which could include size and speed of the generated code. Also provide an evaluation of LLVM: how well did it serve your purpose? What might be improved?

3 Deadlines and Deliverables

Compiler Files (due 11:59pm on Thu Dec 4)

As for all labs, the files comprising the compiler itself should be collected in a directory `compiler/` which should contain a `Makefile`. **Important:** You should also update the `README` file and insert a roadmap to your code. This will be a helpful guide for the grader.

Issuing the shell command

```
% make l4c
```

should generate the appropriate files so that

```
% bin/l4c --safe --llvm <args>
% bin/l4c --unsafe --llvm <args>
% bin/l4c --safe --x86_64 <args>
% bin/l4c --unsafe --x86_64 <args>
```

will run your L4 compiler in safe and unsafe modes, generating LLVM or direct x86-64 native code, respectively. For backwards compatibility, the default is `--safe --x86_64`. The driver will only use the `--llvm` flag to run your code.

In order for you to be able to provide a runtime system or library functions, the compiler expects a file `l4lib.c` at the top-level of your compiler directory and compile and link this against your file when compiled in `--llvm` mode.

The command

```
% make clean
```

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

The compiler should be handed in to Autolab, under `lab6llvm`

Term Paper (due 11:59 on Tue Dec 9)

Submit your term paper in PDF form via Autolab 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!

4 Notes and Hints

- Apply regression testing. It is very easy to get caught up in writing a back end for a new target. Please make sure your native code compiler continues to work correctly!
- Read the LLVM code. Just looking at the LLVM code that your compiler produces will give you useful insights into what you may need to change.
- Study LLVM code generated by other compilers. Particularly, on the Andrew Linux machines you can run

```
clang -S -emit-llvm -O2 <file>.c
```

to generate `<file>.ll` in the LLVM assembly language.

- The intermediate form accepted by LLVM must be in SSA form. However, it is possible to allocate all variables on the stack and use a script provided with LLVM in order to convert into SSA form. See [Chapter 4.7](#) of the LLVM Tutorial.
- LLVM, like C, may leave the result for certain operations undefined (e.g., division by 0). Make sure to use appropriate keyword modifiers and/or turn some basic operations into function calls to ensure correct behavior.