

CS 745, Spring 2011

Homework Assignment 1

Assigned: Thursday, January 13
Due: Thursday, January 27, 9:00AM

Welcome to the Spring 2011 edition of Optimizing Compilers (15-745). We will be using the Low Level Virtual Machine (LLVM) Compiler infrastructure from University of Illinois Urbana-Champaign (UIUC) for our programming assignments. While LLVM is currently supported on a number of hardware platforms, we expect the assignments to be completed on x86 machines, since that is where they will be graded. Although LLVM works quite well on both Mac OS X and Windows, it is recommended that assignments be done in Linux, to increase the chances of getting technical support from the teaching staff.

The objective of this first assignment is to introduce you to LLVM and some ways that it could be used to make your programs run faster. In particular, you will use LLVM to learn interesting properties about your program and to perform local optimizations.

Policy

You will work in groups of two people to solve the problems for this assignment. Turn in a single writeup per group, indicating all group members.

Logistics

Any clarifications and revisions to the assignment will be posted on the “assignments” page on the class web page.

In the following, *HOMEDIR* refers to the directory:

```
/afs/cs.cmu.edu/academic/class/15745-s11/public
```

and *ASSTDIR* refers to the subdirectory *HOMEDIR/asst/asst1*.

1 Install LLVM

First download, install, and build LLVM 2.8 source code from <http://llvm.org>. Do not get the source from SVN, instead get the 2.8 release. To get started, follow the instructions at <http://llvm.org/docs/GettingStarted.html> for your particular machine configuration. You do not need to build the gcc and g++ frontends from source, instead follow the provided instructions to install prebuilt binaries. Be careful not to overwrite the current gcc install on your system (either `/usr/bin/gcc` or `/usr/local/bin/gcc`), rather install your LLVM gcc in a private directory and update your PATH environment variable appropriately so that you have `llvm-gcc`, `llvm-g++`, `opt`, etc., in your PATH. Note that in order to use a debugger on the LLVM binaries you will need to pass `--enable-debug-runtime --disable-optimized` to the configure script.

Peruse through the documentation at <http://llvm.org/docs>. The LLVM Programmer’s Manual (<http://llvm.org/docs/ProgrammersManual.html>) and Writing an LLVM Pass Tutorial (<http://llvm.org/docs/WritingAnLLVMPass.html>) are particularly useful.

```

int g;
int g_incr (int c)
{
  g += c;
}
int loop (int a, int b, int c)
{
  int i;
  int ret = 0;
  for (i = a; i < b; i++) {
    g_incr (c);
  }
  return ret + g;
}
(a)

@g = common global i32 0

define i32 @g_incr(i32 %c) nounwind {
entry:
  %0 = load i32* @g, align 4
  %1 = add nsw i32 %0, %c
  store i32 %1, i32* @g, align 4
  ret i32 undef
}

define i32 @loop(i32 %a, i32 %b, i32 %c) nounwind {
entry:
  %0 = icmp slt i32 %a, %b
  %1 = load i32* @g, align 4
  br i1 %0, label %bb.nph, label %bb2

bb.nph:
  %tmp = sub i32 %b, %a
  %tmp7 = mul i32 %tmp, %c
  %tmp8 = add i32 %1, %tmp7
  store i32 %tmp8, i32* @g, align 4
  ret i32 %tmp8

bb2:
  ret i32 %1
}
(b)

```

Figure 1: (a) A simple loop source code, and (b) its optimized LLVM bytecode.

2 Create a Pass

Assuming the installation directory of your LLVM source is `llvm/`, create a directory (e.g, named `FunctionInfo`) within the `llvm/lib/Analysis` directory. Copy `FunctionInfo.cpp` (provided with the assignment) into the new directory. `FunctionInfo.cpp` contains a dummy LLVM pass for analyzing the functions in a program. Currently it only prints out “15-745 Functions Information Pass”. In the next section, you will extend `FunctionInfo.cpp` to print out more interesting information. For now, we will use the dummy LLVM pass to demonstrate how to build and run LLVM passes on programs. First, create a `Makefile` to build the `FunctionInfo` pass as follows:

```

LEVEL = ../../..
LIBRARYNAME = FunctionInfo
LOADABLE_MODULE = 1
include $(LEVEL)/Makefile.common

```

(Note: you can also copy this code from `ASSTDIR/FunctionInfo/Makefile`.) Before moving on to the next section, make sure you can run this dummy pass properly. Copy the `loop.c` source code (shown in Figure 1(a)) from `ASSTDIR/FunctionInfo/loop.c` into your local directory. Compile it to an optimized LLVM bytecode (`loop.o`) as follows:

```

llvm-gcc -O -emit-llvm -c loop.c

```

Inspect the generated bytecode using `llvm-dis` as follows:

```

llvm-dis loop.o

```

This will create a disassembly of the testcase named `loop.ll` that should look very similar to Figure 1(b).

Name	# Args	# Calls	# Blocks	# Insts
g_incr	1	0	1	4
loop	3	0	3	9

Table 1: Expected FunctionInfo output for the optimized bytecode of `loop.c`

Now try running the dummy `FunctionInfo` pass on the bytecode using the `opt` command. (If you did not compile with debug information, the shared library (`FunctionInfo.so`) will be in the `Release` directory). Note the use of the command line flag “`-function-info`” to enable this pass. (See if you can locate the declaration of this flag in `FunctionInfo.cpp`).

```
opt -load llvm/Debug/lib/FunctionInfo.so -function-info loop.o -o out
```

If all goes well, “15745 Functions Information Pass” should be printed to `stderr`.

3 Meet The Functions

Program analysis is an important prerequisite to applying correct optimizations: i.e. without breaking the code. For example, before the optimizer can remove some piece of code to make a program run faster, it must examine other parts of the program to determine whether the code is truly redundant. A compiler pass is the standard mechanism for analyzing and optimizing programs.

You will now extend the dummy `FunctionInfo` pass from the previous section to learn interesting properties about the functions in a program. Your pass should report the following information about all functions that are used in a program:

1. Name.
2. Number of arguments.
3. Number of call sites (i.e. locations where this function is called).
4. Number of basic blocks.
5. Number of instructions.

To assist you in writing this pass, the expected output of running `FunctionInfo` on the optimized bytecode (Figure 1(b)) is shown in Figure 1. As you can see, the output in Figure 1 is not interesting, since `loop.c` is a trivial piece of code. It is therefore recommended that you debug your pass with more complex source files, as you can imagine grading will be done with complex programs.

3.1 Using FunctionInfo Results

Examine the results of your pass across different programs and explain any interesting observations.

1. You should notice that some functions have zero basic blocks: why is that?

2. Describe how the results could be used by an optimizer to determine which functions should be inlined.
3. What other ways can your pass results be used?

4 Optimize The Block (New Dragon Book 8.5)

Now that you are an expert writing LLVM passes, it is time to write a pass for making programs faster. You will implement optimizations on basic blocks as discussed in class. More details on local optimizations are available in Chapter 8.5 of the new Dragon book. While there are many types of local optimizations, we will keep things quite simple in this section and focus only on the algebraic optimizations discussed in Section 8.5.4 of the book. Specifically, you will implement the following local optimizations:

1. Algebraic identities: e.g, $x + 0 = 0 + x = x$
2. Constant folding: e.g, $2 * 4 \Rightarrow 8$
3. Strength reductions: e.g, $2 * x \Rightarrow (x + x)$ or $(x \ll 1)$

4.1 Implementation Details

You should create a new LLVM pass named `LocalOpts.cpp` following the steps in Section 2. While it is possible to implement more than one pass in the same directory or file, it is probably much easier at this point to simply create a new `LocalOpts` directory in `llvm/lib/Analysis`. Since the `llvm-gcc` optimizer performs local optimizations, your `LocalOpts` pass should be run on unoptimized LLVM bytecodes. Unoptimized bytecode of `loop.c` can be prepared as follows:

```
llvm-gcc -O0 -emit-llvm -c loop.c
```

Now assuming the command line flag for enabling your local optimization pass is `-my-local-opts`, then you can run your pass as follows:

```
opt -load llvm/Debug/lib/LocalOpts.so -my-local-opts loop.o -o out
```

In addition to transforming the bytecode, your pass should also print out a summary of the optimizations it performed: e.g., how many constants were folded. We will provide toy source files with unrealistic amounts of local optimization opportunities for you to debug your pass in: `ASSTDIR/LocalOpts/test-inputs`. In addition to using these test inputs, we recommend that you test your pass on more realistic programs.

5 Hand In

Electronic submission:

- The source code for your passes (`FunctionInfo` and `LocalOpts`), the associated `Makefiles`, and a `README` describing how to build and run them. Do this by creating a tar file with the last name of at least one of your group members in the filename, and copying this tar file into the directory

`/afs/cs.cmu.edu/academic/class/15745-s11/public/asst/asst1/handin`

Include as comments near the beginning of your source files the identities of all members of your group. Also remember to do a good job of commenting your code.

Hard-copy submission:

1. A report that briefly describes the implementations of both passes.
2. A listing of your source code.
3. Answers to the short questions in 3.1.