<div align="center">

15-745, Spring 2015
Homework Assignment 1

Assigned: Thursday, January 15
Due: Thursday, January 29, 9:00AM

</div>

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

All clarifications (if any) to this assignment will be posted on the class discussion board on Piazza. Any revisions will be uploaded to the "assignments" page on the class web page.

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

`/afs/cs.cmu.edu/academic/class/15745-s15/public`

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

1 Install VirtualBox and the 15-745 System Image

To keep you from having to build LLVM yourself and to ensure that all assignments are graded in the same environment, we are distributing a system image for VirtualBox based on 32-bit Ubuntu 14.04.1 LTS Desktop (Trusty Tahr). You must ensure that all of your code works in this image, but you are of course not required to do all of your development with it.

The VirtualBox software is available on several platforms from `http://www.virtualbox.org`. We will use version 4.3.20. You may need to enable your machine's virtualization extensions in your BIOS setup (on some office machines, reboot, press F12 to get the boot menu, choose System Setup, then Virtualization Support, then make sure that the box is checked).

Once you have VirtualBox installed, you can retrieve the virtual machine image from

`/afs/cs.cmu.edu/academic/class/15745-s15/www/vm-images/15-745-s15.ova`

There is a checksum file in the same directory, `15-745-s15.ova.sha1`, that you may use (`sha1sum -c 15-745-s15.ova.sha1`) to verify that the image transferred correctly.

If you prefer, you may also download the file over HTTP via:

`http://www.cs.cmu.edu/afs/cs.cmu.edu/academic/class/15745-s15/www/vm-images/15-745-s15.ova`

The machine name is `15-745-s15`; you may log in with username `user` and password `user`. LLVM binaries are in `/home/user/llvm/llvm-3.5-install` (or `LLVM_ROOT`) and source files are in `/home/user/llvm/llvm-3.5-src`. `LLVM_ROOT/bin` is also added to the `PATH`.

We built LLVM 3.5 in the following way:

Command	Options	Meaning
`clang`		Compile high-level source code, such as C
	`-O` (the letter)	Set level of optimization performed by `clang` (to the default)
	`-O0` (the letter, followed by zero)	Direct `clang` not to perform any optimization
	`-emit-llvm`	Output an LLVM bytecode object
	`-c`	Output object code, do not fully compile
`llvm-dis`		Generate disassembly of LLVM bytecode
`opt`		Run LLVM passes
	`-load`	Load the pass found in the corresponding file (file path must be specified)
	`-{passname}`	Load the pass with this name
	`-mem2reg`	Load the mem2reg pass, which simplifies LLVM bytecode
	`-o`	Specify the output filename

Table 1: A simplified summary of some of the `LLVM` commands you will use in this assignment.

```
mkdir ~/llvm/llvm-3.5-build ; cd ~/llvm/llvm-3.5-build
../llvm-3.5-src/configure --prefix=/home/user/llvm/llvm-3.5-install
make -j 4
make install
```

Moving files between the VM and your host machine is easy using VirtualBox. If you set up a shared directory in the VirtualBox GUI under Settings, Shared Folders (pointing somewhere on your local filesystem). You can mount it in the VM by using the auto-mount option when you set up the shared folder. You can then find the shared folder at

```
/media/sf_foldername
```

where foldername is the name of the shared folder. If you run into permissions errors with your shared folder, you may need to add `user` to the `vboxsf` group by running

```
sudo usermod -a -G vboxsf user
```

Look through the documentation at `http://llvm.org/releases/3.5.0/docs/`. The LLVM Programmer's Manual (`http://llvm.org/releases/3.5.0/docs/ProgrammersManual.html`) and Writing an LLVM Pass Tutorial (`http://llvm.org/releases/3.5.0/docs/WritingAnLLVMPass.html`) are particularly useful. Note that while we are using the latest released version of LLVM, the primary LLVM documentation (`http://llvm.org/docs/`) corresponds to the latest source code in the LLVM SVN repository. While there should not be any changes that affect the sorts of things we will be doing in this class, you may want to stick to the 3.5 documentation just to be on the safe side.

2   Create a Pass

Create a directory named `FunctionInfo` and 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 prints out "15-745 Function Information Pass". In the

```
                                    @g = common global i32 0, align 4

                                    ; Function Attrs: nounwind
                                    define i32 @g_incr(i32 %c) #0 {
                                    entry:
                                      %0 = load i32* @g, align 4, !tbaa !1
                                      %add = add nsw i32 %0, %c
 int g;                             store i32 %add, i32* @g, align 4, !tbaa !1
 int g_incr (int c)                 ret i32 %add
 {                                  }
   g += c;
   return g;                        ; Function Attrs: nounwind
 }                                  define i32 @loop(i32 %a, i32 %b, i32 %c) #0 {
 int loop (int a, int b, int c)     entry:
 {                                    %cmp2 = icmp sgt i32 %b, %a
   int i;                             %0 = load i32* @g, align 4, !tbaa !1
   int ret = 0;                       br i1 %cmp2, label %for.body.lr.ph, label %for.end
   for (i = a; i < b; i++) {
    g_incr (c);                     for.body.lr.ph:                    ; preds = %entry
   }                                  %1 = sub i32 %b, %a
   return ret + g;                    %2 = mul i32 %1, %c
 }                                    %3 = add i32 %0, %2
                                      store i32 %3, i32* @g, align 4, !tbaa !1
                                      br label %for.end

                                    for.end:                           ; preds = %for.body.lr.ph, %entry
                                      %.lcssa = phi i32 [ %3, %for.body.lr.ph ], [ %0, %entry ]
                                      ret i32 %.lcssa
                                    }

              (a)                                            (b)
```

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

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, use the `Makefile` we have provided or create a `Makefile` to build the `FunctionInfo` pass as follows (these instructions assume that your passes are the only `.cpp` files in the directory. Make sure that there are tabs on lines 6 and 8 below):

```
all: FunctionInfo.so

CXXFLAGS = -rdynamic $(shell llvm-config --cxxflags all) -g -O0

%.so: %.o
        (CXX) -dylib -shared $^ -o $@
clean:
        rm -f *.o *~ *.so
```

Before moving on, make sure you can `make` this dummy pass.

Next, copy the `loop.c` source code (shown in Figure 1(a)) from *ASST-DIR*/`FunctionInfo/loop.c` into your local `FunctionInfo` directory. Compile it to an

optimized LLVM bytecode object (`loop.bc`) as follows:

```
clang -O -emit-llvm -c loop.c
```

(`clang` is the LLVM project's frontend for the C language family.)

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

```
llvm-dis loop.bc
```

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

Now, try running the dummy `FunctionInfo` pass on the bytecode. To do this, use the `opt` command listed below. 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`). Note that you must provide the correct path to FunctionInfo.so. Use "./" if you are in the same directory.

```
opt -load path/to/FunctionInfo.so -function-info loop.bc -o out
```

If all goes well, "15745 Function Information Pass" should be printed to stderr.

3   Meet The Functions

Program analysis is an important prerequisite to applying correct optimizations: we want to improve code, not break it. 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 appear in a program:

1. Name.

2. Number of arguments (or `*` if variable).

3. Number of direct call sites in the same LLVM module (i.e. locations where this function is explicitly called, ignoring function pointers).

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 Table 2. As you can see, the output in Table 2 is not that interesting, since `loop.c` is a fairly trivial piece of code. Note, however, that although the source code for `loop.c` has a call to `g_incr` in `loop`, this call is optimized away in the LLVM bytecode, even when using the `-O0` flag. When reporting the number of calls, please count the number

Name	# Args	# Calls	# Blocks	# Insns
g_incr	1	0	1	4
loop	3	0	3	10

Table 2: Expected FunctionInfo output for the optimized bytecode of `loop.c`

that appear in the bytecode, even if this does not match the number of calls in the original source code.

It is recommended that you debug your pass with more complex source files, as you can imagine grading will be done with complex programs. Feel free to hand in your additional testing source files in a separate directory together with your source code.

Hint: It may be useful to look at the documentation at the following URL, along with other documentation:

`http://llvm.org/doxygen/classllvm_1_1Function.html`

Again, note that this is not precisely the 3.5 documentation, although it should be sufficient for our work in this class.

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 => 8`

3. Strength reductions: e.g, `2 * x => (x + x) or (x << 1)`

This is a somewhat open-ended question. Please handle at least the above cases, as well as one more in each category that you come up with, for (scalar) integer types.

For the purposes of this assignment, please do not worry about divide-by-zero issues.

4.1   Implementation Details

You should create a new LLVM pass in a file named `LocalOpts/LocalOpts.cpp` following the steps in Section 2. Because this will be an optimization pass rather than an analysis pass, there will be some small differences from the set up of the FunctionInfo pass. Provide an appropriate makefile at `LocalOpts/Makefile`. (Note that it is possible to implement more than one pass in the same directory or file, but we're trying to keep things clean.) `clang` may apply these kinds of local optimizations during the course of regular compilation. To better test your pass, you should build mostly unoptimized LLVM bytecode from the test cases:

```
clang -O0 -emit-llvm -c loop.c
opt -mem2reg loop.bc -o loop-m2r.bc
```

You may assume that all input to your pass will first go through `mem2reg` as shown above.

We should be able to run your local optimization pass in the following way, from the location of the shared library:

```
opt -load ./LocalOpts.so -some-local-opts loop-m2r.bc -o out
```

In addition to transforming the bytecode, your pass should also print to standard out a summary of the optimizations it performed. There is no canonical format for this output, but you should at least try to categorize and count the transformations your pass applies:

```
Transformations applied:
  Algebraic identities: 2
  Constant folding: 1
  Strength reduction: 3
```

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   Questions

5.1   CFG Basics

For the code provided below (i) find (maximal) basic blocks (ii) build the CFG (Control Flow Graph). Be sure to give your basic blocks clear labels (and label the original code to match). You may assume that the `print` function always returns.

```
    x = 50
    y = 8
    if (x > y) goto L2
    x = 50
    goto L3
L1: x = 27
    return x
L2: y = x + 1
    if (y < x) goto L2
    x = 24
L3: z = x + y
    switch (z) { 26 => L1 | 32 => L4 | default => L5 }
L4: print("success")
L5: x = 50
    return x
```

5.2   Available Expressions, New Dragon Book 9.2.6

An expression x op y is *available* at a point p if every path from the entry node to p evaluates x op y, and after the last such evaluation prior to reaching p, there are no subsequent assignments to x or y. For the *available-expressions* data-flow schema we say that a block *kills* expression x op y if it assigns (or may assign) x or y and does not subsequently recompute x op y. A block *generates* expression x op y if it definitely evaluates x op y and does not subsequently define x or y.

Based on this definition and the corresponding data flow analysis description (See Table 3 from New Dragon Book 9.2.7) perform Available Expressions analysis on the code in Figure 2.

Domain	**Direction**	**Transfer Function**	**Boundary**
Sets of expressions	Forwards	$gen_B \cup (x - kill_B)$	$OUT[entry] = \oslash$
Meet $\wedge$	**Equations**	**Equations**	**Initialize**
$\cap$	$OUT[B] = f_B(IN[B])$	$IN[B] = \bigwedge_{P,pred(B)} OUT[P]$	$OUT[B] = \mathscr{U}$

Table 3: Available Expressions Analysis.

In the following tables, list the EVAL and KILL sets, then the final IN and OUT sets after AE is performed. You may ignore expressions inside conditional statements (e.g., $(z < c)$).
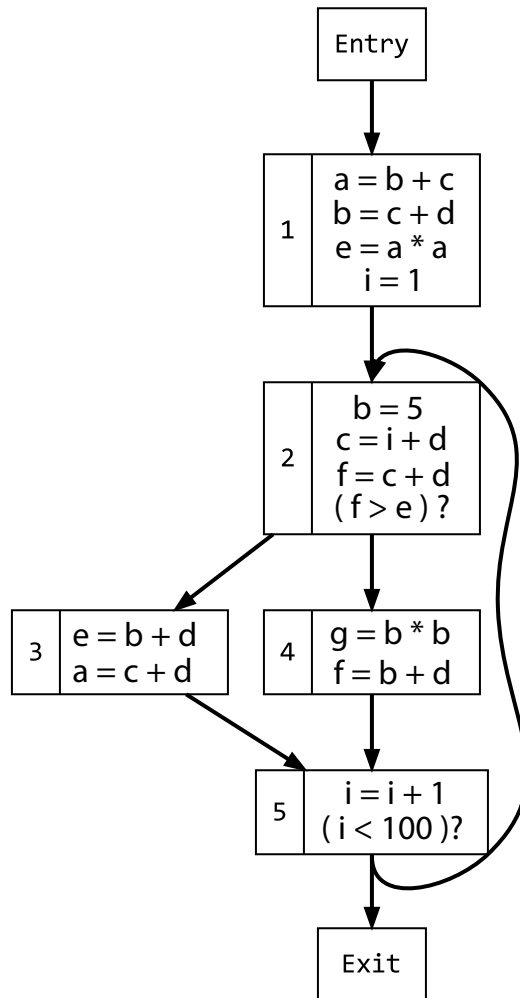
Figure 2: Code for Available Expressions Analysis.

BB	EVAL	KILL
1		
2		
3		
4		
5		

BB	IN	OUT
1		
2		
3		
4		
5		

5.3 New Dataflow Analysis: Faint Analysis

You have been hired to help develop a software analysis package that will detect *faint* expressions, which are useful to perform Dead Code Elimination (*DCE*). The idea behind DCE is that an assignment of the form "x = t" can be eliminated if its LHS variable x is not live (i.e dead) at the program point P immediately following the assignment. One of the limitations of DCE is that it cannot directly eliminate the assignment "x = x + 1" in the two examples shown below:
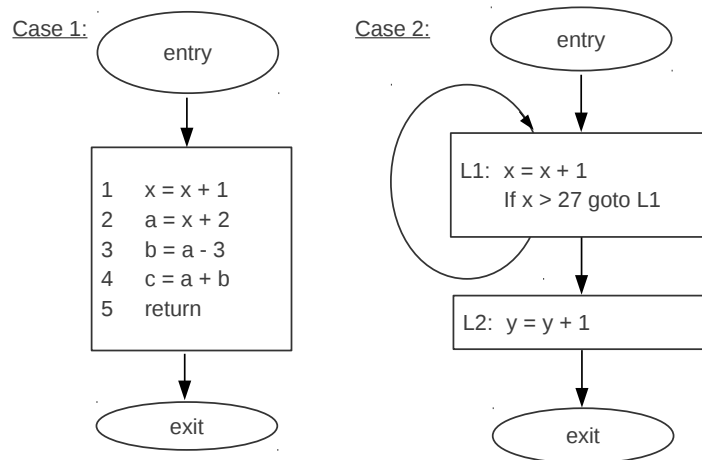
In the first case, x is not dead after the "x = x + 1" assignment (instruction 1) because it is used in instruction 2. Instruction 2 is also not dead because its LHS variable (a) is used in instructions 3 and 4. However, instruction 4 is in fact dead. If we applied DCE repeatedly to this code, we could eventually eliminate instruction 1. However, it would be more desirable to eliminate "x = x + 1" in a single data flow pass.

In the second case, the LHS of "x = x + 1" is not dead because it is used by its own RHS due to the cycle in the flow graph. However, since the ultimate value of x is never used, this instruction could in fact be safely eliminated from the loop body.

We say that the LHS variable x in an assignment "x = t" is *faint* if along every path following the assignment, x is either dead or is only used by an instruction whose LHS variable is also faint.

Your mission in this assignment is to design a new dataflow analysis pass specifically for determining whether the LHS variable of an expression is faint, in both of these cases.

Your analysis should be as simple as possible (i.e., it should not gather unnecessary information), and as fast as possible. Your analysis will be plugged into a generic dataflow framework (e.g., New Dragon Book 9.2-9.3).

1. What is the set of elements that your analysis operates on?

2. What is the direction of your analysis?

3. What is your transfer function? Be sure to clearly define any other sets that your transfer function uses (eg., GEN or KILL etc).

4. What is your meet operator? Give the equation that uses the meet operator.

5. To what value do you initialize exit and/or entry?

6. To what values do you initialize the in or out sets?

7. Does the order that your analysis visits basic blocks matter? What order would you implement and why?

8. Will your analysis converge? Why (in words, not a proof)?

9. Clearly describe in pseudo-code an algorithm that uses the result of your analysis to identify faint expressions

6   Hand In

Hard-copy submission:

1. A report that briefly describes the implementations of both passes. Explanations of your thought process and/or any issues you had in the development process would be welcome.

2. A listing of your source code. One possible way to generate this is by using `enscript`:

   ```
   enscript -q -DDuplex:true -r -2 -E -fCourier7 --tabsize=2 -p listing.ps
   ps2pdf listing.ps listing.pdf
   ```

3. Listings of additional tests that you used for verification of your passes, as well as their expected results. We expect you to provide at least one of these, and we will use some of them to test others' solutions.

4. Answers to the questions in Section 5.

Note: please include the Andrew IDs of group members on each page or on each stapled group of pages.

Electronic submission:

- A PDF of your writeup report and answers to the questions, named `writeup.pdf`.
- The source code for your passes (`FunctionInfo` and `LocalOpts`), the associated `Makefile`s, and a `README` describing how to build and run them (especially if you for some reason diverge significantly from what the assignment requires). Place all of these files in a directory with the same name as the Andrew ID of one of your group members. Archive this directory and name the file with the same Andrew ID (`bovik.tar.gz`):

  ```
  tar czvf bovik.tar.gz bovik
  ```

  When the file is extracted with `tar xf`, we expect to see these required files in these locations:

  ```
  ./bovik/README
  ./bovik/FunctionInfo/FunctionInfo.cpp
  ./bovik/FunctionInfo/Makefile
  ./bovik/LocalOpts/LocalOpts.cpp
  ./bovik/LocalOpts/Makefile
  ./bovik/writeup.pdf
  ```

  It is fine if there are other files included; please also include any additional tests you used for verification.
  Copy the tar.gz file into the directory

  ```
  /afs/cs.cmu.edu/academic/class/15745-s15/public/asst/asst1/handin
  ```

  Include as comments near the beginning of your source files the identities of all members of your group. Please comment your code.
  Note: You will not be able to modify or replace your submission once you put it in the `handin` folder. If you need to update your submission, please add a new tgz file with a version number at the end, such as `bovik-v2.tar.gz`.