

EECS 583 – Homework 1

Fall 2023

Assigned: Fri, September 8, 2023

Due: Mon, September 18, 2023 (11:59pm Eastern USA)

Late Submission Policy

Submissions are accepted a maximum of **two days** after the specified deadline. For each day late, a **10% penalty** will be deducted from your score.

Statistics Computation Pass

The goal of this homework is to learn to write your first real LLVM pass. As part of this, you will learn to use the profiler which provides dynamic execution frequencies for the compiler to make use of.

Write a statistics computation pass in LLVM that computes several dynamic operation counts for each function. First, the total number of dynamic operations should be computed along with the percentages in the following categories: integer ALU, floating-point ALU, memory, biased-branch, unbiased branch, and all other operations. Use the following rules when categorizing the operations:

- Branch: br, switch, indirectbr
- Integer ALU: add, sub, mul, udiv, sdiv, urem, shl, lshr, ashr, and, or, xor, icmp, srem
- Floating-point ALU: fadd, fsub, fmul, fdiv, frem, fcmp
- Memory: alloca, load, store, getelementptr, fence, atomiccmpxchg, atomicrmw
- Others: everything else

Every operation should be placed in one of the above categories. Print this information out in a text file named **benchmark.opcstats** in the following format (comma separated, one function in each line):

FuncName, DynOpCount, %IALU, %FALU, %MEM, %Biased-Br, %Unbiased-Br, %Others

FuncName is the name of the function, DynOpCount refers to the dynamic operation count, i.e. the total number of all instructions that were executed as part of that function, when we ran the program. Note that this differs from the Static Operation count, which will be the number of instructions present in the static IR of the program. The remaining fields are percentage values computed to the third decimal place. All percentage values are computed as the targeted instruction type over the DynOpCount (e.g. %Biased-Br = biased_branch_inst_count / total_dyn_inst_count).

For example, if 50% of the function's operations are integer ALU operations, then its %IALU should be 0.500. You can use the following to print these values: `#include "llvm/Support/Format.h"` and `errs() << format("%.3f", val)`. Print all zeros if a function has never been executed, i.e. (func_name, 0, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000)

The bias of a branch is: (frequency_of_most_likely_target / total_execution_frequency). For branches with only a single target, the branch bias is considered to be 100%. Branches are considered biased if the bias > 80% and unbiased otherwise.

To get started, download the compressed file from the course website. This file contains 3 benchmarks, skeleton code for your pass and a run script for running LLVM. To unpack the file, run `tar -zxvf <file_name>.tgz`

Benchmarks to Run

There are 3 benchmarks that you should profile and collect statistics for: *simple* (*benchmark1*), *anagram* (*benchmark2*), and *compress* (*benchmark3*). Each is progressively larger and more complex. Each benchmark contains directories for the source code (src) and the input file (input).

The *simple* benchmark is just a C code that contains a single main function with loops and array accesses. The *anagram* benchmark takes a few input words and a dictionary of words and computes all possible anagrams of the input words. Note that this benchmark takes time to run (30-60 sec). The *compress* benchmark runs a compression algorithm on the given input file.

We recommend that you use the run script (run.sh) that we have provided to execute your pass on the benchmarks. Instructions for using the script are given in comments at the beginning of the script.

If running it manually remember to copy the input file to the same directory as your binary executable to ensure that your profile counts match ours. You can compile each benchmark with gcc to make sure they work before running LLVM. The compress benchmark may check if an output with the same name already exists (compress.in.Z). This will make a difference in program execution, and lead to different statistics. Therefore, please make sure you delete outputs from the previous run before collecting your statistics.

Submission

To submit your homework, put a single .tgz (gzipped tar file) into the directory **/hw1_submissions/** on **eeecs583a.eecs.umich.edu** via scp (later versions will automatically overwrite the earlier versions if you submit multiple times):

```
$ tar cvzf ${uniquename}_hw1.tgz ${uniquename}_hw1
$ scp ${uniquename}_hw1.tgz {uniquename}@eeecs583a.eecs.umich.edu:/hw1_submissions/
```

Please name your tar file `${uniquename}_hw1.tgz`, and organize the directory as follows:

```
${uniquename}_hw1/
  src/
    hw1_pass.cpp (your source code)
  results/
    simple.opcstats
    anagram.opcstats
    compress.opcstats
```

Since you will be using cp/scp to submit, there is no validation on submitting. Use `ls` on the submission folder to verify that your submission is present.