

# ECE571: Advanced Microprocessor Design – Homework 1

**Due: Friday 13 September 2024, 12:00pm**

## 1. Background

- For this assignment, log into the Haswell-EP machine using the username and password on the account slip that I handed out in class.
- On Linux or OSX you will do the following (replace `username` with the one on the slip):  

```
ssh -p 2131 username@weaver-lab.eece.maine.edu
```
- On a Windows machine you'll want to get a program such as `putty`, some directions can be found here:  
[https://web.eece.maine.edu/~vweaver/classes/ece571\\_2013s/using\\_ssh.html](https://web.eece.maine.edu/~vweaver/classes/ece571_2013s/using_ssh.html)  
Be sure you are connecting to port 2131 (not the default ssh port).
- Be sure to change your password using the `passwd` command once you log in.
- We will use the `401.bzip2` benchmark from the SPEC CPU 2006 benchmark suite.
- Create a document that answers the questions / contains the data from the sections below. A `.pdf` or `.txt` file is preferred but I can accept MS Office format if necessary.

## 2. Aggregate Event Counts

### (a) perf tool

- First copy the input file to your local directory: (be sure to include the period in the command, it means copy to current directory)  

```
cp /opt/ece571/401.bzip2/input.source .
```
- Use the `perf` tool to gather user instruction counts for `bzip2`:  

```
perf stat -e instructions:u /opt/ece571/401.bzip2/bzip2 -k -f ./input.source
```

#### i. Run the benchmark 5 times.

Report the instruction count for each, as well as the overall average.

- ii. Note! Your result should be more than 1 million instructions! If you get less and can't figure out why, let me know so we can make sure you are running things properly.
- iii. Run the benchmark 5 more times, but this time measure user cycles (`cycles:u`) rather than instructions.  
Report the cycle count for each, as well as the overall average.

#### iv. Now gather and report the results for `bzip2.reverse` which is the same benchmark, but compiled with the link order reversed (reverse alphabetical rather than alphabetical).

```
perf stat -e instructions:u,cycles:u /opt/ece571/401.bzip2/bzip2.reverse -k -f ./input.source
```

Gather results for instructions and cycles (5 times) and report the individual and overall average results.

### (b) Questions to Answer

- i. Are the instruction counts deterministic, or do they vary? How large is the variation?
- ii. Are the cycle counts deterministic, or do they vary? How large is the variation?
- iii. Does changing the link order change the instructions or cycle metrics?

### 3. Sampled Results

#### (a) perf

- i. Use `perf` to gather sampled data on the benchmark:

```
perf record -e instructions:u /opt/ece571/401.bzip2/bzip2 -k -f ./input.source
```

Note how long this took to run (try running with `time` before the command to get wallclock time).

- ii. Use `perf report` and report the top 5 most used functions.
- iii. Use `perf annotate` to report which assembly instruction caused the most CPU use, as well as a few instructions on either side. `Perf annotate` will center around the most frequent instruction.

#### (b) Valgrind DBI tool

- i. Use `valgrind` to gather sampled data.

```
time valgrind --tool=callgrind /opt/ece571/401.bzip2/bzip2 -k -f ./input.source
```

Note how long it takes (note: it may take a while).

- ii. Use `callgrind_annotate` for a report on the most used functions; report the top 5.

#### (c) gprof

- i. The `bzip2.gprof` binary was compiled with `-pg` profiling support. Gather profiling data with it, note how long it took to run.

```
time /opt/ece571/401.bzip2/bzip2.gprof -k -f ./input.source
```

- ii. Get a report on the most used functions, report the top 5

```
gprof /opt/ece571/401.bzip2/bzip2.gprof
```

#### (d) Questions to Answer

- i. Did the three different methods of gathering function CPU use return the same results?
- ii. What were the relative speeds of the various methods of gathering the information?

### 4. Skid

- Run the `perf record / perf annotate` results, but use `cycles` this time.

```
perf record -e cycles /opt/ece571/401.bzip2/bzip2 -k -f ./input.source
```

note which instruction has the highest percent from `perf annotate`

- Run again, but this time use special event `cycles:ppp` instead of just `cycles` and note the highest percent
- Questions to Answer:
  - (a) Which instruction was reported as taking the most time for the two cases?
  - (b) Was it the same in each case?
  - (c) What might be the cause of this difference?

### 5. Submitting your work.

- Create the document containing the data as well as answers to the questions asked. (Text of pdf is best, Word/LibreOffice if you must).
- Please make sure your name appears in the document.
- e-mail the file to me by the homework deadline.