

Homework 2

Software Systems
Spring 2005

Allen B. Downey

Due: Thursday 10 February

The purpose of this assignment is to get you to explore the use of memory caches on real machines, and to give you some practice with Linux utilities.

2.1 Get the program and compile it

1. Pick up the “Cache-testing program”:

```
wget wb/ss/code/cache.tar
tar -xf cache.tar
cd cache
make
```

2. Read the code. There are some system calls and things here that we will probably want to use later, so keep this program around for future reference.
3. Compile the program by typing `make`, which reads `Makefile` and performs the appropriate steps. You might want to check out the documentation of `make` and `gcc`.
4. Run the program by typing

```
./cache
```

After a few lines you will get a sense of what the output looks like; you can stop the program by typing Control-C. Take a look at the code and figure out what the units are for Size and Stride.

5. When you run the program again, you will probably want to redirect the output into a file. For example,

```
./cache > data
```

This puts the output in a file named `data`.

2.2 Start poking around the data

1. You can use `more`, `vi` or `emacs` to read the data file.
2. You can use `awk` to extract particular lines from the file. For example:

```
awk '$2 == 4096' < data
```

will print only the lines where the second column is 4096.

```
awk '$2 == 4096 {print $4}' < data
```

selects the same lines as the previous command, but instead of printing the whole line, it only prints the fourth column.

3. You can pipe the output of one command into another. For example

```
awk '$2 == 4096 {print $4, $6}' < data | xgraph
```

Finds the lines with Size 4096, prints Stride and access time, and sends the result to `xgraph`, which plots the data.

2.3 Figure out the cache system

1. Make whatever plots you need to infer the cache size and the block size of the cache. As a JFFE, try to infer the replacement policy (associative, direct-mapped, something else?). You can use whatever program you like to make plots, but I recommend either `xgraph` or `gnuplot`.

2.4 What to turn in

I would like a single report from each pair, with both names on it. The report should include whatever graphs you used, but a bunch of graphs do not make a report. Here is an outline you might want to use to organize your report:

- What you set out to discover.
- How the experimental apparatus works.
- How, in theory, you expect the data to look, given an idealized model of the system.
- How, in fact, the data look.
- What, if anything, we can infer from the data.
- What anomalies there are in the data (things that deviate from our theoretical expectation), and any explanation you have for them. By the way, it is ok to acknowledge the existence of these anomalies even if you can't explain them.

Report data sparingly. Do not overwhelm me with uninterpreted graphs. Use more words than numbers. Never report a number without units (unless it is truly dimensionless). Label the axes on all graphs. Write on the graphs.