

# CSCI 4320 (Principles of Operating Systems), Fall 2010

## Homework 5

**Credit:** 40 points.

### 1 Reading

Be sure you have read Chapter 3, sections 3.4 through 3.9.

### 2 Problems

Answer the following questions. You may write out your answers by hand or using a word processor or other program, but please submit hard copy, either in class or in my mailbox in the department office.

- (5 points) The operating system designers at Acme Computer Company have been asked to think of a way of reducing the amount of disk space needed for paging. One person proposes never saving pages that only contain program code, but simply paging them in directly from the file containing the executable. Will this work always, never, or sometimes? If “sometimes”, when will it work and when will it not? (*Hint:* Search your recollections of CSCI 2321 — or another source — for a definition of “self-modifying code”.)
- (5 points) How long it takes to access all elements of a large data structure can depend on whether they’re accessed in contiguous order (i.e., one after another in the order in which they’re stored in memory), or in some other order. The classic example is a 2D array, in which performance of nested loops such as

```
for (int r = 0; r < ROWS; ++r)
  for (int c = 0; c < COLS; ++c)
    array[r][c] = foo(r,c);
```

can change drastically for a large array if the order of the loops is reversed. Give an explanation for this phenomenon based on what you have learned from our discussion of memory management. For extra credit, give another explanation that is actually probably likelier to be true of current systems.

- (5 points) Consider (imagine?) a very small computer system with only four page frames. Suppose you have implemented the aging algorithm for page replacement, using 4-bit counters and updating the counters after every clock tick, and suppose the  $R$  bits for the four pages are as follows after the first four clock ticks.

Time	$R$ bit (page 0)	$R$ bit (page 1)	$R$ bit (page 2)	$R$ bit (page 3)
after tick 1	0	1	1	1
after tick 2	1	0	1	1
after tick 3	1	0	1	0
after tick 4	1	1	0	1

What are the values of the counters (in binary) for all pages after these four clock ticks? If a page needed to be removed at that point, which page would be chosen for removal?

4. (5 points) A computer at Acme Company used as a compute server (i.e., to run non-interactive jobs) is observed to be running slowly (turnaround times longer than expected). The system uses demand paging, and there is a separate disk used exclusively for paging. The sysadmins are puzzled by the poor performance, so they decide to monitor the system. It is discovered that the CPU is in use about 20% of the time, the paging disk is in use about 98% of the time, and other disks are in use about 5% of the time. For each of the following, say whether it would be likely to increase CPU utilization (i.e., the percentage of time the CPU is in use) and why.
- (a) Installing a faster CPU.
  - (b) Installing a larger paging disk.
  - (c) Increasing the number of processes (degree of multiprogramming).
  - (d) Decreasing the number of processes (degree of multiprogramming).
  - (e) Installing more main memory.
  - (f) Installing a faster paging disk.

### 3 Programming Problems

Do the following programming problems. You will end up with at least one code file per problem. Submit your program source (and any other needed files) by sending mail to [bmassing@cs.trinity.edu](mailto:bmassing@cs.trinity.edu), with each file as an attachment. Please use a subject line that mentions the course number and the assignment (e.g., “csci 4320 homework 5”). You can develop your programs on any system that provides the needed functionality, but I will test them on one of the department’s Fedora Linux machines, so you should probably make sure they work in that environment before turning them in.

1. (10 points) Write a program or programs to demonstrate the phenomenon described in problem 2. Turn in your program(s) and output showing differences in execution time. (It’s probably simplest to just put this output in a text file and send that together with your source code file(s).) Try to do this in a way that shows a non-trivial difference in execution time (so you will likely need to make the arrays or other data structures large). I’d prefer programs in C, C++, or Java, but anything that can be compiled and executed on one of the Fedora lab machines is fine, as long as you tell me how to compile and execute what you turn in, if it’s not C/C++ or Java. You don’t have to develop and run your programs on one of the lab machines, but if you don’t, (1) tell me what system you used instead, and (2) be sure your programs at least compile and run on one of the lab machines, even if they don’t necessarily give the same timing results as on the system you used.

Possibly helpful hints:

- An easy way to measure how long program `mypgm` takes on a Linux system is to run it by typing `time mypgm`. Another way is to run it with `/usr/bin/time mypgm`. (This gives more/different information — try it.) If you’d rather put something in the program itself to collect and print timing information, for C/C++ programs you could use the

function in `timer.h`<sup>1</sup> to obtain starting and ending times for the section of the code you want to time, or for Java programs you could use `System.currentTimeMillis`.

- Your program doesn't have to use a 2D array (you might be able to think of some other data structure that produces the same result). If you do use a 2D array, though, keep in mind the following:
    - To the best of my knowledge, C and C++ allocate local variables on the stack, which may be limited in size. Dynamically allocated variables (i.e., those allocated with `malloc` or `new`) aren't subject to this limit.
    - Dynamic allocation of 2D arrays in C is full of pitfalls. It may be easier to just allocate a 1D array and fake accessing it as a 2D array (e.g., the element in `x[i][j]`, if `x` is a 2D array, is at offset `i*ncols+j`).
2. (10 points) The starting point for this problem is a Java program that simulates execution of a page replacement algorithm. Currently the program simulates only the FIFO algorithm. Your mission is to make it simulate one or more of the other algorithms mentioned in the text (and listed in commented-out code in the main program). You will get full credit for simulating one algorithm, extra points for simulating additional algorithms. The program consists of several classes collected in a package called `pagingsimulator`.
- [API documentation](#)<sup>2</sup>.
  - [Code](#)<sup>3</sup>. (Class `PageReplacerTest` contains a `main` method to test all algorithms.)
  - [Sample input](#)<sup>4</sup>. ([Output from my sample solution](#)<sup>5</sup>.)
  - [Another sample input](#)<sup>6</sup>. ([Output from my sample solution](#)<sup>7</sup>.)

Feel free to rewrite anything about this program, including starting over in a language of your choice. Just remember that the program has to run on one of the department Linux machines, and it needs to accept input from command-line arguments and files — no GUIs, Web-based programs, etc. The latter requirement is to make it easier for me to automate testing your code. If you make changes to the format of the input — and I prefer that you don't — change the comments so they describe the changed requirements.

Make the following assumptions:

- Initially memory is empty.
- All memory references are valid — if the page is not in memory, it can be read in from disk. (You don't have to simulate that part, just count how often it happens.)

---

<sup>1</sup>[http://www.cs.trinity.edu/~bmassing/Classes/CS4320\\_2010fall/Homeworks/HW04/Problems/timer.h](http://www.cs.trinity.edu/~bmassing/Classes/CS4320_2010fall/Homeworks/HW04/Problems/timer.h)

<sup>2</sup>[http://www.cs.trinity.edu/~bmassing/Classes/CS4320\\_2010fall/Homeworks/HW04/Problems/pagingsimulator/docs/](http://www.cs.trinity.edu/~bmassing/Classes/CS4320_2010fall/Homeworks/HW04/Problems/pagingsimulator/docs/)

<sup>3</sup>[http://www.cs.trinity.edu/~bmassing/Classes/CS4320\\_2010fall/Homeworks/HW04/Problems/pagingsimulator/source/pagingsimulator/](http://www.cs.trinity.edu/~bmassing/Classes/CS4320_2010fall/Homeworks/HW04/Problems/pagingsimulator/source/pagingsimulator/)

<sup>4</sup>[http://www.cs.trinity.edu/~bmassing/Classes/CS4320\\_2010fall/Homeworks/HW04/Problems/pagingsimulator/sample.in](http://www.cs.trinity.edu/~bmassing/Classes/CS4320_2010fall/Homeworks/HW04/Problems/pagingsimulator/sample.in)

<sup>5</sup>[http://www.cs.trinity.edu/~bmassing/Classes/CS4320\\_2010fall/Homeworks/HW04/Problems/pagingsimulator/sample.out](http://www.cs.trinity.edu/~bmassing/Classes/CS4320_2010fall/Homeworks/HW04/Problems/pagingsimulator/sample.out)

<sup>6</sup>[http://www.cs.trinity.edu/~bmassing/Classes/CS4320\\_2010fall/Homeworks/HW04/Problems/pagingsimulator/another.in](http://www.cs.trinity.edu/~bmassing/Classes/CS4320_2010fall/Homeworks/HW04/Problems/pagingsimulator/another.in)

<sup>7</sup>[http://www.cs.trinity.edu/~bmassing/Classes/CS4320\\_2010fall/Homeworks/HW04/Problems/pagingsimulator/another.out](http://www.cs.trinity.edu/~bmassing/Classes/CS4320_2010fall/Homeworks/HW04/Problems/pagingsimulator/another.out)