

# CSCI 4320 (Principles of Operating Systems), Fall 2007

## Homework 5

**Assigned:** November 9, 2007.

**Due:** November 19, 2007, at 5pm.

**Credit:** 20 points.

### 1 Reading

Be sure you have read Chapter 5.

### 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.

1. (5 points) Consider the following two I/O devices. For each device, say whether you think programmed I/O or interrupt-driven I/O makes the most sense, and justify your answer. (*Hint:* Consider the time required for interrupt processing versus the time needed for the actual input/output operation.)
  - (a) A printer that prints at a maximum rate of 400 characters per second, connected to a computer system in which writing to the printer's output register takes essentially no time, and each character printed requires an interrupt that takes a total of 50 microseconds (i.e.,  $50 \times 10^{-6}$  seconds) to process.
  - (b) A memory-mapped video terminal, connected to a system where interrupts take a minimum of 100 nsec to process and copying a byte into the terminal's video RAM takes 10 nsec.
2. (5 points) The textbook divides the many routines that make up an operating system's I/O software into four layers, as shown in Figure 5-10. In which of these layers should each of the following be done? Why? (Assume that in general functionality should be provided at the highest level at which it makes sense — e.g., in user-level software rather than device-independent software.)
  - (a) Converting floating-point numbers to ASCII for printing.
  - (b) Computing the track, sector, and head for a disk read operation.
  - (c) Writing commands to a printer controller's device registers.
  - (d) Detecting that an application program is attempting to write data from an invalid buffer address. (Assume that detecting an invalid buffer address can only be done in supervisor mode.)

3. (5 points) Suppose at a given point in time a disk driver has in its queue requests to read cylinders 10, 22, 20, 2, 40, 6, and 38, received in that order. If a seek takes 5 milliseconds (i.e.,  $5 \times 10^{-3}$  seconds) per cylinder moved, and the arm is initially at cylinder 20, how much seek time is needed to process these requests using each of the three scheduling algorithms discussed (FCFS, SSF, and elevator)? Assume that no other requests arrive while these are being processed and that for the elevator algorithm the initial direction of movement is outward (toward larger cylinder numbers).
4. (5 points) Student H. Hacker installs a new disk driver that its author claims improves performance by using the elevator algorithm and also processing requests for multiple sectors within a cylinder in sector order. Hacker, very impressed with this claim, writes a program to test the new driver's performance by reading 10,000 blocks spread randomly across the disk. The observed performance, however, is no better than what would be expected if the driver used a first-come first-served algorithm. Why? What would be a better test of whether the new driver is faster? (*Hint*: The test program reads the blocks one at a time. Think about how many requests will be on the disk driver's queue at any one time.)

### 3 Programming Problems

(Optional) Do the following programming 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"). As always, 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 7 Linux machines, so you should probably make sure they work in that environment before turning them in.

1. (Up to 5 extra-credit points). The Linux lab machines have special files `/dev/random` and `/dev/urandom` that generate sequences of "random" bytes. (Read the man page for `urandom` for an explanation of the difference between them.) Write a program that compares the results of generating  $N$  integers using one of these special files to the results of generating  $N$  integers using function `rand()`. (It's up to you to decide how to compare them. A simple test might be to count how many are even and how many are odd. You may have a better idea!) Submit your source code and a text file containing output of one or more executions.

*Hints:*

- I recommend getting  $N$  as a command-line argument, so you can experiment with different values. If you don't remember how to access command-line arguments in C, here is a simple example:

```
#include <stdio.h>
#include <stdlib.h>
int main(int argc, char *argv[]) {
    int n = -1;
    if ((argc < 2) || ((n = atoi(argv[1])) <= 0)) {
        fprintf(stderr, "usage: %s n (positive number)\n", argv[0]);
        return EXIT_FAILURE;
    }
    printf("%d\n", n);
    return EXIT_SUCCESS;
}
```

- You will probably need to use `open` and `read` rather than `fopen` and `fscanf` to read from the special file. `man` pages for these two functions can be found via `man 2 open` and `man 2 read`.