Final Examination

COP 4600 Operating Systems
May 2, 1997

Instructions

1. **Read and follow all instructions.** Failure to do so will result in penalty.
2. **Print your last name in the upper right corner of this page.**
3. This is a closed-book examination.
4. You are permitted three 8.5 by 11 inch sheets of notes, both sides, that you have prepared.
5. **Answer any four (4) questions**. All questions are of equal value. Points for each part of each question are given in parentheses.
6. **Leave sufficient room in the upper lefthand corner for the staple.**
7. **Use no more than two pages** of paper (both sides OK) to hold the answer to each question.
8. **Write legibly.**
9. **Start** the answer to each question on a **new page** (i.e., do **not** put the answer to more than one question on the same page).
10. Assemble your answers in **numerical order** of the questions when you submit them.
11. **Read and sign the following statement.** You may tear this page off the front of your exam and attach it to the front of your answers if you wish to take the questions home with you today.

On my honor, I have neither given nor received unauthorized aid on this examination.
Signed:

1. In this question, you will be asked to compare various types of operating system by comparing certain components. Your answers should contain not only a description of issues and methods used for each component of each system type, but reasons why these are relevant to that system (and not others, if that is the case). The components of interest are: process management, I/O management, memory management, file systems, and security. Describe and compare how these components are approached in the following types of systems.

   (a) (5) First generation systems.

   (b) (5) Uniprogrammed batch systems.

   (c) (5) Multiprogrammed batch systems.

   (d) (5) Time-sharing systems.

   (e) (5) Personal computers

2. (a) (8) Describe the boot process: all the steps that occur from when a computer is first turned on until the first process is entered into the process table for a typical OS. Include the locations of these instructions at each stage.

   (b) (4) What is a process tree and why is it important?

   (c) (3) Classify the following scheduling algorithms as preemptive or non-preemptive: SJF, SRTF, RR(1), FCFS.

   (d) (5) In what sense does an appropriately configured multi-level feedback queue approximate SRTF? In what ways does it differ (in terms of the effect on scheduling)?

   (e) (5) Describe exponential average method of predicting the length of the next CPU burst of a process. For what types of scheduling is this useful? Given an initial estimate of 5 and observed burst lengths of 5, 9, 9, 4, 12, 10, and 3, what is the next predicted burst length if alpha=0.5?

3. (a) (8) Describe a typical overlay system. Be sure to include each type of component and to describe what it does and how it works with the rest of the system. Why were overlays used? Why are they rarely used any more?

   (b) (8) Draw a diagram of the hardware used to support a pure segmentation system; list and describe the steps (corresponding to labels on your diagram) taken from the time a process generates a logical address until the corresponding physical address is asserted on the system bus (or something else occurs).

   (c) (4) Suppose that the physical memory access time in a system is 50 microseconds, that average disk access time is 20 milliseconds, and that the page table is always in memory. What is the effective access time for this system if the page fault rate is .01? What if the PFR is .005? Show your work.

   (d) (5) Consider the same system as in the previous part, except that the page table is not always completely in memory and the system has a translation lookaside buffer. Compute the effective access times for a PFR of .01 given TLB hit rates of .9 and .95 (assume that the TLB access is instantaneous). Show your work.

4. (a) (5) What is the working set model? How is it used? How may the working set be determined from a process trace?

   (b) (5) Describe and compare local and global page replacement. How do these relate to the working set model?

   (c) (5) How does the CLOCK algorithm attempt to approximate the working set? When does it fail? Give examples.

   (d) (5) Why is LRU considered impractical? Describe the reference vector LRU approximation method and indicate why this is considered more practical.

   (e) (5) What is the page fault frequency based approach to allocation, and how does it relate to the working set model? Use diagrams as necessary to describe the PFF approach and state how the method works.

5. (a) (6) What are cylinder groups and how are they used in file systems? Why are they used?

   (b) (6) How does an operating system know what is on a disk? Use the Unix or the Minix operating systems as an example. Be specific in describing the structure and types of information contained.

   (c) (8) Describe and compare the DOS FAT approach for file structure to the inode approach used in Unix or Minix. What are the advantages or disadvantages of each approach?

   (d) (5) Describe and compare two distinct ways of managing free disk blocks.

6. (a) (6) What is the access control matrix? Describe its structure and function.

   (b) (4) What are identification and authentication and how do they relate to the ACM?

   (c) (5) What are access control lists, how are they used, and why?

   (d) (5) What are capability lists, how are they used, and why?

   (e) (5) Describe how Unix combines ACLs and CLs in file access.