CIS 25 Final Exam

1. What are the four necessary conditions for a deadlock?
2. What three criteria must a critical section solution satisfy?
3. Once you found a deadlock, state at least two ways of handling it? (you already have a deadlock, it's too late to try to avoid it).
4. Under what circumstances do page faults occur? Describe the actions taken by the operating system when a page fault occurs. (q10.1)
5. How does the operating system communicate with a device controller? (a hard-drive controller for example) Describe the roles played by interrupts, I/O ports, and DMA.
6. Describe how non-blocking I/O is implemented. How a process initiates it, who controls the I/O while it is happening, and how the process knows when the I/O has completed. Typically, in which parts of the process is the CPU involved, and why?
7. Name two differences between logical and physical addresses. Why are page sizes always power of 2?
8. What are the five major activities of an operating system in regard to process management? (q3.1).
9. MS-DOS provided no means of concurrent processing. Discuss three major complications that concurrent processing adds to an operating system. (q4.1)
10. Define the difference between preemptive and non-preemptive scheduling.
11. What is the meaning of the term *busy waiting*? What other kinds of waiting are there in an operating system? Can busy waiting be avoided altogether? Explain your answer. (q7.1)
12. In distributed systems context, what is the difference between vertical scalability and horizontal scalability? What are the benefits of each.
13. What is the difference between authentication and authorization? Why do you need them? List three different ways of authenticating someone.
14. Show that if the `wait` and `signal` operations are not executed atomically, then mutual exclusion may be violated. (q7.7)
15. (15 points) Consider the following page-reference string: 6,3,2,1,2,3,6,7,3,2,1,2,6,5,1,2,4,3,2,1. How many page faults could occur for the following page replacement algorithms, assuming one, four, and seven frames. Remember that all frames are initially empty, so your first unique pages will all cost one fault each (q10.11).

    (a) LRU replacement
    (b) FIFO replacement
    (c) Optimal replacement

16. (15 points) Five batch jobs, A through E, arrive at a computer center at almost the same time (but in order A through E). They have estimated running times of 7, 4, 12, 11 and 8 minutes. For each of the following scheduling algorithms, draw a Gantt chart showing the order in which the jobs will run. Then, for each algorithm, determine the average waiting time. Ignore process switching overhead. All jobs are completely CPU bound.

    (a) Round Robin (time slice = 3)
    (b) First-Come, First-Served
    (c) Shortest Job First