

Due in class on Friday 30 November 2001

1. (Exercise 8.3 from OSC) People have said that proper spooling would eliminate deadlocks. Certainly, it eliminates contention card readers, plotters, printers, and so on. It is even possible to spool tapes (called staging them), which would leave the resources of CPU time, memory, and disk space. Is it possible to have a deadlock involving these resources? If it is, how could such a deadlock occur? If it is not, why not? What deadlock scheme would seem best to eliminate these deadlocks (if any are possible), or what condition is violated (if they are not possible)?

For there to be deadlock involving CPU time, memory, and disk space, the four conditions for deadlock would have to hold for those resources. If the CPU scheduler is preemptive, then deadlock can not occur for CPU time. Since memory is usually considered preemptable, deadlock over memory space does not generally occur. However, a non-preemptive CPU scheduler can deadlock, as can a memory manager that does not support swapping. Deadlock over disk space is more likely, as a disk can easily fill up, and is usually not prevented as the overhead for preemption is too high.

2. Consider the following page reference string: **b c b a e b d e c b e b**. If there are 3 frames in physical memory, how many pages faults will occur for (a) FIFO replacement, and (b) LRU replacement. Draw a diagram for each to show your work.

For three frames, FIFO = 9 page faults

B	C	B	A	E	B	D	E	C	B	E	B
B		-		E			-	C			
	C				B				-	E	
			A			D					B

For three frames, LRU = 7 page faults

B	C	B	A	E	B	D	E	C	B	E	B
B		-			-			C			
	C			E			-			-	
			A			D			B		-

3. (Exercise 11.9 from OSC) Give an example of an application in which data in a file should be accessed in the following order (but give at least 3 examples for part (a)).
- (a) Sequentially: text editor reading in a file, compiler reading source code, web browser reading a web page, media player playing music or a video, burning a file onto a CDROM.
 - (b) Randomly: database program reading part of a database file, reading one of a set of files packed together (e.g., one file in a set of ZIP files), reading one work out of a dictionary file.

- 4. My numbers for a “typical disk today” in Lecture 31 are somewhat outdated. Go online, find a good disk for about \$400, and then find similar specifications for that disk. You probably won’t find all the numbers reported exactly the same way, but at a minimum I want to know the number of platters, capacity in GB, and some data transfer numbers. Include a URL to document your information. (NOTE: this question counts double.)**

Most of the class seemed to find the Western Digital Special Edition “Caviar” 200 GB hard drive, \$320 to \$400 from CDW, Best Buy, or some other electronics / computer supplier. This drive has 3 platters (6 physical heads), formatted capacity 200,049 MB (approximately 200 GB), 7200 RPM rotational speed, 8.9 ms average seek time, and 100 MB/s data transfer speed (Mode 5 Ultra ATA).