

Homework 9*Reading: Tanenbaum, Sections 5.1-5.5, Skim Chapter 6**Due: None***Problem 1** (4 Points)

Complete the course evaluation for CSE 422S (<http://evals.wustl.edu>).

Problem 2 (0 Points)

- Suppose that the the disk queue has the following cylinder requests: 27, 129, 110, 186, 147, 41, 10, 64, 120. Compute the average seek distance for the scheduling policies FIFO, SSTF, LOOK, and C-LOOK assuming that the seek arm is initially positioned over cylinder 100 and is moving in the direction of *decreasing* cylinder numbers.
- Sketch the seek trajectory for the C-LOOK policy. The horizontal axis should be cylinder numbers (increasing to the right). The vertical axis should be time (increasing downward).
- How can seek scheduling still be used if the real disk geometry is different than the virtual disk geometry?

Problem 3 (0 Points)

Consider a basic (doesn't use zoned recording) disk drive that has 512-byte sectors, 96 sectors/track, 110 tracks per surface, and 8 usable surfaces.

- What is the raw capacity of the disk drive?
- Suppose that logical records are 120 bytes long and are blocked 10 per physical record. What is the capacity of the disk drive in number of records if physical records can not span across tracks or cylinders, but can span sectors? Ignore file headers, track indexes, and other management overhead in your calculation.
- What is the volume (in megabytes) and percentage of fragmentation due to the no-span limitations in Part b?
- On average, how much time will it take to read every record (from Part a) on the disk if the seek time is 0, and every read spends an average amount of rotational latency time (assuming uniform rotational delay)? Assume that the disk rotates at 3600 RPM.
- Suppose that the system is dedicated to handling just our I/O so that there is no interference from other users. Furthermore, assume that the total software overhead (e.g., system call time, interrupt handling time) is 200 μ sec. How much time will it take this system to read every record (from Part a) on the disk if the seek time is 0?

Problem 4 (0 Points)

Modern RAID storage systems are known for their resilience to single drive failures, high performance (high effective data rate and/or high transaction rate), and flexibility.

- a) How does RAID-5 technology handle a single drive failure?
- b) A typical RAID subsystem consists of multiple drives sitting on a SCSI bus. Although, only one block can be transferred over the bus at one time, the other disks can be concurrently seeking or transferring data to their internal disk cache. What average aggregate rate must this bus be able to handle? Assume that there are 15 drives that rotate at 7200 RPM; strips are 8 KB; the formatted track capacity is 64 KB; and the typical seek time is negligible.