DM510 Operating Systems
Spring 2026

Exercise Sheet 7

Complete before tutorial on Thursday, April 9th

Learning goals

Be able to

- explain the functioning, advantages, and disadvantages of HDDs, NVMs, and tape devices.
- explain disk scheduling algorithms: SSTF, FCFS, SCAN, C-SCAN
- explain the functioning of different RAID organizations and their effect on performance and reliability

Chapter 11 (types of storage)

Exercise 1. Compare HDDs, NVMs, and tapes as storage devices and discuss their application areas, advantages and disadvantages.

It is sometimes said that tape is a sequential-access medium, whereas a hard disk is a random-access medium. In fact, the suitability of a storage device for random access depends on the transfer size. The term streaming transfer rate describes the rate for a data transfer that is underway, excluding the effect of access latency. In contrast, the effective transfer rate is the ratio of total bytes to total seconds, including overhead time such as access latency.
For the following exercises, consider a computer with the following characteristics:

- the level-2 cache has an access latency of 8 nanoseconds and a streaming transfer rate of 800 megabytes per second,
- the main memory has an access latency of 60 nanoseconds and a streaming transfer rate of 80 megabytes per second,
- the hard disk has an access latency of 15 milliseconds and a streaming transfer rate of 5 megabytes per second,
- and a tape drive has an access latency of 60 seconds and a streaming transfer rate of 2 megabytes per second.

Exercise 2. Random access causes the effective transfer rate of a device to decrease, because no data are transferred during the access time. For the hard disk described above, what is the effective transfer rate if an average access is followed by a streaming transfer of:

a. 512 bytes

b. 8 kilobytes

c. 1 megabyte

d. 16 megabytes

The utilization of a device is the ratio of effective transfer rate to streaming transfer rate.

Exercise 3. Calculate the utilization of the disk drive for each of the four transfer sizes given before.

Exercise 4. Suppose that a utilization of 25 percent (or higher) is considered acceptable. Compute the minimum transfer sizes that give acceptable utilization for cache, memory, and tape.

Chapter 11 (disk scheduling)

Exercise 5. Explain why SSTF (Shortest Seek Time First) scheduling tends to favor middle cylinders over the innermost and outermost cylinders.

Exercise 6. None of the disk scheduling algorithms, except FCFS, is truly fair (starvation may occur).

a. Explain why this is true.

b. Describe a way to modify SCAN to ensure fairness.

Exercise 7. Suppose that a disk drive has 5000 cylinders, numbered 0 to 4999. The drive is currently serving a request at cylinder 2150, and the previous request was at cylinder 1805. The queue of pending requests, in FIFO order, is:

$$2069, 1212, 2296, 2800, 544, 1618, 356, 1523, 4965, 3681$$

Starting from the current head position, what is the total distance (in cylinders) that the disk arm moves to satisfy all the pending requests for each of the following disk-scheduling algorithms?

a. FCFS

b. SCAN

c. C-SCAN

Chapter 11 (RAID)

Exercise 8. RAID can be used to protect against data loss. Give two examples where data loss may happen despite using RAID

Exercise 9. Discuss whether it would be sensible to use RAID on several partitions of the same hard disk.

Exercise 10. Could a RAID level 1 organization achieve better performance for read requests than a RAID level 0 organization? If so, how? What about write requests?

Exercise 11. The following picture shows two RAID configurations: Configuration 1 uses RAID 1 (mirrored disks), where Disk 3 mirrors Disk 1 and Disk 4 mirrors Disk 2. Configuration 2 uses RAID 4 (block-interleaved parity), where Disk 3 is the parity disk.

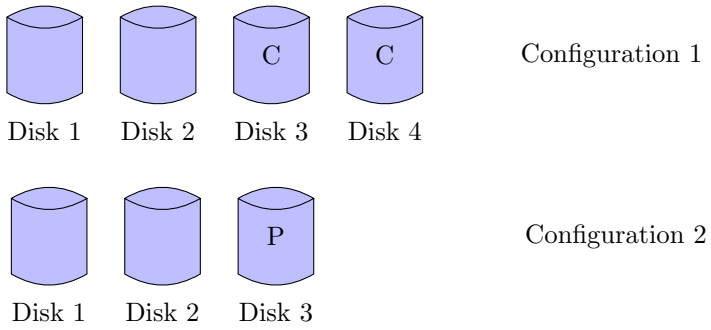

For each configuration: list all combinations of disks that would lead to data loss if they fail simultaneously.

Exercise 12. Consider a RAID level 5 organization comprising five disks, with the parity for sets of four blocks on four disks stored on the fifth disk. How many blocks are accessed in order to perform a write of seven continuous blocks of data?

Exercise 13. Assume that you have a mixed configuration comprising disks organized as RAID level 1 and RAID level 5 disks. Assume that the system has flexibility in deciding which disk organization to use for storing a particular file. Which files should be stored in the RAID level 1 disks and which in the RAID level 5 disks in order to optimize performance?

Exercise 14. The reliability of a storage device is typically described in terms of mean time between failures (MTBF). If a system contains 1000 disk drives, each of which has a 750000 hour MTBF, which of the following best describes how often a drive failure will occur in that disk farm: once per thousand years, once per century, once per decade, once per year, once per month, once per week, once per day, once per hour, once per minute, or once per second?