The Storage Layer

Exercise 9.5 Consider a disk with a sector size of 512 bytes, 2000 tracks per surface, 50 sectors per track, five double-sided platters, and average seek time of 10 msec.

1. What is the capacity of a track in bytes? What is the capacity of each surface? What is the capacity of the disk?
2. How many cylinders does the disk have?
3. Give examples of valid block sizes. Is 256 bytes a valid block size? 2048? 51200?
4. If the disk platters rotate at 5400 rpm (revolutions per minute), what is the maximum rotational delay?
    a. What is the average rotational delay?
5. If one track of data can be transferred per revolution, what is the transfer rate?

Exercise 9.12 What is *sequential flooding* of the buffer pool?

Exercise 9.14 Explain the term prefetching. Why is it important?

Solutions

Answer 9.5

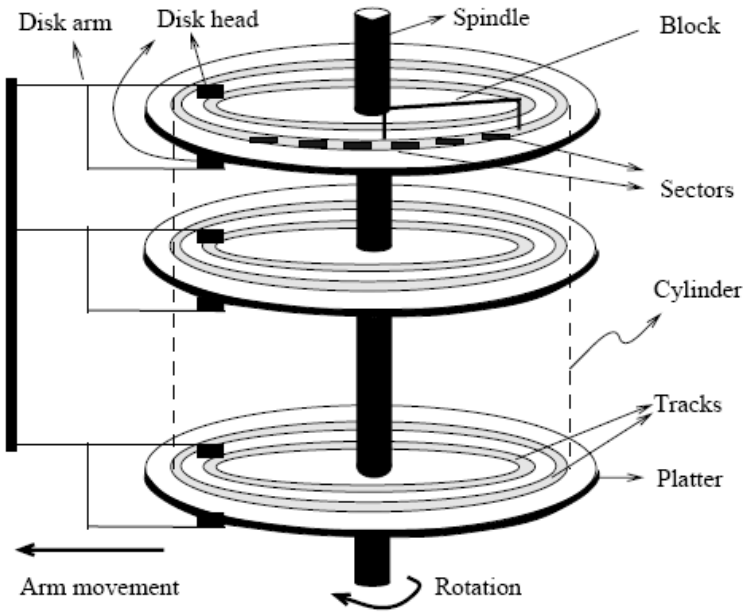

1. bytes/track = bytes/sector * sectors/track = 512 * 50 = 25K
   bytes/surface = bytes/track * tracks/surface = 25K * 2000 = 50, 000K
   bytes/disk = bytes/surface * surfaces/disk = 50, 000K * 5 * 2 = 500, 000K
2. The number of cylinders is the same as the number of tracks on each platter, which is 2000.
3. The block size should be a multiple of the sector size. We can see that 256 is not a valid block size while 2048 is. 51200 is not a valid block size in this case because block size cannot exceed the size of a track, which is 25600 bytes.
4. If the disk platters rotate at 5400rpm, the time required for one complete rotation, which is the maximum rotational delay, is:
   (1/5400) * 60 = 0.011 seconds.
   a. The average rotational delay is half of the rotation time, 0.0055 seconds.
5. The capacity of a track is 25K bytes. Since one track of data can be transferred per revolution, the data transfer rate is:
   25K/0.011 = 2, 250Kbytes/second

Answer 9.12 Some database operations (e.g., certain implementations of the join relational algebra operator) require repeated sequential scans of a relation. Suppose that there are 10 frames available in the buffer pool, and the file to be scanned has 11 or more pages (i.e., at least one more than the number of available pages in the buffer pool). Using LRU, every scan of the file will result in reading in every page of the file! In this situation, called *"sequential flooding"*, LRU is the worst possible replacement strategy.

Answer 9.14 Because most page references in a DBMS environment are with a known reference pattern, the buffer manager can anticipate the next several page requests and fetch the corresponding pages into memory before the pages are requested. This is prefetching.

Benefits include the following:

1. The pages are available in the buffer pool when they are requested.
2. Reading in a contiguous block of pages is much faster than reading the same pages at different times in response to distinct requests.