

Exercise 1

Released: April 9, 2018 · Discussion: April 16, 2018

1 Storage Disks

The lecture discussed the characteristics of hard disks and solid state disks.

1. Describe the functioning of the two storage media.
2. What are the advantages and disadvantages respectively?
3. What are the resultant consequences for a DBMS?
4. Given a hard disk with blocks of 512 bytes in size and following characteristics: read transfer rate of 50 MB/s; write transfer rate of 40 MB/s; seek time 10 ms; rotational speed: 7200 rpm.

Additionally a NAND flash drive is given with 128 KB sized blocks of 64 pages each. The access times are as follows: 25 μ s to begin a read operation; 100 μ s to read a page; 2 ms to delete a block and 200 μ s to write a page.

How long do the following operations take?

- (a) Sequentially reading 1 GB.
- (b) Reading 1 GB where all data is split into randomly distributed blocks of 4 KB.
- (c) Sequentially writing 10 MB.
- (d) Writing a single byte.

2 Query processing

- Consider two plans p_1 and p_2 for a database query. The query shall be processed by a DBMS with 4 KB pages.

p_1 reads 100000 pages sequentially and reads 2500 records with a length of 512 byte from positions equally distributed among a large number of pages.

p_2 reads 250000 pages sequentially.

- Calculate the time needed to process the queries p_1 and p_2 for a hard disk and solid state disk using the parameters from the previous task respectively. It is sufficient to determine the access times only. Data buffering does not need to be considered. Discuss the influence of different access patterns and storage media on the processing of queries.