

Storage-Take Your Time



Credits to:

Katariina Martikainen

Giorgos Tagkoulis

KTH Kista

Stockholm

03.12.2019



File Operations

- Open: `int open(const char *pathname, int flags, mode_t mode);`
 - Returns a file description

- Read: `ssize_t read(int fd, void *buf, size_t count);`
 - Returns the number of bytes read

- Write: `ssize_t write(int fd, const void *buf, size_t count);`
 - Returns the number of bytes written is returned



lseek

read and write at user specified location

`lseek(fd, pos, SEEK_SET);` ← read/write at absolute value given by pos

`lseek (fd , step , SEEK_CUR);` ←read/write at current place + value given by step



Page Cache

- When a file is read from disk, the operating system will place the read blocks in what is called the page cache.
- If we read from the file once it is very likely that we will read from it again so we keep a copy of the blocks that we read in a cache.
- Improves the time performance of a read operation within nearby addresses



Memory speed

- Accessing RAM is much faster than accessing storage on hard-disk drive or solid-state drive
- Solid-state drive is faster than hard-disk because it does not have physically moving elements.

Exam Questions





Exam Question 1

9.6 read performace [2 points]

Assume that we have an ordinary hard drive that is attached over a SATA connection of 6 Gb/s that gives us a read time of a random 4 KByte block of 12 ms. How is this changed if we replace the disk to one that is attached by a SAS connection of 12 Gb/s? Motivate.

Answer:



9.6 read performace [2 points]

Assume that we have an ordinary hard drive that is attached over a SATA connection of 6 Gb/s that gives us a read time of a random 4 KByte block of 12 ms. How is this changed if we replace the disk to one that is attached by a SAS connection of 12 Gb/s? Motivate.

Answer:

The time it takes to read a 4KByte block is dominated by the time to it takes to move the arm and rotate the disc to the correct position to read from. The new hard drive can have better values for these features, but the fact that the connection increases in capacity gives a marginal difference for individual blocks.



Exam Question 2

7.1 a regular HDD [2 points]

If a hard disk drive has a average seek time of 10 ms, a rotation speed of 7200 rpm (rounds per minute) and a read performance of 200 MiB/s. Then what is the average time to read a random sector on 4KiB

Answer:



7.1 a regular HDD [2 points]

If a hard disk drive has a average seek time of 10 ms, a rotation speed of 7200 rpm (rounds per minute) and a read performance of 200 MiB/s. Then what is the average time to read a random sector on 4KiB

Answer:

Approx. 14 ms. If it takes 10 ms to position the arm, and in average 4 ms (half the time of $60/7200$ ms to rotate one turn) then we will find the sector in 14 ms. Reading a sector of 4KiB is negligible in this context, when we have positioned the head, the reading itself takes 20 microseconds.



Exam Question 3

7.1 rpm [2 points]

There are hard drives with different performance, one thing that differ is the rotation speed. Why is the rotation speed important, what is improved and what could possibly decrease?

7.1 rpm [2 points]

There are hard drives with different performance, one thing that differ is the rotation speed. Why is the rotation speed important, what is improved and what could possibly decrease?

Answer: The so-called the access time, the time it takes for the disk to spin the right sector, decreases with the rotation speed. The reading speed also increases because we can move across more sectors at the same time.

When the speed of rotation increases, it may be that limitations in the reading head mean that we cannot store information equally tightly on your hard drive. This would mean that we have fewer sectors on a disk and overall a smaller amount of data. The disk becomes probably even more expensive because better precision in mechanics is required.