

Uni Freiburg, Web Science Group
Prof. Peter Fischer
Systems Infrastructure for Data Science - Winter 2012/13

Exercise Sheet #1: Data Storage

Due/discussion: Nov 2, 2012

Exercise 1.1 : Hard Disks and SSDs

Suppose you have implemented a rudimentary record store that stores fixed-size records, 64 bytes each, in a single flat file (disk-based array of records). A record ID is just an offset into the file. The file resides on a regular hard-disk with the following characteristics:

Average Seek Time (T_{seek}) = 10 ms
Rotational Speed = 7200 RPM
Average Transfer Rate = 50 MB/s
Blocksize = 8 KB

Assume the file is 100 MB big and the application requests to read $x = 1,000$ random records. The operating system's file caching is disabled. You should ignore the track to track seek time.

- A. Which is the better strategy: i) for every single record the application asks for, probe the file at the appropriate offset; or ii) perform a sequential scan of the entire file.
- B. What is the fraction of records requested compared to the capacity of the file, above which scanning the entire file becomes cheaper (faster) than random probing? How does doubling the transfer rate to 100 MB/s change the threshold?
- C. Assume the records in which the application is interested in are not truly random, but are in fact heavily skewed (clustered around a specific portion of the file). Can you think of an alternative to random probing or a full sequential scan that does not involve heavy caching?
- D. Now assume that we are using an Solid State Disk (SSD) instead of a mechanical HDD with the following cost calculation model:

$t = t_{setup} + t_{transfer}$ where $t_{setup} = 0.5ms$ and Average Transfer Rate = 100 MB/s for our SSD.

Under these assumptions, consider the first part of question B again.

Exercise 1.2 : RAID

You are working for a big Swiss bank and considering the acquisition of a state-of-the-art disk array storage system. You have the following RAID levels to choose from:

1. RAID 0
2. RAID 1

3. RAID 5

Discuss which of these RAID levels are better in terms of: (a) read vs. write performance and (b) reliability.

Exercise 1.3 : Buffer Manager

In this exercise, you have three blocks on disk and a buffer which can hold two pages. Consider the following sequence of accesses to blocks:

1, 3, 2, 2, 1, 3, 2, 1, 2

- A. What is the hit ratio using LRU as the replacement strategy?
- B. How about when MRU is used instead?

Exercise 1.4 : Database Files

Consider the two internal organizations for heap files: using lists of pages and a directory of pages.

- A. Describe them briefly and explain the trade-offs. Which organization would you choose if your records are variable in length?
- B. Describe different methods for finding a page for a given new record. Discuss their complexities briefly.
- C. Can you suggest a single page format to implement both internal file organizations?