

Name	Signature	Student Id Nr	Points

Operating Systems, mini exam 6, 24.4.2017 (6p)

Write your answer on this exam paper in the space given. Please notice, that the exam paper is 2-sided.

Disk management

a) [2 p] Which problem in SCAN algorithm (elevator) is solved with C-SCAN? How does it work?

Which (C-SCAN) problem is solved with N-STEP-SCAN scheduler? How does it work?

Which (C-SCAN) problem is solved with Linux Anticipatory I/O Scheduler? How does it work?

Which (C-SCAN) problem is solved with Linux Deadline Scheduler? How does it work?

b) [1 p] In disk cache, the next block to be replaced can be chosen with Frequency Based Replacement (FBR) algorithm. The blocks in disk cache are kept in an LRU stack that has been split into two parts: New Section and Old Section. In a more advanced form the LRU stack has also a Middle Section. Which problem in the original algorithm is solved with this Middle Section, and how does the FBR algorithm now work?

File Management. File Staff has 10 000 records, and it is stored on a disk. Field Corp-ID has values in range 1 – 100 000 000. Field Union-ID has values in range 1 – 1 000 000. The records are large and they have also many other fields.

- c) [1.5 p] Assume that Staff is implemented as an indexed sequential file. What does this mean in practice? Why would have one selected specifically indexed sequential file structure for Staff? In which order are the records in file Staff? How many indexes are there? Where are the indexes located? What do the indexes contain? When do you use the indexes and how?

Assume here, that the indexes are not implemented as B-trees.

- d) [1.5 p] The records are accessed one at a time from the disk, but in practice most of the records are found in disk cache. Assume here that the average time to read one record is T.

How do you access a random record with given Corp-ID field (e.g., Corp-ID = 23 456 789)?

How long will it take (using average record read time T as your time unit)?

How do you access a random record with given Union-ID field (e.g., Union-ID = 456 789)?

How long will it take (using average record read time T as your time unit)?