

Name	Signature	Student Id Nr	Points

Operating Systems, miniexam 4, 27.2.2019 (12p)

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

a) [total 4 p] Disk Scheduling

[1 p] How does the SCAN (elevator) algorithm work and what data is it based on?

[1 p] Which problem in SCAN algorithm is solved by C-SCAN algorithm? How is it solved?

[1 p] Which problem in SCAN algorithm is solved by N-step-SCAN algorithm? How is it solved?

[1 p] Which problem in SCAN algorithm is solved by Linux Anticipatory I/O Scheduler? How is it solved?

b) [total 4 p] Disk Cache

[1 p] Which problem is solved with disk cache? How does the solution work in main principles?

[2 p] Frequency-Based Replacement –algorithm (FBR) is used to select which disk block in disk cache will be replaced next. How does the FBR basic solution work? What data is the solution based on?

[1 p] The FBR basic solution sometimes replaces blocks too early. Why does this happen?

TURN

c) [total 4 p] File Organization

[1 p] Give a practical example on file use case, where it would be best to organize the file as sequential file, and not indexed sequential file or indexed file. Explain.

[1 p] Give a practical example on file use case, where it would be best to organize the file as indexed file, and not sequential file or indexed sequential file. Explain.

[2 p] Indexed sequential file has 100 000 records. It has normal, not B-tree based, index with 10 000 entries, and the index is kept in memory all the time. Keys have values in range 0 - 9 999 999. The file is large and it is kept on disk. Draw a picture of this system.

How do you find the record with key 23 456 from disk?

How many records do you need to read from disk in average to find the record with given key?