Name	Signature	Student Id Nr	*Points*

Operating Systems, mini exam 6, 18.4.2016 (6p)

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

Disk scheduling algorithms

a) [2 p] Why is SCAN better than FIFO?

Why is C-SCAN (Circular SCAN) better than SCAN (elevator)?

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

b) [1 p] How does RAID 5 solution work in main principles?

How do you write a (relatively small) data item using RAID 5 system?

TURN

File Management. Example file MyFile has 100 000 records, and field Id has possible values 1-100 000 000. It has also other fields.

c) [2 p] In what kind of circumstances would it be specifically useful to use an indexed file for file MyFile?

Why would *pile*, *sequential file*, or *indexed sequential file* not be a good solutions in your example case above?

How many indexes are there? How large are the indexes?

d) [1 p] How do you locate a record using B-tree index?
Give a concrete example using file MyFile with Id=23 456 789.