Operating Systems (5 cr), Exam 12.9.2019

Write in each answer sheet course name, date, your name, signature and student id.

No calculators. Exam paper is 2-sided. For each question, it is sufficient to give a 1-2 page answer.
This is ordinary final (separate) exam and covers the whole course. Answer all questions 1-4.

1. [9 p] **Critical section**
    a. [3 p] What is the critical section problem?
    How does it differ from the usual synchronization problem between two processes (e.g., P and Q)?
    What happens, if critical section problem is not noticed or it is solved erroneously?

    b. [3 p] How is critical section problem solved with special machine instructions (Test-And-Set, Compare-And-Swap, etc.)?
    When can you not use this solution? Why?

    c. [3 p] How is critical section problem solved with semaphores?
    When can you not use this solution? Why?

2. [9 p] **File management**
    a. [3 p] Give an example on a situation (use pattern), where sequential file system would be the best way to organize a large file system.
    What is bad with this file organization model?
    How many indexes are there in your example, and how large are they?
    How is some given record (e.g., customer number 123-456-789) found in you example?

    b. [3 p] Give an example on a situation (use pattern), where indexed sequential file system would be the best way to organize a large file system.
    What is bad with this file organization model?
    How many indexes are there in your example, and how large are they?
    How is some given record (e.g., customer number 123-456-789) found in you example?
    Why would sequential file system not be suitable for your example?

    c. [3 p] Give an example on a situation (use pattern), where indexed file system would be the best way to organize a large file system.
    What is bad with this file organization model?
    How many indexes are there in your example, and how large are they?
    How is some given record (e.g., customer number 123-456-789) found in you example?
    Why would sequential file system or indexed sequential file system not be suitable for your example?

TURN

3. [9 p] **Virtual memory**
    a. [2 p] Page size is 4KB. How do you find the main memory address for referenced data (byte address 0x11223344) using 2-level virtual memory (VM), when TLB is not used.

    b. [1 p] What advantage is there with 2-level VM have as compared to 1-level VM?

    c. [3 p] Which VM-related problem does Translation Lookaside Buffer (TLB) solve, and how does the solution work in main principles?
    What data is there in TLB?
    How do you find the main memory address for referenced data (byte address 0x11223344) with using 2-level virtual memory (VM), when TLB is also used?

    d. [3 p] Explain the concept "trashing", relating to virtual memory systems.
    What causes trashing?
    What problems does thrashing cause to the system?
    Give one method to avoid trashing, and explain how the solution works in main principles?

4. [9 p] **Scheduling**
    a. [3 p] How does Shortest Process Next (SPN) work in main principles?
    What data is SPN based on and how do you get that data?
    When can you use SPN? When can you not use it?

    b. [3 p] How does Fair Share Scheduling (FS) work in main principles?
    What data is FS based on and how do you get that data?
    When can you use FS? When can you not use it?

    c. [3 p] How does Rate Monotonic Scheduling (RMS) work in main principles?
    What data is RMS based on and how do you get that data?
    When can you use RMS? When can you not use it?