

# TDIU11 Exam

Ahmed Rezine

2019-03-22 kl:14-18

## Admitted material

Dictionary from English to your chosen language.

## Jour

Ahmed Rezine (013-281938) visiting after about one hour.

## Instructions

- Fill in the exam wrapper and read the instructions on it before you start.  
**Read instructions and all assignments carefully and completely before you begin.**
- You may answer in either English or Swedish.
- State your interpretation of the question and all assumptions you make.
- Write clearly. Unreadable text will be ignored.
- Be precise in your statements. **Prove your point when possible.**  
Ambiguous formulations will lead to reduction of points.
- **Motivate clearly and in depth all statements and reasoning.**  
**Explain calculations and solution procedures.**
- The exam is 40 points and graded U, 3, 4, 5 (**preliminary** limits: 20p, 30p, 35p).  
Points are given for motivations, explanations, and reasoning.

## Definitions

Unless you are more specific, the correcting team will interpret the following terms as follow:

**memory** Volatile random access memory (DRAM), about 100ns access time.

**disk** Permanent storage, about 10ms access time.

**page** A fixed size region of virtual memory, possibly on disk.

**frame** A fixed size region of physical memory (DRAM).

**block** A data block located on disk.

## Problem 1 (12pt)

You are given the following processes:

Process	Arrival time	Execution time
$P_1$	0	8
$P_2$	1	4
$P_3$	2	1

Table 1: Processes and their arrival and execution times

- Assume started processes will execute until their burst time is finished (nonpreemptive scheduling). Draw the Gantt diagram and give the average waiting time:
  - for the FCFS scheduling algorithm. (2pt)
  - for the SJF scheduling algorithm. (2pt)
  - for the SJF scheduling algorithm if the two first processes are “forced” to wait until time unit 2 (i.e. arrival of  $P_3$ ). The forced wait time should be counted in the waiting time. (2pt)
  - Which, among the above three approaches, yields the least average waiting time? what disadvantages does it have? (2pt)
- Draw the Gantt diagram and give the average waiting time for the preemptive SJF scheduling algorithm. (2pt)
- Between tick 0 and tick 1 some process is running on the CPU (not the kernel). This process could continue until completion, but instead another process with higher priority might get scheduled instead. Observe the kernel was not running at the preemption time. Explain how the kernel manages to regain control of the CPU in order to schedule other (previously preempted) processes. (2pt)

## Problem 2 (10 p)

Assume 28 bits (logical blocks) pointers and a 4TiB (i.e.,  $2^{42}$  bytes) hard drive.

- Give the minimal size of individual logical blocks in order for the pointers to be able to access any logical block of the the hard drive (2pt)
- Assume 16KiB logical blocks. What is the size of the FAT table with 4 bytes per entry and corresponding to having a single 4TiB volume? Recall that FAT is a form of linked allocation but where all links are centralized in a single table for the whole volume. (2pt)
- Still assuming 4 bytes per FAT entry and a 4TiB volume, how large (in bytes) should a single logical block be in order to obtain a FAT table of 128 MiB (i.e.,  $2^{27}$  bytes) ? (2pt)
- Compare advantages and disadvantages of the two logical block sizes (i.e, the block sizes adopted/obtained in questions 2 and 3). (2pt)
- Give one advantage and one disadvantage of using FAT based allocation compared to indexed allocation. (2pt)

### Problem 3 (10pt)

Assume an architecture with 32 bits addresses (both physical and logical addresses are 32 bits). The architecture can be configured to use either 4 KiB ( $2^{12}$ B) or 1 MiB ( $2^{20}$ B) page sizes. Assume each page entry takes up 4 bytes.

1. Suppose we adopt 1MiB pages.
  - a) How many page table levels are needed? Explain. (2pt)
  - b) What information is cached in the TLB (translation look-aside buffer) of this system? Explain what for? (2pt)
  - c) Suppose access to main memory takes 10 nanoseconds. You can neglect TLB access time. what should the TLB hit ratio be in order to not suffer more than a 1% performance degradation due to address translation? Explain (2pt)
2. Suppose we adopt 4KiB pages.
  - a) How many page table levels are needed? Explain. (2pt)
  - b) What is the smallest size taken by all the page tables of a process that has access to one valid page from which it can read and write data. (2pt)
  - c) For each level, how many bits are available in each entry for storing information other than the frame number? (2pt).

### Problem 4 (4pt)

1. Describe an attack that exploits that a UNIX search path includes "", for instance where "echo \$PATH" returns "PATH=.:usr/local/bin:/usr/bin:/home/me/bin". (2pt)
2. what does it mean to salt passwords? describe a scenario where salting passwords is useful against security attacks. (2pt)

### Problem 5 (4p)

1. Given a resource (e.g., a file), you want to revoke the access to it for a number of domains. Is revocation easier with access-lists or with capability-lists? Explain. (2pt)
2. Examples of capabilities in current Linux versions include "CAP\_CHOWN: Make arbitrary changes to file UIDs and GIDs" or "CAP\_SYS\_CHROOT: Use chroot". Explain a scenario where it makes sense to create a setuid program with Linux capabilities. (2pt).