TDIU25 Exam

Ahmed Rezine

2017-06-07 14-18
Hus G

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: 21p, 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 fix size region of virtual memory, possibly on disk.

frame A fix size region of physical memory (DRAM).

block A data block located on disk.

Problem 1 (12p)

Part A. Assume a paged virtual memory with 16 bits logical addresses where the 9 least significant bits are used for the page offsets (i.e., byte position within a page). Assume one level paging is used.

1. What is the size, in bytes, of a single page? (1pt)
2. What is the maximum number of pages that can be addressed by a process? For this question, assume page table size and page table entry size are not an issue. (2pt).
3. What is the maximal size, in bytes, of a single page table entry that would still allow a process to fit its page entry in a page? explain. (2pt)

Part B. Assume now a paged virtual memory with 32 bits logical addresses and 4KiB pages.

1. How many pages can a process address? (1pt)
2. Assume one level paging. What is the size of a process page table assuming 4 bytes page entries? Does the table fit in a single page (2pt)
3. Assume two levels paging and 4 bytes page entries. How many bits of the virtual address should you allocate to each level in order to fit each page table (regardless of the level) in a single page? explain. (2pt)
4. Assume one level paging and 4 bytes page table entries, what is the minimum page size that would allow to fit a page table in a page? (2pt)

Problem 2 (4p)

1. Describe, in no more than half a page (excluding sketches or code snippets), a buffer overflow security attack. Use for this a small code snippet that is vulnerable to such an attack and describe how the attack proceeds. Do not hesitate to make use of a simple sketch to illustrate your description. (2pt)
2. How can segmentation and the underlying hardware help making such attacks more difficult? (2pt)

Problem 3 (8p)

Assume an idle system. Consider the workload depicted in Table (1).
Assume a one level scheduler using a Round Robin queue with a 4 time units time quantum.

Job	Arrival time	Execution time
J_1	0	7
J_2	1	5
J_3	2	3
J_4	3	1

Table 1: A workload to be scheduled

1. Draw a Gantt diagram for processes' execution and queue contents. (1pt)
2. What is the waiting time of each job? (1pt)
3. What is the turnaround time of each job? (1pt)
4. Is the scheduler preemptive? Justify (1pt)

Assume a one level scheduler using Preemptive Shortest Job First (SJF).

1. Draw a Gantt diagram for processes' execution and queue contents. (1pt)
2. What is the waiting time of each job? (1pt)
3. What is the turnaround time of each job? (1pt)
4. General question: give an advantage of using SJF instead of a round Robin scheduler, and an advantage of using round Robing instead of SJF (1pt).

Problem 4 (12 p)

Assume a 4TiB (i.e., 2^{42} bytes) hard drive. Assume a filesystem that makes use of 32 bits pointers to select logical blocks.

1. Would 512 B logical blocks allow addressing the totality of the drive? argue. (2pt)
2. What is the minimal size of a logical block? (2pt)
3. Assume 2KiB logical blocks. What is the size of the FAT table corresponding to having a single 4TiB volume? (2pt) recall that FAT is a form of linked allocation but where all links are centralized in a single table for the whole volume.
4. How large (in bytes) should a single logical block be in order to obtain a FAT table of 64 MiB (i.e., 2^{26} bytes) ? (2pt)
5. Compare advantages and disadvantages of the two logical block sizes (i.e, the block sizes adopted/obtained in questions 2 and 4). (2pt)
6. Give an advantage of FAT as opposed to contiguous allocation, and an advantage of contiguous allocation as opposed to FAT. (2pt)

Problem 5 (4p)

1. Give two differences between using two processes or using two threads when implementing concurrent applications. (2pt).
2. Give an advantage of adopting the "many user threads to one kernel thread" model instead of the "one user thread to one kernel thread" model and vice versa, (i.e., an advantage of the "one to one" versus the "many to one"). (2pt).