UNIVERSITY OF MORATUWA

FACULTY OF ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING

MSc in Computer Science
2015 Intake Semester 1 Examination

CS5202 ADVANCED OPERATING SYSTEMS

Time allowed: 2 Hours April 2015

ADDITIONAL MATERIAL: *None*

INSTRUCTIONS TO CANDIDATES:

1. This paper consists of **4** questions in **6** pages.
2. Answer **all** questions.
3. Please note that the total marks allocated to questions are different and indicated at the beginning of each question.
4. Start answering each of the main questions on a new page.
5. The maximum attainable mark for each question is given in brackets.
6. This examination accounts for 50% of the module assessment.
7. This is a closed book examination.

   NB: It is an offence to be in possession of unauthorized material during the examination.
8. Only calculators approved by the Faculty of Engineering are permitted.
9. Assume reasonable values for any data not given in or with the examination paper. Clearly state such assumptions made on the script.
10. In case of any doubt as to the interpretation of the wording of a question, make suitable assumptions and clearly state them on the script.
11. This paper should be answered only in English.

Question 1 (25 marks)

A Select the most appropriate answer, and write the corresponding sub question number and the answer number in your answer book. [10 × 1]

(i) Instruction processing consists of 3 steps:

A) fetch, decode, and execute
B) fetch, execute, and halt
C) instruction, decode, and execute
D) instruction, execute, and halt

(ii) The component of the operating system that selects one of the ready processes to run is called the

A) dispatcher
B) long-term scheduler
C) monitor
D) short-term scheduler

(iii) There are 3 processes *P*, *Q*, and *R* with execution times of 10 ms, 8 ms, and 4 ms, respectively. All the processes arrive at the same time in the order of *P*, *Q*, and *R*. What is the waiting time for process *Q*, if the scheduling algorithm is shortest job first?

A) 0
B) 4
C) 8
D) 10

(iv) For the same set of processes in question (iii), what is the average turnaround time, if the scheduling algorithm is first come first serve?

A) 4
B) 38/3
C) 50/3
D) 22

(v) Which of the following statements are TRUE about Multithreading Models?

(*p*) One-to-One model can benefit from hardware-level parallelism
(*q*) Many-to-Many model allows the OS to create any no of kernel threads
(*r*) Two-level model allows a user thread to be bound to a kernel thread

A) (*p*) and (*q*) only
B) (*p*) and (*r*) only
C) (*q*) and (*r*) only
D) All three

(vi) Which of the following statements are TRUE about Semaphores?

(*p*) Enable process synchronization
(*q*) Usually value of a semaphore cannot be read, even if you read it is not useful
(*r*) Combination of multiple semaphores will never lead to a deadlock

A) (*p*) and (*q*) only
B) (*p*) and (*r*) only
C) (*q*) and (*r*) only
D) All three

(vii) One technique for overcoming external fragmentation of memory is _____ .

A) compaction B) loading

C) partitioning D) relocation

(viii) The _____ structure indexes page table entries by frame number rather than by virtual page number.

A) hash table B) inverted page table

C) page table D) segment table

(ix) A _____ chooses only among the resident pages of the process that generated the page fault in selecting a page to replace.

A) global replacement policy B) local replacement policy

C) page placement policy D) page replacement policy

(x) Which of the following statements is True about real time operating systems?

A) Hard real-time systems may occasionally miss a deadline

B) In Rate Monotonic Scheduling longer duration jobs have higher priority

C) In Earliest Deadline First Scheduling highest priority is given to the job with the shortest deadline

D) Only mission critical systems have real-time processing requirements

B State whether the following statements are TRUE or FALSE. Give one sentence justification for your answer. [5 × 2]

(i) Multithreading is useful for applications that perform a number of essentially independent tasks that need not be serialized.

(ii) Scheduler activations provide upcalls to communicate from kernel to the thread library allowing the creation of appropriate number of kernel-level threads.

(iii) A physical address is the location of a word relative to the beginning of the program and the processor translates that into a logical address.

(iv) A good scheduling algorithm should reduce the throughput while maximizing the latency.

(v) In general, real-time operating systems must provide preemptive kernels.

C Write the most appropriate short answer (word/phrase) for following questions. [5 × 1]

(i) _____ is a large kernel, including scheduling file system, networking, device drivers, memory management and more.

(ii) _____ refers to the ability of multiple process (or threads) to share code, resources or data in such a way that only one process has access to shared object at a time.

(iii) _____ brings a page into physical memory only when it is essential.

(iv) Copy-on-Write (COW) allows both parent and child processes to initially share the _____ pages in memory.

(v) _____ fragmentation is caused when a file block is not fully utilized.

Question 2 (25 marks)

(i) Discuss whether a micro or monolithic kernel based design is more suitable for the following requirements. Justify your selection.

a) A dedicated server designed to run only a specific type of application which requires high throughput. [3]

b) An open source mobile operating system to run on many different smartphones by multiple manufacturers. [3]

c) A real time operating system. [3]

(ii) Following 2 options are suggested for a 64-bit memory system. Discuss the advantages and disadvantages of each of the following design choices.

a) A logical address with a 32-bit page offset and a 2-level page table. [4]

b) A logical address with a 20-bit page offset and an inverted page table. [4]

(iii) Briefly discuss 3 advantages of virtual memory. [3]

(iv) Solaris uses *adaptive mutexes* for efficiency when protecting data from short code segments. An adaptive mutex tries a spin lock a few times, if it fails the process/thread goes and sleep.

Do you agree or disagree that adaptive mutex has better efficiency than using only a spin lock or a mutex? Discuss. [5]

Question 3 (25 marks)

One of your clients recently a deployed a new Customer Relationship Management (CRM) system. While the CRM provides the necessary business functions, client is experiencing significant performance issues during certain times of the day. According to the client following are the symptoms observed:

"Our server has 4 Intel Xeon cores running at 3.2 GHz. Server has 16 GB of RAM and 4 TB of disk space. CRM database is also on the same node. Server runs Ubuntu 12.04. Web and database servers are Apache and MySQL, respectively.

Sometimes it takes several seconds to search for the contact details of a client. It takes tens of seconds for activities like generation of invoices, sales and stock updates, and customer and supplier matching. We have realized the CPU utilization of the server is mostly below 40%. It occasionally jumps to 100% when we run batch updates. However, those updates are shot lived and carried out mostly after typical office hours. Average number of processes or threads in the ready queue is 0.82 per CPU core. However, average number of process or threads in the waiting queue is 13.3 per CPU core. Average memory utilization is 48% while peak memory utilization is 100%. Memory peaks tends to occur at random times of the day; however, they do not last more than 1.6 seconds in average. Disk utilization is 25%. However, we see peak read and write IOPS (I/O Operations per Second) as high as 1,243,400 and 456,700, respectively. Most of the disk reads and writes are short lived. However, long reads and writes that last tens of seconds are not uncommon. We have also observed that several locks on certain database entries last few seconds."

They want your feedback to diagnose the issue(s) of the CRM system and want your input on how to resolve the issue(s). Answer following questions based on this scenario.

(i) Is the workload CPU bound or I/O bound? Justify. [4]

(ii) To improve performance of the CRM, should we give high priority to CPU bound processes/treads or I/O bound processes/threads? Justify. [3]

(iii) Using the given case as an example, briefly describe what causes a process to move from waiting queue to ready queue. [3]

(iv) Using the given case as an example, discuss how the following pairs of scheduling criteria/objectives conflict in certain settings.

a) CPU utilization and response time [3]

b) I/O device utilization and CPU utilization [3]

(v) One of your colleagues suggests to increase the swap space to improve the performance as the server has enough free disk space. Do you agree or disagree to this suggestion? Briefly explain. [4]

(vi) Another suggestion is to increase the number of CPU cores to 8, RAM to 32 GB, and also have a fast disk system with RAID (Redundant Array of Independent Disks). Will this make the workload change its characteristics from CPU bound to I/O bound or vice versa? Discuss. [5]

Question 4 (25 marks)

Answer **any 5** questions. These questions are based on papers discussed in the class. [5 × 5]

(i) Following is a diagram of a priority-based multilevel feedback model. Describe what happens, if the execution time t of a newly arriving process with *priority 1* is much larger than the quantum size q (i.e., $t >> q$).

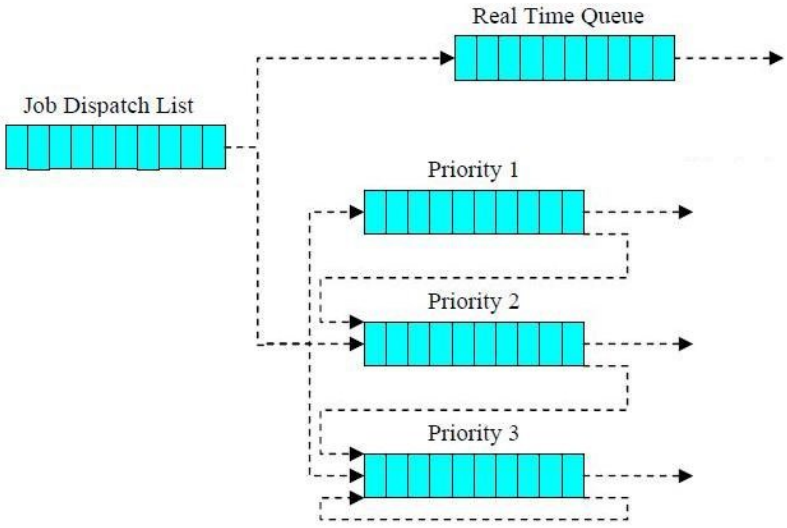

(ii) Frequent allocation and freeing of memory creates holes within the memory allocated to a process. This leads to inefficient memory utilization. Briefly describe how the log-structured memory addresses this issue.

(iii) Following statement is extracted from "Google File System" paper.

"After a file is deleted, GFS does not immediately reclaim the available physical storage. It does so only lazily during regular garbage collection at both the file and chunk levels."

Discuss the advantages and disadvantages of this approach compared to a typical file system like FAT32, NTFS, and Ext3 which reclaims storage as soon as a file is deleted.

(iv) Following statement is extracted from "Apollo: Scalable and Coordinated Scheduling for Cloud-Scale Computing" paper.

"To balance scalability and scheduling quality, Apollo adopts a distributed and (loosely) coordinated scheduling framework, in which independent scheduling decisions are made in an optimistic and coordinated manner by incorporating synchronized cluster utilization information."

Discuss the advantages and disadvantages of such a design.

(v) "Historically Windows is known to have better multimedia performance than Linux."

By considering the architectural aspects of the two operating systems discuss the reason for this behavior.

(vi) Mobile processors with heterogeneous CPU cores create new problems of cache coherence. Outline a solution to address this problem where a mobile app can seamlessly switch between different CPU cores to gain power-performance advantages.

---------- END OF THE PAPER ----------