UNIVERSITY OF MORATUWA

FACULTY OF ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING

MSc in Computer Science
2014 Intake Semester 1 Examination

CS5202 ADVANCED OPERATING SYSTEMS

Time allowed: 2 Hours April 2014

ADDITIONAL MATERIAL: *None*

INSTRUCTIONS TO CANDIDATES:

1. This paper consists of **4** questions in **5** pages.
2. Answer all questions.
3. Please note that the total marks allocated to questions are different.
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 unauthorised 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) The _____ routine determines the nature of the interrupt and performs whatever actions are needed.

A) interrupt handler
B) instruction signal
C) interrupt signal
D) program handler

(ii) The component of the operating system that selects the next process to run is called the

A) dispatcher
B) kernel
C) monitor
D) process control block

(iii) A common strategy to give each process in the queue same time in turn is referred to as a _____ technique.

A) multithreading
B) round-robin
C) serial processing
D) time slicing

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

(*p*) One-to-One model enables the creation of any number of kernel-level threads
(*q*) Many-to-One model cannot benefit from hardware-level parallelism
(*r*) Many-to-Many model allows the OS to create a sufficient no of kernel threads

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

(v) Which of the following statement is NOT true about semaphores?

A) Utilize priority inversion
B) Enforce mutual exclusion
C) Enable process synchronization
D) Can be implemented without busy waiting

(vi) Relocation register value is 14,000 and limit register value is 2,048. Logical address is 2,046. What is the physical address?

A) 2,048
B) 2,046
C) 16,046
D) This address is invalid

(vii) Main memory divided into a number of static partitions at system generation time is referred to as

A) dynamic partitioning
B) fixed partitioning
C) simple paging
D) simple segmentation

(viii) _____ is transparent to the programmer and eliminates external fragmentation providing efficient use of main memory.

A) Hashing B) Paging

C) Segmentation D) Thrashing

(ix) In general, real-time operating systems must provide:

(*p*) Preemptive, priority-based scheduling

(*q*) Preemptive kernels

(*r*) Latency must be minimized

A) (*p*) and (*q*) only B) (*p*) and (*r*) only

C) (*q*) and (*r*) only D) All three

(x) In rate monotonic scheduling

A) longer duration job has higher priority

B) shorter duration job has higher priority

C) priority does not depend on the duration of the job

D) none of the above

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

(i) Multiprogramming was designed to keep the processor and I/O devices, including storage devices, simultaneously busy to achieve maximum efficiency.

(ii) A process that is waiting for access to a critical section consumes processor time.

(iii) Swapping is not an I/O operation so it will not enhance performance.

(iv) Virtual memory reduces the I/O required to load and unload user programs into memory.

(v) The principle of locality states that program and data references within a process tend to cluster.

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

(i) 2 processes can exchange data/information using _____ .

(ii) _____ allows a sequence of instructions to execute as a group or not execute at all.

(iii) In _____ policy, the scheduler always chooses the process that has the shortest expected remaining processing time.

(iv) _____ is a facility that allows programs to address memory from a logical point of view, without regard to the amount of main memory physically available.

(v) A _____ task has an associated deadline that is desirable, but not mandatory.

Question 2 (25 marks)

(i) Briefly discuss long-term scheduling vs. short-term scheduling. [5]

(ii) Discuss how the following pairs of scheduling criteria/objectives conflict in certain settings.

a) CPU utilization and throughput [3]

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

(iii) Several of your friends are planning to develop an open source, geographically distributed cloud storage called *community cloud*. They are planning to use a peer-to-peer approach, where each user of *community cloud* contributes part of his/her free disk space to form one large distributed file system. This file system is to be named CFS (Community cloud File System). They want your feedback on the following design options of CFS.

While your OS class did not cover file systems in detail, you can still contribute to the CFS design by drawing ideas from concepts learned in other topics.

a) Recommend a suitable block size for CFS. Justify. [3]

b) Files that are larger than the block size need to be broken into multiple blocks. Therefore, recommend a suitable disk block allocation method for CFS. You need to explain where to place the blocks and how to keep track of them. [5]

c) List the necessary metadata to include in a file control block of CFS.

Hint: A file control block is used to keep track the metadata related to a file. [3]

d) Would you recommend allowing concurrent access to the same file? Explain. [3]

Question 3 (25 marks)

(i) Briefly describe an example that will cause a process to perform each of the following state transitions.

a) Transition from Blocked to Blocked/Suspend state. [2]

b) Transition from Ready/Suspend to Running state. [2]

(ii) Using a suitable example explain what is a Memory-Mapped File. [3]

(iii) It is known that memory access time of a particular system is 100 ns and average page-fault service time is 4 ms. If one access out of 25,000 causes a page fault under demand paging, what is the effective access time? [3]

(iv) Answer any **3** of the following questions. These questions are based on papers discussed in the class. [3 × 5]

a) Compare and contrast (i.e., similarities and dissimilarities) exokernel, microkernel, and monolithic kernel based operating systems.

b) Linux kernel gives a relatively high priority for I/O-bound programs. Describe the 2 key reasons behind this decision?

c) Explain Symmetric Multiprocessing in the context of Windows operating system.

d) Briefly describe the roles of Android Runtime and Linux Kernel in the Android architecture.

Question 6 (25 marks)

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

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

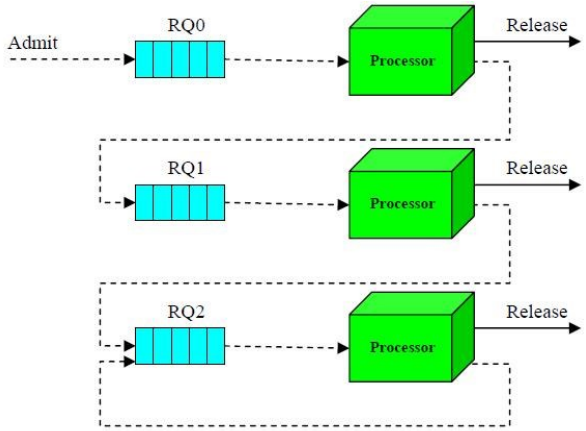

(ii) Using suitable example(s) briefly describe issues associated with thread pools.

(iii) How does dynamically changing the size of Superpages (i.e., promotion and demotion) help to improve Translation Lookaside Buffer (TLB) coverage?

(iv) By default, Google File System replicates a chunk on 3 chunk servers. Discuss the pros and cons of creating only 3 replicas.

(v) Explain how a shadow driver can mask the failure of a device driver from the applications.

(vi) Describe Space Multiplexing in the context of multi-core and cloud computing systems.

(vii) Why real-time capabilities are required in multimedia devices? How can such requirements be supported while reducing the energy consumption of a multimedia device?

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