



UNIVERSITY OF MORATUWA

FACULTY OF ENGINEERING

DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING

MSc in Computer Science
Semester 1 Examination: 2013/2014

CS5202 ADVANCED OPERATING SYSTEMS

Time allowed: 2 Hours

April 2013

ADDITIONAL MATERIAL: *None*

INSTRUCTIONS TO CANDIDATES:

1. This paper consists of 6 questions in 6 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 (15 marks)

Select the most appropriate answer, and write the corresponding sub question number and the answer number in your answer book.

[15 × 1]

- (i) Instruction processing consists of two steps:
A) execute and halt B) instruction and execute
C) instruction and halt D) fetch and execute
- (ii) A _____ is a unit of activity characterized by the execution of a sequence of instructions, a current state, and an associated set of system resources.
A) identifier B) kernel
C) process D) state
- (iii) Which is NOT a multithreading model?
A) One-to-One B) One-to-Many
C) Many-to-One D) Many-to-Many
- (iv) Which of the following occurs when multiple processes or threads read and write data items so that the final result depends on the order of execution of instructions in the multiple processes?
A) Atomic operation B) Deadlock
C) Livelock D) Race condition
- (v) Which of the following is a high-level abstraction that provides a convenient and effective mechanism for process synchronization?
A) Atomic operation B) Co-routine
C) Critical section D) Monitor
- (vi) Relocation register value is 14,000 and limit register value is 2,048. Logical address is 2,346. What is the physical address?
A) 2,048 B) 2,346
C) 16,346 D) This address is invalid
- (vii) In Dynamic Partitioning technique of memory management, the placement algorithm that chooses the block that is closest in size to the request is called _____ .
A) first-fit B) best-fit
C) last-fit D) next-fit
- (viii) The page table for each process maintains _____ .
A) physical memory location of the process
B) logical memory location of the process
C) frame location for each page of the process
D) page location for each frame of the process

- (ix) A _____ chooses only among the resident pages of the process that generated the page fault in selecting a page to replace.
- A) frame replacement policy B) global replacement policy
 C) local replacement policy D) page placement policy
- (x) Suppose 64 frames are to be allocated to 2 processes. Process P requests 30 pages while process Q requests 90 pages. How many frames will be allocated to process P under the Proportional Allocation scheme?
- A) 16 B) 30
 C) 48 D) 90
- (xi) Which of the following is a technique that smoothes out peaks in I/O demands?
- A) Blocking B) Buffering
 C) Smoothing D) Tracking
- (xii) Which is an advantage of contiguous allocation of disk space?
- (p) Simple to implement
 (q) Support sequential access
 (r) Support random access
- A) (p) and (q) only B) (p) and (r) only
 C) (q) and (r) only D) All three
- (xiii) Which layer in a layered file system handles read and write operations to a given physical block.
- A) Basic file system B) File organization module
 C) I/O control D) Logical file system
- (xiv) Which of the following tasks must meet its deadline, otherwise it will cause unacceptable damage or a fatal error to the system?
- A) Aperiodic task B) Hard real-time task
 C) Periodic task D) Soft real-time task
- (xv) In general, real-time operating systems must provide:
- (p) Preemptive, priority-based scheduling
 (q) Nonpreemptive kernels
 (r) Latency must be minimized
- A) (p) and (q) only B) (p) and (r) only
 C) (q) and (r) only D) All three

Question 2 (10 marks)

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

- (i) If a process is swapped out, all of its threads are necessarily swapped out.
- (ii) Peterson's solution can be applied to enforce mutual exclusion among m threads.
- (iii) All segments of all programs must be of the same length.
- (iv) Any protection mechanism must have the flexibility to allow several processes to access the same portion (i.e., address range) of main memory.
- (v) Linux and Windows retain symbolic links when a file is deleted.

Question 3 (10 marks)

Write the most appropriate short answer (word/phrase) for the following questions. [10 × 1]

- (i) A process in the _____ state is in main memory and available for execution.
- (ii) It is necessary to _____ the activities of various threads so they do not interfere with each other or corrupt data structures.
- (iii) A _____ is a mutual exclusion mechanism in which a process executes in an infinite loop waiting for the value of a lock variable to indicate availability.
- (iv) The _____ is a directed graph that depicts a state of the system of resources and processes, with each process and each resource represented by a node.
- (v) Logical memory address is generated by the _____.
- (vi) As time goes on, memory becomes more and more fragmented and memory utilization declines, creating a phenomenon referred to as _____.
- (vii) The portion of a process that is actually in main memory at any time is defined to be the _____ of the process.
- (viii) _____ records each update to the file system as a transaction.
- (ix) In _____ systems correctness depends not only on logical result but also time at which results are produced.
- (x) In rate monotonic scheduling, shorter the period of a process, _____ the priority.


Question 4 (15 marks)

- (i) What is scheduling and why is it important? [3]
- (ii) Discuss how throughput and maximum waiting time scheduling criteria conflict in certain settings. [3]

- (iii) There are three processes P , Q , and R with execution times of 10 ms, 6 ms, and 4 ms, respectively. All the processes arrive at the same time in the order of P , Q , and R .
- A) Calculate the turnaround time for each process, if the scheduling algorithm is first come first serve. [2]
- B) Calculate the turnaround time for each process, if the scheduling algorithm is shortest job first. [2]
- C) For (A) and (B) calculate the average execution time. [2]
- D) Based on answers to (C), what is the more appropriate scheduling algorithm? [1]
- E) Can you generalize the answer for (D) above for any scheduling problem? Explain your answer. [2]

Question 5 (20 marks)

- (i) Briefly explain how logical to physical address translation happens with hierarchical page tables. [4]
- (ii) Briefly explain the steps involved in handling a page fault. [3]
- (iii) It is known that memory access time of a particular system is 100 ns and average page-fault service time is 6 ms. If one access out of 2,500 causes a page fault under demand paging, what is the effective access time? [3]
- (iv) Following graph illustrates CPU utilization vs. degree of multiprogramming. Briefly describe why the graph show such a behaviour while focusing on regions (X) and (Y) on the graph. [4]



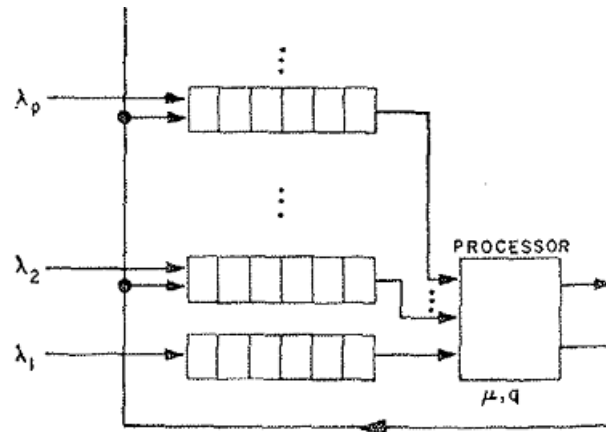
- (v) Compare and contrast consistency semantics of Unix File System (UFS), Network File System (NFS), and Andrew File System (AFS). [6]

Question 6 (30 marks)

Answer any **six (6)** of the following questions.

[6 × 5]

- (i) How do exokernels differ from traditional microkernel and monolithic kernel based operating systems?
- (ii) What are thread pools? Describe issues associated with those.
- (iii) Following is a diagram of multilevel feedback model with priorities. Briefly describe what happens if a higher priority process ($p > 2$) arrives while a process in the 2nd queue is being executed. Assume high p means higher priority.



- (iv) How do Superpages improve Translation Lookaside Buffer (TLB) coverage?
- (v) The Google File System uses 64 MB chunks. Discuss pros and cons of such a large chunk size.
- (vi) Briefly describe the concept of shadow drivers.
- (vii) Briefly describe how the Factored Operating System (FOS) addresses some of the major issues in multi-core systems and clouds.
- (viii) Briefly describe how the scheduling is performed in Linux 2.6.8.1 scheduler and its performance.
- (ix) How do Protected Virtual Machines provide resource partitioning and retain quality of service?
- (x) What is Dalvik virtual machine in Android and how does it enhance performance?

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