

UVA WELASSA UNIVERSITY
FACULTY OF SCIENCES AND TECHNOLOGY
COMPUTER SCIENCE & TECHNOLOGY DEGREE PROGRAM
FIRST SEMESTER EXAMINATION – JAN/FEB – 2011
CST307-2 OPERATING SYSTEMS
CST302-2 OPERATING SYSTEMS

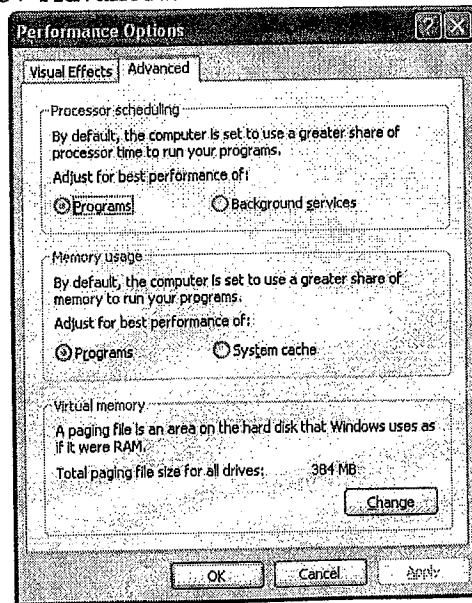
Time Allowed: Two hours

Answer all questions in part A and part B. wrong answer carries negative points
And attach your all question papers to the answer script

Part A

Q1.

- a. Windows XP (and Windows 2000) allows users to set different quantum for application programs and background services in the following dialog box: Control Panel > System > Advanced > Performance > Advanced.



Explain the effect of the two options. Compare this method with the approach of setting higher priority for application programs. (Suggestion: consider that case of a CPU-bound application program running with a CPU-bound background service.)

- b. Provide rough definitions for the following terms:
- i. process
 - ii. process state
 - iii. processor
 - iv. multiprogramming
 - v. time sharing
 - vi. non-determinism

- c. On all current computers, some parts of the interrupt handlers are always written in assembly language. Explain why?
- d. Which of the following should require instructions that are allowed only in kernel mode? Explain why?
 - i. Disable all interrupts.
 - ii. Read the time-of-day clock.
 - iii. Set the time-of-day clock.
 - iv. Change the memory map.
- e. When an interrupt or system call transfers control to the OS, execution stops using the stack of the interrupted process and switches to a separate kernel stack. Explain why?
- f. Can a process make a transition from the Ready state to the Blocked state? Explain why? or why not?

(40 marks)

Q2.

Determine the sequence of execution of processes for each of the following scheduling algorithms.

- A. First Come First Served
- B. Shortest Job First
- C. Shortest Remaining Time First
- D. Round Robin with quantum size = 5
- E. Round Robin with quantum size = 10

The four columns give (a) Process number, (b) Arrival time, (c) Burst time (how long the process needs the CPU), (d) Earliest possible completion time.

1	0	14	14
2	2	12	14
3	5	10	15
4	7	4	11
5	19	7	26

For parts (c) and (d), assume that an arriving process joins the queue before the quantum expires if such events occur simultaneously.

- a. What is the turnaround time for each job, the waiting time for each job, and the average waiting time for all jobs for each of the scheduling algorithms?
- b. What are other factors must be taken into consideration?
- c. In the round-robin scheduling, why can't you start new processes at the front of the queue instead of at the end?
- d. What is the difference between a preemptive and a non-preemptive scheduling algorithm?
- e. Using above five algorithms (A) through (E) in the question, which are preemptive and which are non-preemptive? Why is a non-preemptive scheduling algorithm unlikely to be used in a computer centre?

- f. Explain the operation of a multilevel feedback queue. What advantages arise from having different quantum sizes on different levels of a multilevel queue?
- g. Give an example reason why a real-time thread may require different amounts of CPU time at different times.

(25 marks)

Q3.

- a. Explain with the aid of a diagram how files managed in:
- (i) The UNIX file-system
 - (ii) The FAT32 file-system
 - (iii) The NTFS file-system
- b. A researcher suggests using non-volatile flash memory to store the NTFS log file. He believes this will improve performance, reduce power consumption and make the system more resilient to failure. Is he wrong? Briefly justify your answer in each case.
- c. Consider the UNIX file system. Assume that disk blocks are 16 kilobytes in size and each pointer to a disk block requires 32 bytes. What is the largest possible file size using this design? (You don't need to calculate the final numeric answer, just provide an expression for evaluating it and explain how it is derived). Explain and justify any assumptions you make.

(15 marks)

