~~Oct~~ Nov-17

TyBSc/04. V sem

Q. P. Code: 24043

(2 ½ Hours) **[Total Marks: 75]**

N.B: (1) All questions are compulsory.
(2) Figures to the right indicate marks.
(3) Illustrations, in-depth answers and diagrams will be appreciated.

Q1. **Attempt the following (any THREE):** **(15)**

(A) List five services provided by an operating system that is designed to make it more convenient for users to use the computer system. In what cases it would be impossible for user-level programs to provide these services? Explain.

(B) Which of the following instructions should be privileged?
a. Set value of timer.
b. Read the clock.
c. Clear memory.
d. Issue a trap instruction.
e. Turn off interrupts.
f. Modify entries in device-status table.
g. Switch from user to kernel mode.
h. Access I/O device.

(C) List five services provided by an operating system, and explain how each creates convenience for users. In which cases would it be impossible for user-level programs to provide these services? Explain your answer.

(D) What is the purpose of the command interpreter? Why is it usually separate from the kernel? Would it be possible for the user to develop a new command interpreter using the system-call interface provided by the operating system?

(E) What are the two models of inter-process communication? What are the strengths and weaknesses of the two approaches?

(F) What is the purpose of interrupts? What are the differences between a trap and an interrupt? Can traps be generated intentionally by a user program? If so, for what purpose?

Q2. **Attempt the following (any THREE):** **(15)**

(A) What are the benefits and the disadvantages of each of the following? Consider both the system level and the programmer level.
a. Synchronous and asynchronous communication
b. Automatic and explicit buffering
c. Send by copy and send by reference
d. Fixed-sized and variable-sized messages

(B) Describe the actions taken by a thread library to context switch between user-level threads.

(C) Describe the differences among short-term, medium-term, and long term scheduling.

(D) Why is it important for the scheduler to distinguish I/O-bound programs from CPU-bound programs?

(E) Consider the following set of processes, with the length of the CPU burst given in milliseconds:

Process	Burst Time	Priority
P1	10	3
P2	1	1
P3	2	3
P4	1	4
P5	5	2

The processes are assumed to have arrived in the order Pi, P 2, P3, P4, P5, all at time 0.

a. Draw four Gantt charts that illustrate the execution of these processes using the following scheduling algorithms: FCFS, SJF, non-preemptive priority (a smaller priority number implies a higher priority), and RR (quantum = 1).

(F) Which of the following scheduling algorithms could result in starvation? Explain.

a. First-come, first-served

b. Shortest job first

c. Round robin

Q3. Attempt the following (any THREE): (15)

(A) What is the meaning of the term busy 'waiting? What other kinds of waiting are there in an operating system? Can busy waiting be avoided altogether? Explain your answer.

(B) Explain why implementing synchronization primitives by disabling interrupts is not appropriate in a single-processor system if the synchronization primitives are to be used in user-level programs.

(C) Show that, if the wait( ) and signal ( ) semaphore operations are not executed atomically, then mutual exclusion may be violated.

(D) Consider the following snapshot of a system:

	Allocation	*Max*	*Available*
	A B C D	*A B C D*	*A B C D*
P_0	0 0 1 2	0 0 1 2	1 5 2 0
P_1	1 0 0 0	1 7 5 0	
P_2	1 3 5 4	2 3 5 6	
P_3	0 6 3 2	0 6 5 2	
P_4	0 0 1 4	0 6 5 6	

Answer the following questions using the banker's algorithm:

a. What is the content of the matrix Need?

b. Is the system in a safe state?

c. If a request from process Pi arrives for (0,4,2,0), can the request be granted immediately?

(E) Given five memory partitions of 100 KB, 500 KB, 200 KB, 300 KB, and 600 KB (in order), how would each of the first-fit, best-fit, and worst-fit algorithms place processes of 212 KB, 417 KB, 112 KB, and 426 KB (in order)? Which algorithm makes the most efficient use of memory?

(F) Explain the difference between internal and external fragmentation.

Q4. **Attempt the following (any THREE):** **(15)**

(A) Consider a file system where a file can be deleted and its disk space reclaimed while links to that file still exist. What problems may occur if a new file is created in the same storage area or with the same absolute path name? How can these problems be avoided?

(B) What are the advantages of the variant of linked allocation that uses a FAT to chain together the blocks of a file?

(C) Suppose that a disk drive has 5,000 cylinders, numbered 0 to 4999. The drive is currently serving a request at cylinder 143, and the previous request was at cylinder 125. The queue of pending requests, in FIFO order, is:
86,1470, 913, 1774, 948, 1509, 1022,1750,130
Starting from the current head position, what is the total distance (in cylinders) that the disk arm moves to satisfy all the pending requests for each of the following disk-scheduling algorithms?
a. FCFS
b. SSTF

(D) Compare the performance of C-SCAN and SCAN scheduling, assuming a uniform distribution of requests. Consider the average response time (the time between the arrival of a request and the completion of that request's service), the variation in response time, and the effective bandwidth. How does performance depend on the relative sizes of seek time and rotational latency?

(E) When multiple interrupts from different devices appear at about the same time, a priority scheme could be used to determine the order in which the interrupts would be serviced. Discuss what issues need to be considered in assigning priorities to different interrupts.

(F) To protect a system, at which levels we must take security measures?

Q5. **Attempt the following (any THREE):** **(15)**

(A) Consider a paging system with the page table stored in memory.
a. if a memory reference takes 200 nanoseconds, how long does a paged memory reference take?
b. If we add TLBs, and 75 percent of all page-table references are found in the TLBs, what is the effective memory reference time?
(Assume that finding a page-table entry in the TLBs takes zero time, if the entry is there.)

(B) What is process? What operations are performed on process under process management?

(C) Assume that we have a demand-paged memory. The page table is held in registers. It takes 8 milliseconds to service a page fault if an empty frame is available or if the replaced page is not modified and 20 milliseconds if the replaced, page is modified. Memory-access time is 100 nanoseconds.
Assume that the page to be replaced is modified 70 percent of the time. What is the maximum acceptable page-fault rate for an effective access time of no more than 200 nanoseconds?

(D) What are the advantages and disadvantages of supporting memory-mapped I/O to device control registers?

(E) What are the various kinds of performance overheads associated with servicing an interrupt?

(F) Discuss security threats and attacks.

****