OPERATING SYSTEMS
(MCA1204)

Time Allotted : 2½ hrs **Full Marks : 60**

Figures out of the right margin indicate full marks.

Candidates are required to answer Group A and any 4 (four) from Group B to E, taking one from each group.

Candidates are required to give answer in their own words as far as practicable.

Group – A

1. Answer any twelve: **12 × 1 = 12**

Choose the correct alternative for the following

(i) The degree of multi-programming is
(a) The number of processes in memory
(b) The number of processes in the I/O queue
(c) The number of processes in the ready queue
(d) The number of processes executed per unit time

(ii) Which among the following is shared by the threads of a multithreaded process?
(a) registers
(b) stack
(c) program counter
(d) data section

(iii) When several processes access the same data concurrently and the outcome of the execution depends on the particular order in which access takes place, is called
(a) Starvation
(b) Mutual exclusion
(c) Race condition
(d) Critical section

(iv) At a particular time of computation, the value of a counting semaphore is 7. Then 20 numbers of P (wait) operations and 'x' numbers of V (signal) operations were completed on this semaphore. If the final value of the semaphore is 5, the value of 'x' is
(a) 13
(b) 15
(c) 18
(d) 22

(v) Which of the following is a hardware based solution to critical section problem?
(a) Peterson's solution
(b) Dekker's solution
(c) Monitors
(d) TestAndSet

(vi) Which page is selected by the Optimal Page Replacement Algorithm?
(a) The page that has been used minimum number of times
(b) The page that has been used maximum number of times
(c) The page that has not been used for the longest time in the past
(d) The page that will not be used for the longest time in the future

(vii) An edge from process P1 to process P2 in a wait-for graph indicates that
(a) P1 is waiting for P2 to release a resource that P1 needs
(b) P2 is waiting for P1 to release a resource that P2 needs
(c) P1 is waiting for P2 to leave the system
(d) P2 is waiting for P1 to leave the system

(viii) If a process has 24 KB of logical address space and the page size is 4096 bytes, the number of pages in the process is
(a) 12.
(b) 6.
(c) 16.
(d) 8.

(ix) Which one of the following is not an attack, but a search for vulnerabilities to attack?
(a) Denial of service
(b) Port scanning
(c) Memory access violation
(d) Dumpster diving

(x) Which one of the following facilitates transfer of bulk data from hard disk to main memory with the highest throughput?
(a) Programmed I/O transfer
(b) Interrupt driven I/O transfer
(c) DMA based I/O transfer
(d) Polling based I/O transfer

Fill in the blanks with the correct word

(xi) The interval from the time of submission of a process to the time of completion is the _____ time.

(xii) The initial program that runs when the computer is powered up is called _____.

(xiii) The technique of gradually increasing the priority of a process that wait in a system for a long time is known as _____.

(xiv) _____ occurs when a system spends more time paging than executing.

(xv) A _____ is the kernel data structure that represents a process in an operating system.

Group - B

2. (a) What resources are used when a thread is created? How do they differ from those used when a process is created?

[(CO2)(Understand/LOCQ)]

(b) The following processes are being scheduled using a preemptive, priority-based, round-robin scheduling algorithm.

Process	Priority	Burst	Arrival
P1	8	15	0
P2	3	20	0
P3	4	20	20
P4	4	20	25
P5	5	5	45
P6	5	15	55

Each process is assigned a numerical priority, with a higher number indicating a higher relative priority. The scheduler will execute the highest priority process. For processes with the same priority, a round-robin scheduler will be used with a time quantum of 10 units. If a process is preempted by a higher-priority process, the preempted process is placed at the end of the queue.

(i) Show the scheduling order of the processes using a Gantt chart.
(ii) What is the turnaround time for each process?
(iii) What is the waiting time for each process?

[(CO2)(Apply/IOCQ)]

4 + (4 + 2 + 2) = 12

3. (a) Describe the actions taken by a kernel to context-switch between processes. *[(CO2)(Understand/LOCQ)]*

(b) Consider the following program:

```
int main (void) {
    int id = fork();
    int x = 5;
    if (id == 0)
        x += 5;
    else {
        id = fork();
        x += 10;
        if(id)
            x += 5;
    }
    return 0;
}
```

How many different copies of the variable x are there? What are their values when their process finishes? *[(CO2)(Analyse/IOCQ)]*

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

Process	Burst	Arrival
P1	8	0
P2	4	1
P3	9	2
P4	5	3

Compute the turnaround time and waiting time of each process for the following scheduling algorithms: Nonpremptive SJF scheduling and preemptive SJF scheduling.

[(CO2)(Apply/IOCQ)]

4 + 4 + 4 = 12

Group - C

4. (a) State the classical definition of wait() and signal() operation of semaphore. How can the classical definitions be modified to have an implementation without busy waiting? *[(CO3)(Apply/IOCQ)]*

(b) Two processes, P1 and P2, need to access critical section of their code. Consider the following synchronization construct used by the processes:

```
/*P1*/
while ( true )
{
    p1WantsToEnter = true;
    while(p2WantsToEnter);
    /* critical section code */
     p1WantsToEnter = false;
    /* remainder section */
}
```

```
/*P2*/
while ( true )
{
    p2WantsToEnter = true;
    while(p1WantsToEnter);
    /* critical section code */
     p2WantsToEnter = false;
    /* remainder section */
}
```

p1WantsToEnter and p2WantsToEnter are shared variables, which are initialized to false.
Analyse the synchronization construct used by the processes to check if

(i) Mutual exclusion is preserved?
(ii) The progress requirement is satisfied?

(iii) The bounded-waiting requirement is met? *[(CO3)(Analyse/IOCQ)]*

(c) Explain why spinlocks are not appropriate for single-processor systems yet are often used in multiprocessor systems. *[(CO3)(Understand/LOCQ)]*

(2 + 3) + 4 + 3 = 12

5. (a) Describe the classical Producer Consumer problem of process synchronization. Deduce a solution of it using semaphore. *[(CO3)(Apply/IOCQ)]*

(b) Explain with an example how mutual exclusion is ensured with test-and-set instruction. *[(CO3)(Apply/IOCQ)]*

(2 + 6) + 4 = 12

Group - D

6. (a) Consider a system with five processes < P1, P2, P3, P4, P5 > and four resource types A, B, C, D. Resource type A has 3 instances; B has 14 instances, C and D has 12 instances respectively. Suppose that at time T0, the following snapshots of the system have been taken.

	Allocation				Max			
	A	B	C	D	A	B	C	D
P1	0	0	1	2	0	0	1	2
P2	1	0	0	0	1	7	5	0
P3	1	3	5	4	2	3	5	6
P4	0	6	3	2	0	6	5	2
P5	0	0	1	4	0	6	5	6

Test whether the system is currently in a safe-state. If a request for process P2 arrives for (0, 4, 2, 0), can the request be granted immediately? *[(CO4)(Analyse/IOCQ)]*

(b) Given six memory partitions of 100 MB, 170 MB, 40 MB, 205 MB, 300 MB, and 185 MB (in order), how would the first-fit, best-fit, and worst-fit algorithms place processes of size 200 MB, 15 MB, 185 MB, 75 MB, 175 MB, and 80 MB (in order)? Indicate which—if any—requests cannot be satisfied. Comment on how efficiently each of the algorithms manages memory. *[(CO4)(Analyse/IOCQ)]*

6 + 6 = 12

7. (a) Consider a paging system with the page table stored in memory. If a memory reference takes 50 nanoseconds, how long does a paged memory reference take? If we add TLBs, and if 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 2 nanoseconds, if the entry is present). *[(CO6)(Apply/IOCQ)]*

(b) Consider the following page reference string: 1, 2, 3, 4, 2, 1, 5, 6, 2, 1, 2, 3, 7, 6, 3, 2, 1, 2, 3, 6. Assuming demand paging with three frames, how many page faults would occur for the following replacement algorithms: LRU replacement, FIFO replacement, Optimal replacement? *[(CO5)(Apply/IOCQ)]*

(c) Explain indefinite postponement. How does it differ from deadlock? *[(CO4)(Understand/LOCQ)]*

(1 + 2) + 6 + 3 = 12

Group - E

8. (a) Contrast the performance of the three techniques for allocating disk blocks (contiguous, linked, and indexed) for both sequential and random file access. *[(CO5)(Analyse/IOCQ)]*

(b) Suppose that a disk drive has 5000 cylinders, numbered 0 to 4999. The drive is currently serving a request at cylinder 2150, and the previous request was at cylinder 1805. The queue of pending requests, in FIFO order, is: 2069; 1212; 2296; 2800; 544; 1618; 356; 1523; 4965; 3681. 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 FCFS, SCAN and C-SCAN disk-scheduling algorithms? *[(CO5)(Apply/IOCQ)]*

6 + 6 = 12

9. (a) Consider a file system where free space is maintained by a linked list. Suppose that the pointer to the free-space list is lost. Devise a scheme to reconstruct the free space list. *[(CO6)(Create/HOCQ)]*

(b) Consider a UNIX based file system where a logical block on the file system holds 4 KB, and that a block number is addressable by a 32-bit integer. The inode has 13 pointers for blocks of which 10 points to direct blocks and one each for a single, double and triple indirect blocks. What is the maximum file size supported by this system? *[(CO5)(Apply/IOCQ)]*

(c) Describe the steps in a DMA (Direct Memory Access) transfer. How does DMA increase system concurrency? *[(CO6)(Understand/LOCQ)]*

3 + 3 + 6 = 12

Cognition Level	LOCQ	IOCQ	HOCQ
Percentage distribution	20.83	76.04	3.13