Roll No. | | | | | | | | | | |

B.E (FT) END SEMESTER EXAMINATIONS – MAY/JUNE 2024

Computer Science and Engineering

IVth Semester

CS6108 & OPERATING SYSTEMS

(Regulation 2018 - RUSA)

Time: 3 Hours | Answer ALL Questions | Max. Marks 100

CO 1	Articulate the main concepts, key ideas, strengths and limitations of operating systems
CO 2	Analyze the structure and basic architectural components of OS
CO 3	Elaborate and design various scheduling algorithms
CO 4	Discuss various memory management schemes and design them
CO 5	Point out the various aspects of storage management

BL – Bloom's Taxonomy Levels

(L1 - Remember, L2 - Understand, L3 - Apply, L4 - Analyze, L5 - Evaluate, L6 -Create)

PART-A (10 x 2 = 20 Marks)

Q.No.	Questions	Marks	CO	BL
1	What are the types of multiprocessors, and highlight their advantages.	2	1	L1
2	Sketch the microkernel system structure.	2	2	L1
3	Highlight the benefits achieved by process cooperation.	2	3	L2
4	Sketch the state diagram of process states.	2	3	L1
5	Consider a program that performs three tasks sequentially: Task A, Task B, and Task C. The sequential execution time for each task is as follows: Task A: 10 seconds, Task B: 15 seconds, Task C: 30 seconds. If you parallelize the execution of Task B using two threads, how much time would it take to complete all tasks?	2	3	L3
6	List the programming challenges for multicore systems.	2	2	L1
7	If the average wait time to acquire a mutex lock is 10 ms and the critical section execution time is 5 ms, what is the total time a process spends if it has to wait once before entering the critical section?	2	3	L3
8	Detail the circular wait leading to a deadlock situation.	2	3	L2
9	Under what circumstances do page faults occur? Describe the actions taken by the operating system when a page fault occurs.	2	4	L4
10	A 1000 KB memory is managed using variable partitions but no compaction. It currently has two partitions of sizes 200 KB and 260 KB respectively. What is the smallest allocation request in KB that could be denied?	2	4	L3

PART – B (8 x 8 = 64 marks)
(Answer any 8 questions)

Q.No.	Questions	Marks	CO	BL
11	a) Sketch the storage device hierarchy, highlighting the tradeoffs between storage capacity and access time. (4) b) List the services offered by the operating systems to the users for efficient operation of the system. (4)	8	5 1	L2, L1
12	a) Consider a uniprocessor system executing three tasks T_1, T_2, and T_3, each of which is composed of an infinite sequence of jobs (or instances) which arrive periodically at intervals of 3, 7, and 20 milliseconds, respectively. The priority of each task is the inverse of its period, and the available tasks are scheduled in order of priority, with the highest priority task scheduled first. Each instance of T_1, T_2 & T_3 requires an execution time of 1, 2, and 4 milliseconds, respectively. Given that all tasks initially arrive at the beginning of the 1^{st} millisecond and task preemptions are allowed. Compute the time required for the first instance of T_3 completes its execution at the end. (4) b) Highlight the types of process scheduling queues and sketch the queuing diagram representing process scheduling. (4)	8	3	L3, L2

13	a) List out any two methods to logically implement a link for interprocess communications with their respective send() and receive() operations. (4)
	
	b) Consider two functions *incr* and *decr* shown below:
	
	incr( ) {
	    wait (s);
	    X = X + 1;
	    signal (s);
	}
	
	decr( ) {
	    wait (s);
	    X = X -1;
	    signal (s);
	}
	
	There are 5 threads each invoking *incr* once, and 3 threads each invoking *decr* once, on the same shared variable X. The initial value of X is 10. Suppose there are two implementations of the semaphore s, as follows:
	I-1: s is a binary semaphore initialized to 1.
	I-2: s is a counting semaphore initialized to 2.
	Let V1, V2 be the values of X at the end of execution of all the threads with implementations I-1, I-2, respectively.
	
	What are the minimum possible values of V1 and V2 respectively? (4)
		8	3
			
			3
				L1, L3

<table>
<tr><td>14</td><td>a) Highlight and compare any two relationship exist between user threads and kernel threads. (4)

b) Consider the following processes arriving in the ready queue and sharing the CPU in a preemptive priority scheduling algorithm. Calculate the average waiting time, average turnaround time, and average response time of the processes. Also, obtain the overall throughput. (4)
<table>
<tr><th>Process No.</th><th>Priority</th><th>Arrival Time</th><th>Burst Time</th></tr>
<tr><td>1</td><td>2 (Low)</td><td>0</td><td>8</td></tr>
<tr><td>2</td><td>7</td><td>1</td><td>2</td></tr>
<tr><td>3</td><td>6</td><td>2</td><td>4</td></tr>
<tr><td>4</td><td>10</td><td>3</td><td>5</td></tr>
<tr><td>5</td><td>8</td><td>4</td><td>2</td></tr>
<tr><td>6</td><td>12 (High)</td><td>5</td><td>4</td></tr>
<tr><td>7</td><td>9</td><td>6</td><td>6</td></tr>
</table></td><td>8</td><td>3</td><td>L2, L3</td></tr>
<tr><td>15</td><td>a) Consider the following set of processes, assumed to have arrived at time 0. Consider the CPU scheduling algorithms Shortest Job First (SJF) and Round Robin (RR). For RR, assume that the processes are scheduled in the order P_1, P_2, P_3, P_4.
<table>
<tr><td>Processes</td><td>P1</td><td>P2</td><td>P3</td><td>P4</td></tr>
<tr><td>Burst time (in ms)</td><td>8</td><td>7</td><td>2</td><td>4</td></tr>
</table>
If the time quantum for RR is 4 ms, then compute the absolute value of the difference between the average turnaround times (in ms) of SJF and RR (round off to 2 decimal places). (4)

b) In a system with three processes (P1, P2, P3) sharing a critical section, if each process takes 15 ms to execute in the critical section and the mutex lock overhead is 5 ms, calculate the total time required for P1 to complete its critical section execution when P2 has just exited the critical section. (4)</td><td>8</td><td>3</td><td>L3</td></tr>
<tr><td>16</td><td>a) Highlight the procedures to Implement a semaphore with no busy waiting. (4)

b) Consider a system consisting of processes P_1, P_2, ..., P_n, each of which has a unique priority number. Write a monitor that allocates three identical printers to these processes, using the priority numbers for deciding the order of allocation. (4)</td><td>8</td><td>3</td><td>L2, L3</td></tr>
</table>

17	a) Consider a system consisting of *m* resources of the same type being shared by *n* threads. A thread can request or release only one resource at a time. Show that the system is deadlock free if the following two conditions hold: (4) i. The maximum need of each thread is between one resource and m resources. ii. The sum of all maximum needs is less than $m + n$. b) Draw a resource-allocation-graph for the scenario given below: (4) Set of resources: (R1, R2, R3, R4, R5), Set of processes: (P1, P2, P3, P4, P5) i. R1 is allotted to P2 and requested by P1; R2 is allotted to P1 and requested by P4; ii. R3 is allotted to P5 and requested by P2, R4 is allotted to P3 and requested by P2; Provide the corresponding wait-for-graph and comment on the presence of deadlock.	8	3	L4, L2
18	a) To eliminate deadlocks, we should abort those processes whose termination will incur the minimum cost. List out the factors to select the processes. (4) b) Consider a disk with 300 cylinders. The requests to access the cylinders arrive in the following order: 143, 86, 147, 913, 177, 948, 150, 102, 175, and 130. Suppose the disk arm is initially at cylinder 100. Implement the LOOK disk scheduling algorithm. Calculate the total head movement. (4)	8	3 5	L2, L4
19	a) Why is rotational latency usually not considered in disk scheduling? How would you modify SSTF, SCAN, and C-SCAN to include latency optimization? (4) b) How relocation and memory protection is achieved by hardware support in Contiguous Memory Allocation? (4)	8	5 4	L3, L2
20	a) Consider a computer system with 40-bit virtual addressing and page size of 16 kilobytes. If the computer system has a one-level page table per process and each page table entry requires 48 bits, then what is the size of the per-process page table (in megabytes)? (4) b) Consider a system with a two-level paging scheme in which a regular memory access takes 150 nanoseconds and servicing a page fault takes 8 milliseconds. An average instruction takes 100 nanoseconds of CPU time and two memory accesses. The TLB hit ratio is 90% and the page fault rate is one in every 10,000 instructions. What is the effective average instruction execution time? (4)	8	4	L3, L4
21	a) Sketch and discuss the hardware implementation of paging. (6) b) Consider a system with byte-addressable memory, 32-bit logical address, 4 kilobyte page size and page table entries of 4 bytes each. What is the size of the page table in the system in megabytes? (2)	8	4 4	L2, L3

<table>
<tr><td>22</td><td>a) Sketch the procedure for handling page fault. (4)

b) Consider a demand-paging system with a paging disk that has an average access and transfer time of 20 milliseconds. Addresses are translated through a page table in main memory, with an access time of 1 microsecond per memory access. Thus, each memory reference through the page table takes two accesses. To improve this time, we have added an associative memory that reduces access time to one memory reference if the page-table entry is in the associative memory. Assume that 80 percent of the accesses are in the associative memory and that, of those remaining, 10 percent (or 2 percent of the total) cause page faults. What is the effective memory access time? (4)</td><td>8</td><td>4</td><td>L2, L4</td></tr>
</table>

PART – C (2 x 8 = 16 marks)

<table>
<tr><th>Q. No.</th><th>Questions</th><th>Marks</th><th>CO</th><th>BL</th></tr>
<tr><td>23</td><td>Consider the following snapshot of a system:
<table>
<tr><th></th><th>Allocation</th><th>Max.</th></tr>
<tr><td></td><td>ABCD</td><td>ABCD</td></tr>
<tr><td>T0</td><td>1202</td><td>4316</td></tr>
<tr><td>T1</td><td>0112</td><td>2424</td></tr>
<tr><td>T2</td><td>1240</td><td>3651</td></tr>
<tr><td>T3</td><td>1201</td><td>2623</td></tr>
<tr><td>T4</td><td>1001</td><td>3112</td></tr>
</table>
Using the banker's algorithm, determine whether or not each of the following states is unsafe. If the state is safe, illustrate the order in which the threads may complete. Otherwise, illustrate why the state is unsafe. (8)
i. Available = (2, 2, 2, 3)
ii. Available = (4, 4, 1, 1)
iii. Available = (3, 0, 1, 4)
iv. Available = (1, 5, 2, 2)</td><td>8</td><td>3</td><td>L4</td></tr>
<tr><td>24</td><td>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.
How many page faults would occur for the following replacement algorithms, assuming one, two, three, four, five, six, and seven frames? Remember that all frames are initially empty, so your first unique pages will cost one fault each. (8)
• LRU replacement
• FIFO replacement
• Optimal replacement</td><td>8</td><td>4</td><td>L3</td></tr>
</table>