MASTER OF COMPUTER SCIENCE & ENGINEERING EXAMINATION 2018
1ST SEMESTER

ADVANCED OPERATING SYSTEMS

Time: 3 hours Full Marks: 100

Answer any **four questions** from the following

(All parts of the same question must be answered together)

1.

a. Suppose you are asked to design an operating system for *portable hand-held devices* including smart phones. What characteristics should this operating system have? Justify your answer.

b. How can *monolithic kernel* be redesigned to adapt the modularity feature of microkernel? What are the problems of having servers in user address space, as in microkernel? Compare monolithic kernel with microkernel.

c. What are the contents of Android software architecture? What is the uniqueness of Dalvik Virtual Machine?

d. What is priority inheritance?

4+(3+3+4)+(6+3)+2=25

2. With reference to Unix, answer the following questions:

a. How is a file located in disk after command for opening the file is given by the user?

b. What does in-core inode mean and what information does it have? Why and when may the in-core copy and disk copy of an inode be different? What does the kernel do then?

c. What support does the kernel provide to a process? What happens after *fork* system call? What are the characteristics of a process?

d. What is the difference between *mode switch* and *context switch*?

4+(4+3+2)+(3+2+3)+4=25

3.

a. How are buffer caches organized in Unix? What are the contents of the buffer header?

b. What are the assumptions in Rate Monotonic Scheduling (RMS)? Consider the following set of four tasks: {(1,3), (1,5), (1,6), (2,10)} where the first number in each subset denotes the execution time and the second number denotes the periodicity of each task respectively. Are these tasks schedulable? Justify your answer.

Q3 continued..

c. What is the *One to One threading* model? What are its advantages and disadvantages? Which *scheduling approach* you feel will be appropriate for this model? Why?

d. If no processor is free, then what will happen to a new incoming thread in a system with *Dynamic scheduling*?

(4+3)+(3+5)+(2+3+2)+3=25

4.

a. Consider the following category of process in a real-time system:

Process	Arrival time	Execution time	Deadline (ending)
P1	10	10	30
P2	25		60
P3	40		90
P4	55		120
P5	70		150
--	--	--	--

However, several other processes, as mentioned below, also arrive at this system.

Process	Arrival time	Execution time	Deadline
A	20	10	30 (starting)
B	25	15	70 (ending)
C	50	10	80 (starting)
D	55	5	70 (ending)

(i) What type of process is P1… ? What type of process is A?

(ii) What scheduling approach will you take and why? Mention any assumption that you may take.

(iii) Develop a scheduling strategy and find out whether all processes could be executed successfully. Show the scheduling decisions according to your adopted approach. Justify your answer.

b. Explain the Static Table-driven approach with a suitable example and mention its advantages and problems.

(3+5+10)+7=25

5.

a. How can a remote data object be located using *Migration algorithm* for Distributed Shared Memory (DSM)? What are the problems with this algorithm?

b. What is *write invalidate* protocol? Why may *write update* protocol be difficult to implement?

c. How does writing happen in *Read Replication* algorithm for DSM? In what situation will this algorithm be advantageous?

Q5 continued..

d. What are the semantics for file sharing in Distributed File Systems? How is consistency maintained using *token* in Open Software Foundation DFS?

(3+3)+(3+2)+(3+3)+(4+4)=25

6.

a. What are the functionalities of Master in Google File Systems (GFS)? How are files stored in GFS? How does the Master ensure that a chunk is a valid chunk? What does a *chunkserver* contain? What is *mutation*?

b. In what ways are *selection policy* and *location policy* implemented?

c. Consider the following situations and suggest and justify your choice of the scheduling algorithm for each:

(i) The distributed system is always lightly loaded.

(ii) The distributed system is always heavily loaded.

(iii) The distributed system has a mix of load – sometimes lightly loaded, sometimes heavily loaded.

(4+2+2+2+2)+5+8=25

- - -