# Question Paper

Exam Date & Time: 17-May-2022 (10:00 AM - 01:00 PM)

MANIPAL INSTITUTE OF TECHNOLOGY
MANIPAL
(A constituent unit of MAHE, Manipal)

SIXTH SEMESTER B.TECH END SEMESTER EXAMINATIONS, MAY 2022

**DISTRIBUTED SYSTEMS [ICT 3254]**

**Marks: 50** **Duration: 180 mins.**

**A**

**Answer all the questions.**

Instructions to Candidates:
1. Answer ALL questions.
2. Missing data may be suitably assumed.

| | | | |
|---|---|---|---|
| 1) | A) | A server program written in one language (for example C++) provides the implementation of a BLOB object that is intended to be accessed by clients that may be written in a different language (for example Java). The client and server computers may have different hardware, but all of them are attached to an internet. Describe the problems due to each of the five aspects of heterogeneity that need to be solved to make it possible for a client object to invoke a method on the server object | (5) |
| | B) | Sun NFS aims to support heterogeneous distributed systems by the provision of an operating system-independent file service. What are the key decisions that the implementer of an NFS server for an operating system other than UNIX would have to take? What constraints should an underlying filing system obey to be suitable for the implementation of NFS servers? | (3) |
| | C) | Discuss whether the following operations are idempotent:<br>i) Pressing a lift (elevator) request button<br>ii) Writing data to a file<br>iii) Is it a necessary condition for idempotent that the operation should not be associated with any state? | (2) |
| 2) | A) | Justify the statement "Interface Definition Languages are used by client for accessing the web services." with suitable example. Consider a scenario in a asynchronous distributed systems, there are two communication services X and Y. In service X, messages may be lost, duplicated or delayed and checksums apply only to headers. In service Y, messages may be lost, delayed or delivered too fast for the recipient to handle them, but those that are delivered arrive order and with the correct contents. Discuss the classes of failure experienced by each of these services and which one of these can be reliable? Why? | (5) |
| | B) | Cloud computing services like Amazon's EC2 assign users virtual machines (VMs) instead of allocating physical machines directly. Doing so provides at least three major benefits to Amazon. Explain what these three benefits are giving a brief motivation for each one. | (3) |
| | C) | How cache coherency will be handled in Distributed Shared Memory? | (2) |
| 3) | A) | Two processes *P* and *Q* are connected in a ring using two channels, and they constantly rotate a message *m*. At any one time, there is only one copy of *m* in the system. Each process's state consists of the number of times it has received *m*, and *P* sends *m* first. At a certain point, *P* has the message, and its state is 101. Immediately after sending *m, P* initiates the snapshot algorithm. Explain the operation of the algorithm in this case, giving the possible global state(s) reported by it. The figure shows events occurring for each of two processes, p1 and p2. Arrows between processes denote message transmission. Draw and label the lattice of consistent states (p1 state, p2 state), beginning with the initial state (0,0). | (5) |

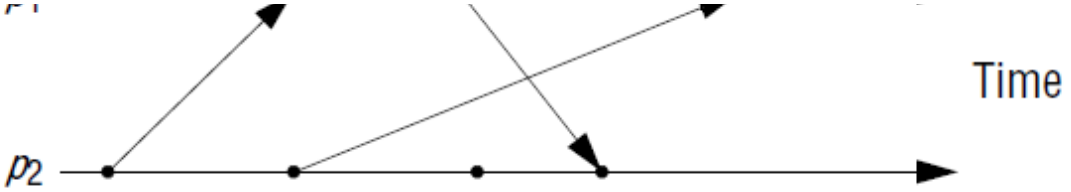


B) A scheme for implementing at-most-once reliable message delivery uses synchronized clocks to reject duplicate messages. Processes place their local clock value (a 'timestamp') in the messages they send. Each receiver keeps a table giving, for each sending process, the largest message timestamp it has seen. Assume that clocks are synchronized to within 100 ms, and that messages can arrive at most 50 ms after transmission. (3)

i) When may a process ignores a message bearing a timestamp T, if it has recorded the last message received from that process as having timestamp T'?

ii) When may a receiver remove a timestamp 175,000 (ms) from its table?

iii) Should the clocks be internally synchronized or externally synchronized?

C) The bank transaction must ensure that a withdrawal operation is executed whenever any user selects that option. Which call semantics is acceptable in this scenario. Justify your answer. (2)

4) A) In a certain system, each process typically uses a critical section many times before another process requires it. Explain why Ricart and Agrawala's multicast-based mutual exclusion algorithm is inefficient for this case and describe how to improve its performance. Does your adaptation satisfy liveness condition ME2 (5)

B) A server manages the objects $a_1$, $a_2$, ..., $a_n$. The server provides two operations for its clients: (3)

read (i) returns the value of $a_i$;

write (i, Value) assigns Value to $a_i$;

The transactions T and U are defined as follows:

T: x = read(j); y = read (i); write (j, 22); write (i, 11);

U: x = read(k); write (i, 33); y = read (j); write (k, 44).

Give three serially equivalent interleavings of the transactions T and U.

C) Explain dirty read and premature write with an example. (2)

5) A) Assuming that strict two-phase locking is in use, describe how the actions of the two-phase commit protocol relate to the concurrency control actions of each individual server. How does distributed deadlock detection fit in? (5)

B) Illustrate Byzantine agreement problem considering four processes where one of them is a faulty process. (3)

C) Google earth has geographical segments and different images are available for that segment. How table abstraction can be achieved in this application to manage the storage and retrieval? (2)

-----End-----