Ref. No.: Ex/PG/SWE/T/128A/2024

ME Software Engineering
First Year, Second Semester Exam 2024
Distributed Software Architecture

Time: Three hours **Full Marks: 100**

Answer any five questions

Q.1
a) What are asynchronous and synchronous message passing systems?
b) Explain the method of finding a DFS spanning tree when the root of the tree is known, and include a neat diagram.
c) What are cycles and knots? How can these knots and cycles affect deadlock detection in distributed systems?
d) Mention two correctness criteria for deadlock detection in distributed systems.
e) Illustrate how Mitchell Merritt's Algorithm detects deadlock in distributed systems? **3+6+4+2+5=20**

Q.2
a) You are synchronizing your clock from a time server using NTP and observe the following times:
   i. timestamp at client when the message leaves the client: 6:22:15.100
   ii. timestamp at which the server receives the message: 7:05:10.700
   iii. timestamp at which the server sends the reply: 7:05:10.710
   iv. timestamp at client when the message is received at client: 6:22:15.250

   To what value do you set the client's clock?
b) What is Lamport's definition of concurrent events using the happened before relation?
c) What are the inherent limitations of Lamport's logical clocks in accurately capturing causality in distributed systems, especially in scenarios involving message delays and concurrent events?
d) How do vector clocks extend the notion of Lamport's logical clocks? Consider the following event diagram for processes P1, P2 and P3 executing in a distributed system. Compute the vector clock that is carried on each message.

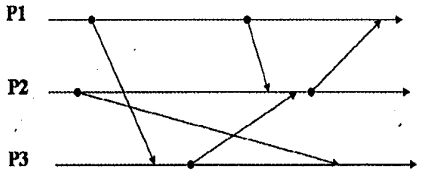

5+3+4+8 =20

Q.3
a) How does a hierarchical namespace improve the scalability and manageability of distributed systems compared to a flat namespace?
b) What are the different methods used for namespace resolution in distributed systems, and how do they compare in terms of efficiency and reliability?
c) How does caching affect the performance of recursive name resolution methods in distributed systems, and what are the potential trade-offs?
d) Briefly describe DNS hierarchy structure. **4+5+5+6=20**

[Turn over

Q.4

a) What is a snapshot state?

b) What is a consistent cut in the context of distributed systems, and why is it important for ensuring system consistency?

c) State the issues that need to be addressed when recording a global state.

d) Give Marker receiving rule and consider a distributed system with three processes P1, P2, and P3, connected as follows:

P1 -> P2
P2 -> P3
P3 -> P1

Apply Chandy Lamport's Snapshot Alogorithm to ensure that a consistent snapshot of the global state got captured. Clearly mention each step of the process. **3+5+4+8=20**

Q.5

a) What are the differences between centralized and distributed mutual exclusion?

b) Describe Lamport's Algorithm for distributed mutual exclusion. Run the algorithm for 4 nodes with all having degree 4. Analyze the algorithm to measure the performance metrics.

c) Describe the limitations of this algorithm and give the optimizations made to enhance it.

5+10+5=20

Q.6

Write short notes on :-

a) Convergecast
b) Authoritative Name Servers
c) The (P/Q) Model of Deadlock
d) Consistent Global State
e) Correctness of Suzuki Kasami's Broadcast Algorithm **5X4=20**