Question Paper

Exam Date & Time: 18-Jun-2024 (02:30 PM - 05:30 PM)

MANIPAL ACADEMY OF HIGHER EDUCATION

SIXTH SEMESTER B.TECH. (INFORMATION TECHNOLOGY) DEGREE EXAMINATIONS - JUNE 2024
SUBJECT: ICT 3254/ICT_3254 - DISTRIBUTED SYSTEMS

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

Answer all the questions.

1A)	How does interoperability ensure seamless communication and collaboration among diverse devices like computers, smartphones, and smart home devices, regardless of brand or type, in a multi-device environment, and what strategies can be employed to effectively overcome interoperability challenges in such a context?	(5)
1B)	Describe how the principles of transparency can be leveraged to overcome challenges in managing smart devices within a home environment to provide seamless operation, efficient resource utilization, and enhanced user experience in the context of smart home automation.	(3)
1C)	What are the key distinctions between Remote Object Reference and Remote Interface in distributed systems, and how do they contribute to enabling remote method invocation?	(2)
2A)	Three distributed systems, P, Q, and R are utilizing Lamport logical clocks for synchronization. At the initial state, all systems start with their logical clocks set to zero. Throughout the process, the following events occur: • System P sends message M1 to Q. • System Q sends message M2 to R. • System R receives M1 before M2. • System R sends message M3 to P. • System P receives M3 before M2. • System Q sends message M4 to P. Illustrate with a time-line diagram and proper stepwise explanation of Lamport timestamp algorithm for the above-mentioned scenario.	(5)
2B)	Discuss how you would choose between soft mounting and hard mounting for configuring NFS mounts on client machines for an environment that experiences occasional connectivity issues due to factors like network congestion or intermittent connectivity. Provide reasons for your choice and explain the potential implications of your decision on system performance and reliability.	(3)
2C)	Calculate the total drift in milliseconds over a 24-hour period for the network time server.	(2)
3A)	Model the central server algorithm for mutual exclusion to handle the crash failure of any client (in any state), assuming that the server is correct and given a reliable failure detector. Comment on whether the resultant system is fault tolerant. What would happen if a client that possesses the token is wrongly suspected to have failed?	(5)
3B)	What are the key components and functionalities of a naming service, and how does it contribute to the efficiency and scalability of distributed systems?	(3)
3C)	Let's consider a distributed system where multiple nodes collaborate to process a large dataset of images for machine learning tasks. Each node is responsible for executing a portion of the machine learning algorithm on a subset of images. The goal is to design a Distributed Shared Memory (DSM) system that efficiently manages shared data while optimizing performance and scalability. Develop the trade-offs if you give a solution when granularity is small.	(2)

4A)	Make use of a Server as an example to illustrate different failure models. Consider a Web browser that returns an outdated cached page instead of a more recent one that had been updated at the server. Is this a failure, and if so, what kind of failure may occur?	(5)
4B)	Explain the significance of the following terminologies in replica management. i) State Transfer. ii) Multicast Communication. iii) Active Replication.	(3)
4C)	Operation conflict rules in timestamp ordering are essential for maintaining transactional correctness. Explain these rules with suitable examples for read-write and write-read conflicts.	(2)
5A)	Consider the following transactions for Two-Phase Commit (2PC) protocol to coordinate transaction commits and ensure data consistency. • Transaction OrderProcessing initiates a nested transaction InventoryUpdate within its scope to fulfill a customer's order by updating product availability and processing shipment details. • During InventoryUpdate, the system performs operations to check product stock levels, reserve items for the order, update inventory records, and confirm shipment logistics across different warehouse databases and inventory systems. • Upon completion of InventoryUpdate, OrderProcessing may continue with additional tasks, such as generating order confirmations and updating customer accounts. • The commit process of both OrderProcessing and InventoryUpdate must be coordinated and executed atomically to ensure that customer orders are processed accurately, and inventory records are updated consistently across all warehouse locations. Compare and contrast the 2PC protocol for flat transactions (single level) and nested transactions (multi-level) within the context of a distributed online shopping platform. What additional complexities arise when managing commit operations for nested transactions, and how does the 2PC protocol mitigate these complexities?	(5)
5B)	Transactions X, Y, and Z share read locks on object A, while transaction W holds a write lock on object B. Transaction Z is waiting to obtain a lock, and then transactions X and W request write locks on object A. Identify the deadlock scenario with the deadlock chain. Determine which transaction should be aborted to resolve the deadlock effectively. Justify your decision based on the dependency chain and potential impact on system operations. Outline the steps involved in resolving the deadlock by aborting the selected transaction. Include the process of releasing locks, notifying affected transactions, and restarting the system to resume transaction processing.	(3)
5C)	Diagrammatically outline the operational workflow of the MapReduce framework.	(2)

-----End-----