

USN 

Internal Assessment Test 3 – August 2024

Sub:	Database Management Systems					Sub Code:	BCS403	Branch:	ISE	
Date:	04/06/2024	Duration:	90 min's	Max Marks:	50	Sem/Sec:	IV A, B & C		OBE	
Answer any FIVE FULL Questions								MARKS	CO	RBT
1	a. What are ACID properties? Explain with an example. b. How do you test conflict serializability of a schedule S? Explain with the help of a graph.					5	5	CO5	L2	
2	Why is concurrency control needed? Demonstrate with an example					10		CO5	L2	
3	Discuss binary locks and unlock operations & shared/ exclusive operations					10		CO5	L2	
4	Explain timestamp ordering algorithm protocol for concurrency control.					10		CO5	L2	
5	a. What is NOSQL? Explain the CAP theorem. b. What is NOSQL Graph database? Explain Neo4j.					5	5	CO6	L2	
6	What are document based NOSQL systems? basic operations CRUD in MongoDB.					10		CO6	L2	

CI

CCI

HOD

USN 

Internal Assessment Test 3 – August 2024

Sub:	Database Management Systems					Sub Code:	BCS403	Branch:	ISE	
Date:	04/06/2024	Duration:	90 min's	Max Marks:	50	Sem/Sec:	IV A, B & C		OBE	
Answer any FIVE FULL Questions								MARKS	CO	RBT
1	a. What are ACID properties? Explain with an example. b. How do you test conflict serializability of a schedule S? Explain with the help of a graph.					5	5	CO5	L2	
2	Why is concurrency control needed? Demonstrate with an example					10		CO5	L2	
3	Discuss binary locks and unlock operations & shared/ exclusive operations					10		CO5	L2	
4	Explain timestamp ordering algorithm protocol for concurrency control.					10		CO5	L2	
5	a. What is NOSQL? Explain the CAP theorem. b. What is NOSQL Graph database? Explain Neo4j.					5	5	CO6	L2	
6	What are document based NOSQL systems? basic operations CRUD in MongoDB.					10		CO6	L2	

CI

CCI

HOD

	<p style="text-align: center;">T1 ↑ ↓ T2</p> <p style="text-align: center;">There is a directed edge from T1 to T2 and from T2 to T1, forming a cycle.</p> <p>3. Analyze the Graph:</p> <ul style="list-style-type: none"> ○ Cycle Detection: The graph contains a cycle (T1 → T2 → T1). 			
2	<p>Why is concurrency control needed? Demonstrate with an example</p> <p>Concurrency control is necessary to ensure the consistency and isolation properties of transactions. (1M)</p> <p>Without it, the following issues can occur:</p> <ul style="list-style-type: none"> • Lost Updates: When two transactions read the same data and then update it, one update can be lost. (1M) • Dirty Read: Happens when a transaction reads data that has been written by another transaction that is not yet committed, leading to potential inconsistencies. (1M) • Incorrect Summary: Arises when a transaction calculates an aggregate summary of data while another transaction is updating some of that data, resulting in an inaccurate summary.. (1M) • Unrepeatable Read: Happens when a transaction reads the same data twice and gets different results each time due to another transaction modifying the data in between the reads.. (1M) <p>Examples of all issues. (5M)</p>	10	CO5	L2
3	<p>Discuss binary locks and unlock operations & shared/ exclusive operations (5M)</p> <p>Binary Locks:</p> <ol style="list-style-type: none"> 1. Definition: <ul style="list-style-type: none"> ○ A binary lock is a simple type of lock that can be in one of two states: locked (1) or unlocked (0). 2. Operations: <ul style="list-style-type: none"> ○ Lock (Acquire): This operation sets the lock on a data item. If the data item is already locked, the transaction requesting the lock must wait until it is released. ○ Unlock (Release): This operation releases the lock on a data item, allowing other transactions to acquire the lock. 3. Behavior: <ul style="list-style-type: none"> ○ When a transaction holds a binary lock on a data item, no other transaction can access that data item until the lock is released. ○ Binary locks provide a straightforward mechanism for ensuring mutual exclusion, where only one transaction can access a data item at a time. 4. Example: <ul style="list-style-type: none"> ○ Suppose Transaction T1 locks data item A using a binary lock. No other transaction can read or write to A until T1 releases the lock. <p>Shared/Exclusive (S/X) Locks</p> <p>Shared (S) Locks: (3M)</p> <ol style="list-style-type: none"> 1. Definition: <ul style="list-style-type: none"> ○ A shared lock allows multiple transactions to read a data item but prevents any transaction from writing to it. 2. Operations: <ul style="list-style-type: none"> ○ Acquire Shared Lock: A transaction can acquire a shared lock on a data item if no other transaction holds an exclusive lock on the same data item. ○ Release Shared Lock: The transaction releases the shared lock after it completes reading the data item. 3. Behavior: <ul style="list-style-type: none"> ○ Multiple transactions can hold shared locks on the same data item simultaneously. ○ Shared locks are used to enable concurrent read operations without interference. 4. Example: <ul style="list-style-type: none"> ○ Transactions T1 and T2 can both hold shared locks on data item A and read it concurrently. <p>Exclusive (X) Locks: (2M)</p> <ol style="list-style-type: none"> 1. Definition: <ul style="list-style-type: none"> ○ An exclusive lock allows a transaction to read and write a data item, but no other transaction can access the data item while the exclusive lock is held. 2. Operations: 	10	CO5	L2

- **Acquire Exclusive Lock:** A transaction can acquire an exclusive lock on a data item if no other transaction holds either a shared or exclusive lock on the same data item.
 - **Release Exclusive Lock:** The transaction releases the exclusive lock after it completes its operations on the data item.
3. **Behavior:**
- Only one transaction can hold an exclusive lock on a data item at a time.
 - Exclusive locks provide mutual exclusion for both read and write operations.
4. **Example:**
- If Transaction T1 holds an exclusive lock on data item A, no other transaction can read or write to A until T1 releases the lock.

Lock Compatibility Matrix

The compatibility matrix helps in determining whether a transaction can acquire a lock based on the locks held by other transactions:

Compatibility	S Lock	X Lock
S Lock	Yes	No
X Lock	No	No

- **S Lock:** If a data item has a shared lock, other transactions can also acquire a shared lock on it.
- **X Lock:** If a data item has an exclusive lock, no other transaction can acquire any type of lock on it.

4 **Explain timestamp ordering algorithm protocol for concurrency control.**

The Timestamp Ordering (TO) protocol is used in database systems to ensure that transactions are executed in a serializable order based on their timestamps. This protocol assigns a unique timestamp to each transaction when it starts, and this timestamp determines the order in which the transactions must be executed. **(2M)**

Key Concepts **(2M)**

1. **Timestamps:**
 - Each transaction T_i is given a unique timestamp $TS(T_i)$ when it starts.
 - These timestamps are used to order the transactions. If $TS(T_i) < TS(T_j)$, then T_i must appear to execute before T_j .
2. **Read and Write Timestamps:**
 - Each data item QQQ in the database has two timestamps:
 - **WTS(Q):** The timestamp of the last transaction that wrote to QQQ .
 - **RTS(Q):** The timestamp of the last transaction that read QQQ .

Protocol Rules **(4M)**

The TO protocol ensures serializability by enforcing the following rules when a transaction T_i issues a read or write operation on a data item QQQ :

1. **Read Operation:**
 - When T_i issues a read(Q), the protocol checks if $TS(T_i) < WTS(Q)$ or $TS(T_i) < RTS(Q)$:
 - If $TS(T_i) < WTS(Q)$, it means that T_i is trying to read a value that has been overwritten by a later transaction T_j (where $TS(T_j) > TS(T_i)$). Hence, T_i 's read operation is rejected, and T_i is aborted.
 - If $TS(T_i) \geq WTS(Q)$ and $TS(T_i) \geq RTS(Q)$, the read operation is allowed to proceed. Additionally, $RTS(Q)$ is updated to $\max(RTS(Q), TS(T_i))$.
2. **Write Operation:**
 - When T_i issues a write(Q), the protocol checks two conditions:
 - If $TS(T_i) < RTS(Q)$, it means that QQQ has already been read by a later transaction T_j (where $TS(T_j) > TS(T_i)$). Allowing T_i to write QQQ would violate the serializability. Hence, T_i 's write operation is rejected, and T_i is aborted.
 - If $TS(T_i) < WTS(Q)$, it means that QQQ has already been written by a later transaction T_j (where $TS(T_j) > TS(T_i)$). Allowing T_i to write QQQ would also violate the serializability. Hence, T_i 's write operation is rejected, and T_i is aborted.

	<p style="text-align: right;">▪ If both conditions are false, the write operation is allowed to proceed, and $WTS(Q)WTS(Q)WTS(Q)$ is updated to $TS(T_i)TS(T_i)TS(T_i)$. (2M)</p> <p>Example Consider three transactions T1T1T1, T2T2T2, and T3T3T3 with timestamps $TS(T1)=1$, $TS(T2)=2$, and $TS(T3)=3$. $TS(T1)=1$, $TS(T2)=2$, and $TS(T3)=3$.</p> <ol style="list-style-type: none"> Initial State: <ul style="list-style-type: none"> Data item QQQ has $WTS(Q)=0$ and $RTS(Q)=0$. Transaction Operations: <ul style="list-style-type: none"> T1T1T1 issues write(Q): <ul style="list-style-type: none"> Since $TS(T1)=1$ and both conditions $TS(T1) < RTS(Q)$ and $TS(T1) < WTS(Q)$ are false, T1T1T1 writes QQQ, and $WTS(Q)$ is updated to 1. T2T2T2 issues read(Q): <ul style="list-style-type: none"> Since $TS(T2)=2$ and $TS(T2) \geq WTS(Q)$, T2T2T2 reads QQQ, and $RTS(Q)$ is updated to 2. T3T3T3 issues write(Q): <ul style="list-style-type: none"> Since $TS(T3)=3$ and both conditions $TS(T3) < RTS(Q)$ and $TS(T3) < WTS(Q)$ are false, T3T3T3 writes QQQ, and $WTS(Q)$ is updated to 3. <p>Advantages</p> <ul style="list-style-type: none"> Ensures Serializability: The TO protocol ensures that the transactions are executed in a serializable order based on their timestamps. Avoids Deadlocks: Since transactions are ordered by their timestamps, there are no circular wait conditions, thus avoiding deadlocks. <p>Disadvantages</p> <ul style="list-style-type: none"> Aborts: The protocol can lead to a high number of transaction aborts, especially in systems with high contention. Overhead: Maintaining and checking timestamps can introduce overhead in the transaction processing system. 			
5	<p>c. What is NOSQL? Explain the CAP theorem.</p> <p>NoSQL refers to a broad class of database management systems that differ from traditional relational database management systems (RDBMS). These databases are designed to handle large volumes of data, provide flexible data models, and enable horizontal scaling.</p> <p>Key Characteristics of NoSQL: (1M)</p> <ol style="list-style-type: none"> Schema-less Data Model: <ul style="list-style-type: none"> Unlike relational databases, NoSQL databases do not require a fixed schema, allowing for more flexibility in how data is stored and managed. Horizontal Scalability: <ul style="list-style-type: none"> NoSQL databases are designed to scale out by adding more servers, making them suitable for handling massive amounts of data and high-traffic applications. Variety of Data Models: <ul style="list-style-type: none"> NoSQL databases support various data models, including document, key-value, column-family, and graph models, each optimized for different types of applications. High Performance: <ul style="list-style-type: none"> These databases are optimized for high read/write performance, often at the cost of strict consistency. <p>Types of NoSQL Databases:</p> <ol style="list-style-type: none"> Document Stores: <ul style="list-style-type: none"> Example: MongoDB, CouchDB Store data in JSON-like documents. Key-Value Stores: <ul style="list-style-type: none"> Example: Redis, DynamoDB Store data as key-value pairs. Column-Family Stores: <ul style="list-style-type: none"> Example: Cassandra, HBase Store data in columns rather than rows. Graph Databases: <ul style="list-style-type: none"> Example: Neo4j, OrientDB Store data in graph structures with nodes and edges. 	5	CO6	L2

CAP Theorem (3M)
 CAP Theorem (Consistency, Availability, Partition Tolerance):
 The CAP theorem, also known as Brewer's theorem, states that a distributed database system can only achieve at most two of the following three guarantees simultaneously:

1. Consistency (C):
 - Every read receives the most recent write or an error.
 - Ensures that all nodes see the same data at the same time.
2. Availability (A):
 - Every request receives a response, without guarantee that it contains the most recent write.
 - Ensures that the database system is always operational and responds to requests.
3. Partition Tolerance (P):
 - The system continues to operate despite network partitions or communication breakdowns.
 - Ensures that the system can handle failures or partitions in the network that prevent some nodes from communicating with others.

CAP Theorem Implications: (1M)

- In a distributed system, when a network partition occurs, the system must choose between consistency and availability.
- CA (Consistency and Availability): These systems ensure consistency and availability but cannot tolerate partitions. They are generally centralized systems rather than distributed.
- CP (Consistency and Partition Tolerance): These systems ensure consistency and can tolerate partitions but may sacrifice availability. Example: HBase.
- AP (Availability and Partition Tolerance): These systems ensure availability and can tolerate partitions but may sacrifice consistency. Example: Cassandra, DynamoDB.

d. What is NOSQL Graph database? Explain Neo4j. (1M)

NoSQL Graph Database: (1M)
 A NoSQL graph database is a type of database designed to handle data structured as a graph, with nodes (entities) and edges (relationships). Graph databases are optimized for traversing and querying these relationships, making them particularly suited for applications with interconnected data, such as social networks, recommendation systems, and network analysis.

Key Characteristics of Graph Databases: (2M)
 Nodes and Edges:
 Nodes represent entities (e.g., people, products) and edges represent relationships between these entities (e.g., friendships, purchases).
 Properties:
 Both nodes and edges can have properties (key-value pairs) that store relevant information.
 Schema-less:
 Similar to other NoSQL databases, graph databases are typically schema-less, allowing for flexible and dynamic data models.
 Efficient Relationship Traversal:
 Graph databases are optimized for traversing relationships, which makes complex queries about connections and paths very efficient.
 Example of a NoSQL Graph Database: Neo4j
 Neo4j:
 Neo4j is one of the most popular and widely used graph databases. It is an ACID-compliant transactional database with native graph storage and processing.

Key Features of Neo4j: (2M)
 Native Graph Storage:
 Neo4j uses a native graph storage format, meaning that it stores data as graphs directly on disk, which optimizes for graph-related operations.
 Cypher Query Language:
 Neo4j uses Cypher, a powerful declarative query language specifically designed for querying and updating graph data.
 ACID Compliance:
 Neo4j ensures transactional integrity with ACID properties (Atomicity, Consistency, Isolation, Durability).
 High Performance:
 Neo4j is optimized for high-performance graph traversal and query execution.
 Scalability:
 Neo4j can scale horizontally by sharding the graph across multiple nodes and vertically by increasing the resources of the existing nodes.

6	<p>What are document based NOSQL systems? basic operations CRUD in MongoDB.</p> <p>Document-Based NoSQL Systems: (5M) Document-based NoSQL systems store data in documents, typically using formats like JSON, BSON (Binary JSON), or XML. Each document is a self-contained unit of data that encapsulates fields and values, allowing for a flexible and hierarchical data structure. Unlike traditional relational databases, document-based systems do not require a fixed schema, making them highly adaptable to changing data models.</p> <p>Key Characteristics of Document-Based NoSQL Systems:</p> <ol style="list-style-type: none"> 1. Schema Flexibility: <ul style="list-style-type: none"> ○ Documents can have varying structures, allowing for different fields and data types within the same collection. 2. Hierarchical Data Representation: <ul style="list-style-type: none"> ○ Documents can represent complex nested data structures, which can directly map to objects in application code. 3. Scalability: <ul style="list-style-type: none"> ○ These systems are designed to scale horizontally by distributing data across multiple servers. 4. Ease of Use: <ul style="list-style-type: none"> ○ Documents are easy to read and understand, making it simpler for developers to interact with the database. <p>Popular Document-Based NoSQL Databases:</p> <ul style="list-style-type: none"> • MongoDB • CouchDB • Amazon DocumentDB • ArangoDB <p>MongoDB: MongoDB is one of the most widely used document-based NoSQL databases. It stores data in BSON format and provides powerful querying capabilities, indexing, and aggregation.</p> <p>Basic CRUD Operations in MongoDB</p> <p>CRUD stands for Create, Read, Update, and Delete. (1M) These are the four basic operations that can be performed on data in a database.</p> <p>1. Create (Insert): (1M) The insertOne and insertMany methods are used to add documents to a collection. Example: # Insert a single document db.collection.insertOne({ "name": "Alice", "age": 25, "city": "New York" })</p> <p># Insert multiple documents db.collection.insertMany([{ "name": "Bob", "age": 30, "city": "Chicago" }, { "name": "Charlie", "age": 35, "city": "San Francisco" }])</p> <p>2. Read (Query): (1M) The find method is used to retrieve documents from a collection. It can take a query object to filter results. Example: # Find all documents db.collection.find()</p> <p># Find documents with a specific condition db.collection.find({ "age": { "\$gt": 30 } })</p> <p># Find a single document db.collection.findOne({ "name": "Alice" })</p> <p>3. Update: (1M) The updateOne, updateMany, and replaceOne methods are used to modify existing documents. Example: # Update a single document db.collection.updateOne({ "name": "Alice" }, { "\$set": { "age": 26 } })</p> <p># Update multiple documents db.collection.updateMany({ "city": "Chicago" }, { "\$set": { "city": "Boston" } })</p> <p># Replace a document db.collection.replaceOne({ "name": "Charlie" }, { "name": "Charlie", "age": 36, "city": "Los Angeles" })</p>	10	CO6	L2
---	---	----	-----	----

	<p>4. Delete: (1M) The deleteOne and deleteMany methods are used to remove documents from a collection. Example: # Delete a single document db.collection.deleteOne({ "name": "Alice" }) # Delete multiple documents db.collection.deleteMany({ "city": "Boston" })</p>			
--	---	--	--	--

CI

CCI

HOD