

--	--	--	--	--	--	--	--	--	--	--

Internal Assessment Test –III, October 2024				
Sub:	Database Management System	Code:	22MCA21	
Answer Any 5 QUESTIONS		Marks	OBE	
			CO	RBT
1)	Discuss various Transaction states with a neat diagram and its pros and cons	10	CO3	L2
OR				
2)	Explain various types of failure that may occur in a system.	10	CO3	L2
3)	Explain types of problems that may occur when two simple transaction run concurrently with examples. Why concurrency control and recovery are needed in DBMS?	10	CO3	L1
OR				
4)	Explain properties, characteristics, advantages and disadvantages of a transaction in detail.	10	CO3	L3
5)	Explain the select and project operation with syntax and examples.	10	CO4	L2
OR				
6)	Explain Union, intersection and minus operation with examples	10	CO4	L3
7)	Consider the following COMPANY database EMP(Name,SSN,Salary,address, SuperSSN,Gender,Dno) DEPT(DNum,Dname,MgrSSN) PROJECT(Pname,Pnumber,Plocation,Dnum) Write the relational algebra queries for the following (i) Retrieve the name, address, salary of employees who work for the Research department. (ii) Find the names of employees who work on all projects controlled by department number 4. iii) Retrieve the SSN of all employees who either in department no :4 or directly supervise an employee who work in dno 4	10	CO4	L3
OR				
8)	Explain briefly about database application development	10	CO4	L2
9)	Explain in detailed about i.)JDBC ii.) Cursors	10	CO4	L3
OR				
10)	Explain in detailed about i.)Embedded SQL ii.) SQLJ	10	CO4	L3

1	<p><b>Key Transaction States in DBMS</b></p> <ol style="list-style-type: none"> <li>1. <b>Active State:</b> <ul style="list-style-type: none"> <li>○ <b>Description:</b> The transaction starts here, performing its operations (e.g., read/write).</li> <li>○ <b>Pro:</b> Allows flexibility to perform all operations.</li> <li>○ <b>Con:</b> No guarantee of completion; may end abruptly if an error occurs.</li> </ul> </li> <li>2. <b>Partially Committed State:</b> <ul style="list-style-type: none"> <li>○ <b>Description:</b> Transaction has completed all operations but not yet committed.</li> <li>○ <b>Pro:</b> Ensures that checks can occur before commitment.</li> <li>○ <b>Con:</b> Can lead to a failed state if checks fail.</li> </ul> </li> <li>3. <b>Committed State:</b> <ul style="list-style-type: none"> <li>○ <b>Description:</b> Changes are saved permanently; the transaction is successful.</li> <li>○ <b>Pro:</b> Guarantees durability of transaction results.</li> <li>○ <b>Con:</b> No changes can be reverted without a new transaction.</li> </ul> </li> <li>4. <b>Failed State:</b> <ul style="list-style-type: none"> <li>○ <b>Description:</b> An error occurred; transaction cannot continue.</li> <li>○ <b>Pro:</b> Prevents inconsistent data from entering the database.</li> <li>○ <b>Con:</b> Requires re-execution if a failure occurs due to recoverable issues.</li> </ul> </li> <li>5. <b>Aborted State:</b> <ul style="list-style-type: none"> <li>○ <b>Description:</b> Transaction has been rolled back, undoing all changes.</li> <li>○ <b>Pro:</b> Restores database consistency by discarding changes.</li> <li>○ <b>Con:</b> May involve performance overhead during rollback.</li> </ul> </li> </ol>		
2	<p><b>Failure Type Impact Recovery Mechanism</b></p> <p>Transaction Failure: Incomplete transaction Rollback using logs</p> <p>System Failure Lost in-memory data Checkpoints, transaction logs</p> <p>Media Failure Permanent data loss Restore from backups or replication</p> <p>Application Failure Inconsistent or incomplete transactions Rollback incomplete transactions</p> <p>Network Failure Distributed transaction inconsistency Distributed recovery protocols</p> <p>Buffer Overflow Possible data corruption Memory management adjustments, rollback</p>		
3	<p><b>1. Lost Update Problem</b></p> <ul style="list-style-type: none"> <li>• <b>Description:</b> Occurs when two transactions update the same data item simultaneously, causing one update to overwrite the other.</li> <li>• <b>Example:</b> <ul style="list-style-type: none"> <li>○ Assume a bank account balance is initially \$500.</li> <li>○ Transaction T1T_1T1 reads the balance, adds \$100, and sets it to \$600.</li> <li>○ Before T1T_1T1 writes \$600 back, Transaction T2T_2T2 reads the same balance of \$500, subtracts \$50, and sets it to \$450.</li> <li>○ Both transactions complete, but T1T_1T1's update to \$600 is overwritten by T2T_2T2's update to \$450, causing the addition of \$100 to be "lost."</li> </ul> </li> </ul> <p><b>2. Dirty Read (Uncommitted Dependency Problem)</b></p>		

- **Description:** Occurs when a transaction reads data modified by another transaction that has not yet been committed. If the modifying transaction rolls back, the reading transaction has read invalid data.
- **Example:**
  - T1T\_1T1 updates an account balance from \$500 to \$600 but does not commit.
  - T2T\_2T2 reads this uncommitted balance of \$600 and uses it for calculations.
  - If T1T\_1T1 rolls back, the original balance of \$500 is restored, but T2T\_2T2 has already used the invalid \$600 balance.

### 3. Unrepeatable Read (Inconsistent Retrieval)

- **Description:** Occurs when a transaction reads the same data multiple times and gets different results due to another transaction's updates.
- **Example:**
  - T1T\_1T1 reads the balance of an account as \$500.
  - T2T\_2T2 updates the balance to \$600 and commits.
  - T1T\_1T1 reads the balance again and sees \$600, which is inconsistent with the initial read.

### 4. Phantom Read

- **Description:** Occurs when a transaction reads a set of rows based on a condition, but another transaction inserts, deletes, or updates rows that affect the result of the original query.
- **Example:**
  - T1T\_1T1 reads a list of accounts with a balance greater than \$500, finding three accounts.
  - Meanwhile, T2T\_2T2 inserts a new account with a balance of \$700.
  - If T1T\_1T1 re-reads the list, it will find four accounts, leading to inconsistencies.

## Why Concurrency Control and Recovery are Needed in DBMS

### 1. Concurrency Control:

- Concurrency control mechanisms manage the execution of transactions to ensure data integrity and isolation. They prevent conflicts like lost updates, dirty reads, unrepeatable reads, and phantom reads by coordinating access to shared data.
- **Reason:** Concurrency control is essential to maintain the ACID properties (especially isolation and consistency) and avoid data inconsistencies. It enables multiple users to access the database concurrently without interference.

### 2. Recovery:

- Recovery mechanisms restore the database to a consistent state in the event of system crashes, transaction failures, or hardware issues. This includes rolling back incomplete transactions or redoing committed ones.
- **Reason:** Recovery mechanisms ensure atomicity and durability, allowing transactions to either complete entirely or leave no trace in case of a failure. This maintains database consistency and helps recover from unexpected issues.

4 A transaction in a Database Management System (DBMS) is a logical unit of work that includes one or more operations on the database (such as reading, writing, updating, and deleting data). It ensures data consistency, reliability, and integrity, especially in multi-user environments where several transactions might occur simultaneously. Here's a detailed explanation of its properties, characteristics, advantages, and disadvantages:

### Properties of a Transaction (ACID Properties)

Transactions follow the **ACID** properties, which are essential to maintaining database integrity:

1. **Atomicity:**

- **Definition:** A transaction is an atomic unit of operation, meaning it either fully completes or does not happen at all. If any part of the transaction fails, the entire transaction is rolled back.
- **Explanation:** If a transaction involves multiple steps (e.g., transferring funds between accounts), either all steps must succeed, or none should take effect.
- **Example:** In an e-commerce system, if a customer's payment fails, the entire transaction (including inventory update and order creation) should be reversed.

2. **Consistency:**

- **Definition:** A transaction must bring the database from one consistent state to another. All integrity constraints must be satisfied both before and after the transaction.
- **Explanation:** Database rules (like foreign keys or constraints) ensure that data remains valid. If any condition fails, the transaction will not complete.
- **Example:** A transaction should not result in a negative balance if the system disallows negative account balances.

3. **Isolation:**

- **Definition:** Transactions should be executed independently and should not interfere with each other. A transaction's intermediate states should not be visible to other transactions.
- **Explanation:** This property prevents issues like dirty reads, unrepeatable reads, and lost updates, maintaining the accuracy of concurrent transactions.
- **Example:** In a banking application, two simultaneous withdrawals from the same account should occur in isolation to avoid incorrect balance calculations.

4. **Durability:**

- **Definition:** Once a transaction is committed, its changes to the database are permanent, even in the event of a system crash or failure.
- **Explanation:** Durability ensures that the database maintains a record of committed transactions, usually through logging and backup mechanisms.
- **Example:** If a power failure occurs right after a transaction is committed, the changes should remain when the system restarts.

### Characteristics of Transactions

1. **Consistency Preservation:** Transactions uphold database integrity constraints, ensuring consistency before and after each transaction.
2. **Concurrency Control:** Transactions are designed to operate in environments where multiple users may access the database

simultaneously. Concurrency control techniques, such as locking or timestamp ordering, are used to maintain isolation.

3. **Recoverability:** A transaction should be able to be rolled back in case of failure, and committed transactions should be persistent. This helps maintain the atomicity and durability of transactions.
4. **Transparency:** The underlying database system handles transactions so that the user only sees the end result, ensuring a seamless user experience regardless of the complexity of the transaction.

### **Advantages of Transactions**

1. **Data Integrity and Consistency:**
  - Transactions help maintain data accuracy and consistency, even in cases of simultaneous updates by multiple users. By enforcing the ACID properties, transactions help protect against data corruption.
2. **Error Handling and Recovery:**
  - Transactions make it easier to handle errors. If any part of a transaction fails, the entire transaction can be rolled back, preventing partial updates that could lead to inconsistencies.
3. **Concurrent Access:**
  - By providing isolation and using concurrency control, transactions allow multiple users to interact with the database simultaneously without conflicting with each other's operations.
4. **Simplifies Complex Operations:**
  - Transactions allow complex sequences of operations to be managed easily, enabling developers to define all or nothing operations that make application development simpler and more reliable.
5. **Guaranteed Durability:**
  - Transactions ensure that once a change is committed, it is permanent. This provides reliability for users, who can trust that completed transactions will not be lost due to system failures.

### **Disadvantages of Transactions**

1. **Performance Overhead:**
  - Implementing transaction management, especially isolation and durability, adds overhead, as mechanisms like locking, logging, and checkpoints are resource-intensive. This can slow down system performance, especially under high concurrency.
2. **Deadlocks:**
  - Concurrency control mechanisms (like locking) can lead to deadlocks, where transactions wait indefinitely for each other to release resources, requiring additional management techniques to resolve.
3. **Increased Complexity:**
  - The algorithms required to manage transactions and enforce the ACID properties are complex, making the DBMS more challenging to implement and maintain, which can increase development time and cost.
4. **Resource Consumption:**
  - Transactions require additional resources, such as memory and storage, to manage logs, backups, and checkpoints. In high-volume transaction environments, this can strain system resources and impact scalability.
5. **Reduced Parallelism:**
  - High levels of isolation (like serializability) restrict parallel

execution of transactions, impacting performance in systems with large numbers of concurrent users.

5 The **Select** and **Project** operations are fundamental operations in relational algebra used to manipulate and query relational databases. Here's a detailed explanation, including syntax and examples:

### 1. Select Operation ( $\sigma$ )

- **Purpose:** The select operation retrieves rows (tuples) from a relation (table) that satisfy a specified condition.
- **Symbol:**  $\sigma$  (sigma)
- **Syntax:**

$\sigma_{condition}(Relation)$   
 $\sigma_{\text{condition}}(\text{Relation})$

Where `condition` is the criterion used to filter rows from `Relation`.

- **Example:** Suppose we have a relation (table) called `Employee` with the following schema and data:

EmpID	Name	Department	Salary
-------	------	------------	--------

101	Alice	HR	5000
102	Bob	Sales	6000
103	Charlie	IT	7000
104	David	Sales	5500
105	Eve	IT	7500

- **Query:** Retrieve all employees in the Sales department.
- **Select Operation:**
- $\sigma_{Department = 'Sales'}(Employee)$   
 $\sigma_{\text{Department} = \text{'Sales'}}(\text{Employee})$
- **Result:**

EmpID	Name	Department	Salary
-------	------	------------	--------

102	Bob	Sales	6000
104	David	Sales	5500

- **Explanation:** The select operation filtered out only those rows where the department is 'Sales'.

### 2. Project Operation ( $\pi$ )

- **Purpose:** The project operation retrieves specific columns (attributes) from a relation, effectively creating a subset of columns in a new relation.
- **Symbol:**  $\pi$  (pi)
- **Syntax:**

$\pi_{column1, column2, \dots}(Relation)$   
 $\pi_{\text{column1, column2, \dots}}(\text{Relation})$

	<p>...}}(\text{Relation})\pi\text{column1, column2, ...}(\text{Relation})</p> <p>Where <code>column1</code>, <code>column2</code>, ... are the attributes you want to retrieve from <code>Relation</code>.</p> <ul style="list-style-type: none"> <li><b>Example:</b> Using the same <code>Employee</code> table, let's say we want to retrieve only the names and salaries of all employees.</li> </ul> <p><b>Project Operation:</b></p> $\pi_{\text{Name, Salary}}(\text{Employee})$ <p><b>Result:</b></p> <table border="1"> <thead> <tr> <th>Name</th> <th>Salary</th> </tr> </thead> <tbody> <tr> <td>Alice</td> <td>5000</td> </tr> <tr> <td>Bob</td> <td>6000</td> </tr> <tr> <td>Charlie</td> <td>7000</td> </tr> <tr> <td>David</td> <td>5500</td> </tr> <tr> <td>Eve</td> <td>7500</td> </tr> </tbody> </table> <p><b>Explanation:</b> The project operation selected only the <code>Name</code> and <code>Salary</code> columns, discarding other attributes.</p>	Name	Salary	Alice	5000	Bob	6000	Charlie	7000	David	5500	Eve	7500		
Name	Salary														
Alice	5000														
Bob	6000														
Charlie	7000														
David	5500														
Eve	7500														
6	<p>In relational algebra, <b>Union</b>, <b>Intersection</b>, and <b>Minus</b> are set operations used to combine or differentiate between two relations (tables) with the same schema. Here's an explanation of each operation with examples.</p> <hr/> <p><b>1. Union Operation ( <math>\cup</math> )</b></p> <ul style="list-style-type: none"> <li><b>Purpose:</b> Combines the results of two relations and returns all unique rows present in either or both relations.</li> <li><b>Symbol:</b> <math>\cup</math></li> <li><b>Syntax:</b></li> </ul> $\text{Relation1} \cup \text{Relation2}$ <p>Both <code>Relation1</code> and <code>Relation2</code> must have the same schema (same number of columns with matching data types).</p> <ul style="list-style-type: none"> <li><b>Example:</b> Suppose we have two tables <code>Employee_A</code> and <code>Employee_B</code>:</li> </ul> <p><b>Employee_A:</b></p> <table border="1"> <thead> <tr> <th>EmpID</th> <th>Name</th> </tr> </thead> <tbody> <tr> <td>101</td> <td>Alice</td> </tr> <tr> <td>102</td> <td>Bob</td> </tr> <tr> <td>103</td> <td>Charlie</td> </tr> </tbody> </table>	EmpID	Name	101	Alice	102	Bob	103	Charlie						
EmpID	Name														
101	Alice														
102	Bob														
103	Charlie														

**Employee\_B:****EmpID Name**

103 Charlie  
104 David  
105 Eve

**Union Operation:**

$Employee\_A \cup Employee\_B$   
 $\text{\text{Employee\_A}} \cup \text{\text{Employee\_B}}$

**Result:****EmpID Name**

101 Alice  
102 Bob  
103 Charlie  
104 David  
105 Eve

**Explanation:** The union operation returns all unique rows from both tables, combining the data without duplicates.

**2. Intersection Operation (  $\cap$  )**

- **Purpose:** Returns only the rows that are present in both relations.
- **Symbol:**  $\cap$
- **Syntax:**

$Relation1 \cap Relation2$   
 $\text{\text{Relation1}} \cap \text{\text{Relation2}}$

As with union, both `Relation1` and `Relation2` must have the same schema.

- **Example:** Using the same tables `Employee_A` and `Employee_B`:

**Intersection Operation:**

$Employee\_A \cap Employee\_B$   
 $\text{\text{Employee\_A}} \cap \text{\text{Employee\_B}}$

**Result:****EmpID Name**

103 Charlie

**Explanation:** The intersection operation returns only the rows that appear in both tables.



### 3. Minus Operation ( - )

- **Purpose:** Returns rows that are present in the first relation but not in the second.
- **Symbol:** -
- **Syntax:**

Relation1 - Relation2

As with the other two operations, both Relation1 and Relation2 must have the same schema.

- **Example:** Using the same tables Employee\_A and Employee\_B:

#### Minus Operation:

Employee\_A - Employee\_B

#### Result:

##### EmpID Name

101 Alice  
102 Bob

**Explanation:** The minus operation returns only the rows that are in Employee\_A but not in Employee\_B.

#### 7. Employees in Research Department:

$\pi_{Name, address, Salary}(\sigma_{Dno=DNum(EMP \times DEPT) \wedge Dname = 'Research'})$   
 $\pi_{Name, address, Salary}(\sigma_{Dno = DNum}(\text{EMP} \times \text{DEPT}))$   
 $\setminus \text{and } \text{Dname} = 'Research')$   
 $\pi_{Name, address, Salary}(\sigma_{Dno=DNum(EMP \times DEPT) \wedge Dname = 'Research'})$

#### Employees working on all projects in department 4:

$\pi_{Name}(EMP) \text{ where } SSN \in (WORKS\_ON \div P_4)$   
 $\pi_{Name}(EMP) \text{ where } SSN \in (WORKS\_ON \div P_4)$

#### Employees in department 4 or supervising someone in department 4:

$E_{Dept4} \cup S_{Dept4}$

8

- **Embedded SQL** is great for static queries in host languages.
- **Dynamic SQL** offers flexibility for runtime query generation.
- **JDBC** provides a standard API for Java applications to interact with databases.
- **Stored Procedures** encapsulate complex logic in the database, improving performance and security.
- **SQLJ** combines the benefits of SQL with Java, offering type safety and integration.

<p>9</p>	<ul style="list-style-type: none"> <li>• <b>DBC Driver:</b> A software component that enables Java applications to interact with a specific database. JDBC drivers can be categorized into four types: <ul style="list-style-type: none"> <li>• <b>Type 1:</b> JDBC-ODBC Bridge Driver</li> <li>• <b>Type 2:</b> Native-API Driver</li> <li>• <b>Type 3:</b> Network Protocol Driver</li> <li>• <b>Type 4:</b> Thin Driver (pure Java driver)</li> </ul> </li> </ul> <p>Type 4 drivers are commonly used because they are platform-independent and don't require any native libraries.</p> <ul style="list-style-type: none"> <li>• <b>Connection Interface:</b> Establishes a connection to the database. This interface includes methods to create statements, manage transactions, and close the connection.</li> <li>• <b>Statement Interface:</b> Used to execute SQL statements against the database. There are three types of statements: <ul style="list-style-type: none"> <li>• <b>Statement:</b> Used for executing simple SQL queries without parameters.</li> <li>• <b>PreparedStatement:</b> Used for executing parameterized SQL queries, which helps prevent SQL injection and improves performance.</li> <li>• <b>CallableStatement:</b> Used to execute stored procedures in the database.</li> </ul> </li> <li>• <b>ResultSet Interface:</b> Represents the result set of a query. It provides methods to retrieve data from the result set, navigate through the rows, and update data if applicable.</li> <li>• <b>Cursors</b> play an essential role in database programming by enabling fine-grained control over the processing of result sets, making them useful for complex data manipulation tasks. However, their performance overhead necessitates careful consideration when designing database interactions.</li> </ul> <pre> DECLARE     emp_cursor IS SELECT * FROM employees;     emp_record emp_cursor%ROWTYPE; BEGIN     OPEN emp_cursor;     LOOP         FETCH emp_cursor INTO emp_record;         EXIT WHEN emp_cursor%NOTFOUND;         DBMS_OUTPUT.PUT_LINE(emp_record.name    ' - '    emp_record.salary);     END LOOP;     CLOSE emp_cursor; END; </pre>		
<p>10</p>	<p>Embedded SQL allows the integration of SQL statements within host programming languages (like C, C++, or Java) to facilitate database interactions.</p> <p><b>Key Features</b></p> <ul style="list-style-type: none"> <li>• <b>Static SQL:</b> SQL statements are defined at compile time.</li> <li>• <b>Preprocessing:</b> A preprocessor translates SQL into calls to a database interface.</li> <li>• <b>Type Safety:</b> Compile-time checks ensure data type consistency.</li> </ul> <p><b>Example (C)</b></p> <pre> c Copy code EXEC SQL SELECT name INTO :emp_name FROM employees WHERE id = :emp id; </pre>		

## Advantages

- **Performance:** Faster execution due to precompilation.
- **Security:** Helps prevent SQL injection attacks.
- **Error Checking:** Syntax errors are caught at compile time.

## Disadvantages

- **Less Flexibility:** Hard to create dynamic queries.
- **Tight Coupling:** Links application logic closely to database schema.

## II. SQLJ

### Overview

SQLJ is an extension of SQL for Java, allowing SQL statements to be embedded directly in Java code while maintaining strong typing.

### Key Features

- **Type Safety:** Compile-time checks for SQL statements and variables.
- **Precompilation:** A precompiler converts SQLJ into standard JDBC code.
- **Java Integration:** Combines Java syntax with SQL statements.

### Example (Java)

```
java
Copy code
#sql {
    SELECT name INTO :empName FROM employees WHERE id = 1
};
```

### Advantages

- **Type Safety:** Errors caught at compile time.
- **Ease of Use:** SQL statements are readable and integrated with Java.
- **Performance:** Optimized as SQLJ compiles to JDBC.

### Disadvantages

- **Precompilation Step:** Requires an additional build step.
- **Learning Curve:** Developers need to learn SQLJ syntax and JDBC.

