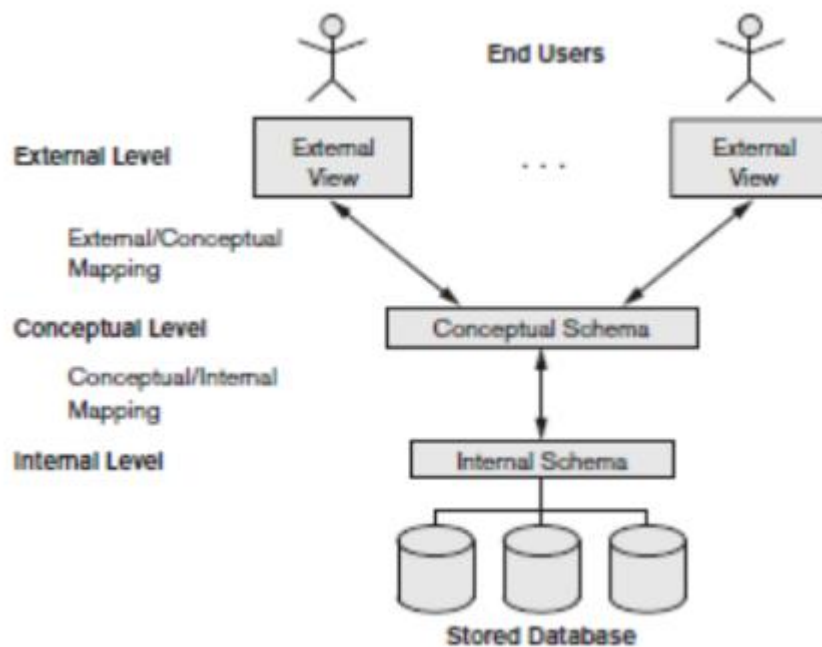


--	--	--	--	--	--	--	--	--	--

Internal Assessment Test 1 – March 2025

Sub:	Database Management System					Sub Code:	BCS403	Branch:	AIML		
Date:	27/3/2025	Duration:	90 min	Max Marks:	50	Sem/Sec:	IV /A, B & C			OBE	
<u>Answer any FIVE FULL Questions</u>									MARKS	CO	RBT
1	<p>Describe Main characteristics of DBMS approach and how it is different from the traditional file system.</p> <p>1. Self-Description: A database system includes—in addition to the data stored that is of relevance to the organization—a complete definition/description of the database's structure and constraints. This meta-data (i.e., data about data) is stored in the so-called system catalog, which contains a description of the structure of each file, the type and storage format of each field, and the various constraints on the data (i.e., conditions that the data must satisfy). The system catalog is used not only by users (e.g., who need to know the names of tables and attributes, and sometimes data type information and other things), but also by the DBMS software, which certainly needs to "know" how the data is structured/organized in order to interpret it in a manner consistent with that structure.</p> <p>2. Insulation between Programs and Data; Data Abstraction: Program-Data Independence: In traditional file processing, the structure of the data files accessed by an application is "hard-coded" in its source code. (E.g., Consider a file descriptor in a COBOL program: it gives a detailed description of the layout of the records in a file by describing, for each field, how many bytes it occupies.)</p> <p>In contrast, DBMS access programs, in most cases, do not require such changes, because the structure of the data is described (in the system catalog) separately from the programs that access it and those programs consult the catalog in order to ascertain the structure of the data</p> <p>3. Data Abstraction: A data model is used to hide storage details and present the users with a conceptual view of the database. Programs refer to the data model constructs rather than data storage details</p> <p>Example by which to illustrate this concept: Suppose that you are given the task of</p>								10	CO1	L1

	<p>developing a program that displays the contents of a particular data file. Specifically, each record should be displayed as follows:</p> <p>Record #i: value of first field value of second field value of last field</p> <p>To keep things very simple, suppose that the file in question has fixed-length records of 57 bytes with six fixed-length fields of lengths 12, 4, 17, 2, 15, and 7 bytes, respectively, all of which are ASCII strings. Developing such a program would not be difficult. However, the obvious solution would be tailored specifically for a file having the particular structure described here and would be of no use for a file with a different structure.</p> <p>4. Multiple Views of Data: Different users (e.g., in different departments of an organization) have different "views" or perspectives on the database. For example, from the point of view of a Bursar's Office employee, student data does not include anything about which courses were taken or which grades were earned. (This is an example of a subset view.)</p> <p>A good DBMS has facilities for defining multiple views. This is not only convenient for users, but also addresses security issues of data access.</p> <p>5. Data Sharing and Multi-user Transaction Processing: As you learned about (or will) in the OS course, the simultaneous access of computer resources by multiple users/processes is a major source of complexity. The same is true for multi-user DBMS's. Arising from this is the need for concurrency control, which is supposed to ensure that several users trying to update the same data do so in a "controlled" manner so that the results of the updates are as though they were done in some sequential order. This gives rise to the concept of a transaction, which is a process that makes one or more accesses to a database and which must have the appearance of executing in isolation from all other transactions (even ones that access the same data at the "same time") and of being atomic (in the sense that, if the system crashes in the middle of its execution, the database contents must be as though it did not execute at all). Applications such as airline reservation systems are known as online transaction processing applications.</p>			
2 a.	<p>Explain the three-schema architecture with neat diagram. Why do we need mapping among the schema levels?</p>	6	CO1	L2



The goal of the three-schema architecture, illustrated in the above Figure , is to separate the user applications and the physical database. In this architecture, schemas can be defined at the following three levels:

1. The internal level has an internal schema, which describes the physical storage structure of the database. The internal schema uses a physical data model and describes the complete details of data storage and access paths for the database.
2. The conceptual level has a conceptual schema, which describes the structure of the whole database for a community of users. The conceptual schema hides the details of physical storage structures and

concentrates on describing entities, data types, relationships, user operations, and constraints. A high-level data model or an implementation data model can be used at this level.

3. The external or view level includes a number of external schemas or user views. Each external schema describes the part of the database that a particular user group is interested in and hides the rest of the database from that user group. A high-level data model or an implementation data model can be used at this level

The three-schema architecture can be used to explain the concept of data independence, which can be defined as the capacity to change the schema at one level of a database system without having to change the schema at the next higher level. We can define two types of data independence:

1. Logical data independence is the capacity to change the conceptual schema without having to change external schemas or application programs. We may change the conceptual schema to expand the database (by adding a record type or data item), or to reduce the database (by

	<p>removing a record type or data item). In the latter case, external schemas that refer only to the remaining data should not be affected. Only the view definition and the mappings need be changed in a DBMS that supports logical data independence. Application programs that reference the external schema constructs must work as before, after the conceptual schema undergoes a logical reorganization. Changes to constraints can be applied also to the conceptual schema without affecting the external schemas or application programs.</p> <p>2. Physical data independence is the capacity to change the internal schema without having to change the conceptual (or external) schemas. Changes to the internal schema may be needed because some physical files had to be reorganized—for example, by creating additional access structures—to improve the performance of retrieval or update. If the same data as before remains in the database, we should not have to change the conceptual schema.</p>			
2 b.	<p>Explain data models and its types with the help of examples.</p> <p>Data Models and Their Types (With Examples) A data model defines how data is organized, stored, and manipulated within a database system. It provides a structured representation of the data and relationships between different elements.</p>	4	CO1	L2
	<p>Types of Data Models:</p> <p>1. Hierarchical Data Model</p> <ul style="list-style-type: none"> • Structure: Data is organized in a tree-like structure with parent-child relationships. • Example: <ul style="list-style-type: none"> ○ Consider an organization structure where the CEO is at the top, followed by managers, then employees. ○ Example Representation: <pre> css CopyEdit CEO ├── Manager 1 │ ├── Employee A │ └── Employee B └── Manager 2 ├── Employee C └── Employee D </pre> <p>2. Network Data Model</p> <ul style="list-style-type: none"> • Structure: Similar to the hierarchical model but allows a many-to-many relationship using pointers. • Example: <ul style="list-style-type: none"> ○ University Database: A student can enroll in multiple courses, and a course can have multiple students. ○ Representation: <pre> css CopyEdit Student A <--> Course X Student A <--> Course Y </pre>			

	<p>Student B <--> Course X</p> <p>3. Relational Data Model</p> <ul style="list-style-type: none">• Structure: Data is stored in tables (relations) with rows (records) and columns (fields).• Example:<ul style="list-style-type: none">◦ Customer Database:<table><thead><tr><th>Customer_ID</th><th>Name</th><th>Email</th></tr></thead><tbody><tr><td>101</td><td>Alice</td><td>alice@email.com</td></tr><tr><td>102</td><td>Bob</td><td>bob@email.com</td></tr></tbody></table>◦ SQL (Structured Query Language) is used to manage relational databases. <p>4. Object-Oriented Data Model</p> <ul style="list-style-type: none">• Structure: Data is stored as objects, similar to object-oriented programming (OOP) concepts.• Example:<ul style="list-style-type: none">◦ Online Shopping System: A "Product" can have multiple attributes like name, price, and category. <pre>arduino CopyEdit class Product { String name; double price; String category; }</pre> <p>5. Entity-Relationship (E-R) Model</p> <ul style="list-style-type: none">• Structure: Uses entities, attributes, and relationships to design the database structure.• Example:<ul style="list-style-type: none">◦ Hospital Management System:<ul style="list-style-type: none">▪ Entities: Patient, Doctor, Appointment▪ Relationships: A patient can have multiple appointments, and each appointment is assigned to a doctor.	Customer_ID	Name	Email	101	Alice	alice@email.com	102	Bob	bob@email.com			
Customer_ID	Name	Email											
101	Alice	alice@email.com											
102	Bob	bob@email.com											
3 a.	<p>Construct an Entity-Relationship (E-R) diagram for a Movie Database considering the following entities, attributes, and relationships. Represent the entities, attributes, relationships, and cardinalities accurately using standard notations.</p> <p>Entities and Attributes:</p> <p>Movie (Movie_ID, Title, Release_Year, Genre, Language, Duration)</p> <p>Production House (Production_ID, Name, Established_Year, Country)</p> <p>Director (Director_ID, Name, Date_of_Birth, Nationality)</p> <p>Actor (Actor_ID, Name, Date_of_Birth, Gender, Nationality)</p> <p>Role (Role_ID, Character_Name, Actor_ID, Movie_ID)</p> <p>Award (Award_ID, Name, Category, Year, Winner_ID)</p>	6	CO2	L3									

3b.	<p>For the Movie database given in question number 3a, Write relational algebra queries for the following.</p> <ol style="list-style-type: none"> 1. Retrieve all movies that belong to the "Action" genre. 2. Retrieve the names of all production houses. 3. Retrieve the title and release year of all movies that were released in 2020. <p>1. Retrieve all movies that belong to the "Action" genre. Relational Algebra Query: $\sigma_{\text{Genre}='Action'}(\text{Movie})$ Explanation:</p> <ul style="list-style-type: none"> • σ (Selection) is used to filter movies where the Genre is "Action". • This will return all movies that belong to the Action genre. 	4	CO2	L3
	<p>2. Retrieve the names of all production houses. Relational Algebra Query: $\pi_{\text{Name}}(\text{Production_House})$ Explanation:</p> <ul style="list-style-type: none"> • π (Projection) is used to select only the Name attribute from the Production_House table. • This will return all unique production house names. 			
	<p>3. Retrieve the title and release year of all movies that were released in 2020. Relational Algebra Query: $\pi_{\text{Title, Release_Year}}(\sigma_{\text{Release_Year}=2020}(\text{Movie}))$ Explanation:</p> <ul style="list-style-type: none"> • $\sigma_{\text{Release_Year}=2020}(\text{Movie})$: Selects movies where Release_Year is 2020. • $\pi_{\text{Title, Release_Year}}$: Extracts only the Title and Release_Year attributes from the filtered movies. 			
4.	<p>Explain the following.</p> <p>i) Database Schema and Database State</p>	10	CO2	L2

<div><div>ii) Participation Constraints</div><div>iii) Recursive Relationships and Role names.</div><div>iv) Cardinality Ratio</div><div>v) Primary Key</div><div>vi) Candidate Key</div><div>vii) Foreign Key</div><div><div>i) Database Schema and Database State</div><div><div>Database Schema:</div><div><div>The <i>logical structure/design</i> of the database.</div><div>Defines tables, attributes, data types, constraints, relationships, etc.</div><div>Example:</div></div></div><div>sql</div><div>Copy</div><div>CREATE TABLE Employee (EmpID INT PRIMARY KEY, Name VARCHAR(50), Dept VARCHAR(20));</div><div><div>Schema remains <i>constant</i> unless altered.</div><div>Database State (Instance):</div><div><div>The <i>actual data</i> stored in the database at a given time.</div><div>Changes with insertions, updates, and deletions.</div><div>Example:</div><div><table><tr><th>EmpID</th><th>Name</th><th>Dept</th></tr><tr><td>101</td><td>Alice</td><td>HR</td></tr><tr><td>102</td><td>Bob</td><td>IT</td></tr></table></div></div></div></div></div>	EmpID	Name	Dept	101	Alice	HR	102	Bob	IT			
EmpID	Name	Dept										
101	Alice	HR										
102	Bob	IT										
<div><div>ii) Participation Constraints</div><div><div>Specifies whether an entity's participation in a relationship is mandatory (total) or optional (partial).</div><div>Total Participation (Double Line in ER Diagram):</div><div><div>Every entity must participate in the relationship.</div><div>Example: Every Student <i>must</i> enroll in at least one Course.</div></div><div>Partial Participation (Single Line):</div><div><div>Some entities may not participate.</div><div>Example: Not every Employee manages a project.</div></div></div></div>												
<div><div>iii) Recursive Relationships and Role Names</div><div><div>Recursive Relationship:</div><div><div>An entity relates <i>to itself</i> in a relationship.</div><div>Example: An Employee supervises other Employees.</div></div></div><div>mermaid</div><div>Copy</div><div>erDiagram</div><div>EMPLOYEE --o{ EMPLOYEE : "supervises"</div><div><div>Role Name:</div><div><div>Clarifies the purpose of an entity in a relationship.</div><div>Example: In "supervises", roles are Supervisor (one side) and Subordinate (many side).</div></div></div></div>												
<div><div>iv) Cardinality Ratio</div><div><div>Defines the <i>numeric relationship</i> between entities in a relationship:</div><div>1:1 (One-to-One):</div><div><div>Example: One Employee has one CompanyCar.</div></div></div></div>												

	<ul style="list-style-type: none">○ 1:N (One-to-Many):<ul style="list-style-type: none">▪ Example: One Department has many Employees.○ M:N (Many-to-Many):<ul style="list-style-type: none">▪ Example: Students enroll in many Courses, and Courses have many Students.																								
	<p>v) Primary Key</p> <ul style="list-style-type: none">• A <i>minimal set of attributes</i> that uniquely identifies a row in a table.• Properties:<ul style="list-style-type: none">○ Unique○ Not NULL○ Immutable (should not change)• Example: <pre>sql Copy CREATE TABLE Student (StudentID INT PRIMARY KEY, -- Primary Key Name VARCHAR(50));</pre>																								
	<p>vi) Candidate Key</p> <ul style="list-style-type: none">• A <i>super key</i> (set of attributes that uniquely identifies a row) <i>without unnecessary attributes</i>.• A table can have multiple candidate keys, but only one becomes the primary key.• Example:<ul style="list-style-type: none">○ In Employee(EmpID, Email, SSN), both EmpID and SSN are candidate keys.																								
	<p>vii) Foreign Key</p> <ul style="list-style-type: none">• An attribute that <i>references the primary key</i> of another table to enforce referential integrity.• Example: <pre>sql Copy CREATE TABLE Orders (OrderID INT PRIMARY KEY, ProductID INT, CustomerID INT, FOREIGN KEY (ProductID) REFERENCES Products(ProductID), FOREIGN KEY (CustomerID) REFERENCES Customers(CustomerID));</pre> <ul style="list-style-type: none">• Ensures that ProductID in Orders must exist in Products.																								
	<table><tr><th>Summary Table Concept</th><th>Definition</th><th>Example</th></tr><tr><td>Database Schema</td><td>Blueprint/structure of the database.</td><td>CREATE TABLE P</td></tr><tr><td>Database State</td><td>Actual data stored at a given time.</td><td>Rows in the Employ</td></tr><tr><td>Participation</td><td>Whether an entity must participate in a relationship (total/partial).</td><td>All students must en</td></tr><tr><td>Recursive Relationship</td><td>An entity relates to itself.</td><td>Employee supervise</td></tr><tr><td>Cardinality</td><td>Numerical relationship between entities (1:1, 1:N, M:N).</td><td>One department has employees.</td></tr><tr><td>Primary Key</td><td>Uniquely identifies a row (unique, NOT NULL).</td><td>StudentID in Studer</td></tr></table>	Summary Table Concept	Definition	Example	Database Schema	Blueprint/structure of the database.	CREATE TABLE P	Database State	Actual data stored at a given time.	Rows in the Employ	Participation	Whether an entity must participate in a relationship (total/partial).	All students must en	Recursive Relationship	An entity relates to itself.	Employee supervise	Cardinality	Numerical relationship between entities (1:1, 1:N, M:N).	One department has employees.	Primary Key	Uniquely identifies a row (unique, NOT NULL).	StudentID in Studer			
Summary Table Concept	Definition	Example																							
Database Schema	Blueprint/structure of the database.	CREATE TABLE P																							
Database State	Actual data stored at a given time.	Rows in the Employ																							
Participation	Whether an entity must participate in a relationship (total/partial).	All students must en																							
Recursive Relationship	An entity relates to itself.	Employee supervise																							
Cardinality	Numerical relationship between entities (1:1, 1:N, M:N).	One department has employees.																							
Primary Key	Uniquely identifies a row (unique, NOT NULL).	StudentID in Studer																							

	<p>Candidate Key A possible primary key (minimal super key).</p> <p>Foreign Key References a primary key in another table.</p>	SSN or Email in Employee.		
5	<p>Explain the select, project, union, intersection, set difference, Cartesian product and join operations in relational algebra with suitable example.</p> <p>Selection Operator(σ)</p> <p>Selection and Projection are unary operators.</p> <p>The selection operator is sigma: σ</p> <p>The selection operation acts like a filter on a relation by returning only a certain number of tuples.</p> <p>$\sigma_C(R)$ Returns only those tuples in R that satisfy condition C</p> <p>A condition C can be made up of any combination of comparison or logical operators that operate on the attributes of R. Comparison operators: $>, <, \leq, \geq, \neq, =$</p> <p>Logical operators</p> <p>\neg - not, \vee - or</p> <p>Example</p> <p>Select only those Employees in the CS department:</p> <p>$\sigma_{\text{Dept} = \text{'CS'}}(\text{EMP})$</p> <p>Projection($\pi$)</p> <p>Projection is also a Unary operator.</p> <p>The Projection operator is π</p> <p>Projection limits the attributes that will be returned from the original relation.</p> <p>The general syntax is: $\pi_{\text{attributes}} R$</p> <p>Where attributes is the list of attributes to be displayed and R is the relation.</p> <p>The resulting relation will have the same number of tuples as the original relation (unless there are duplicate tuples produced).</p> <p>The degree of the resulting relation may be equal to or less than that of the original relation</p> <p>Project only the names and departments of the employees:</p> <p>$\pi_{\text{name, dept}}(\text{EMP})$</p> <p>THETA JOIN: Similar to a CARTESIAN PRODUCT followed by a SELECT. The condition c is called a join condition.</p> <p>$R(A_1, A_2, \dots, A_m, B_1, B_2, \dots, B_n) \bowtie_{c} R_2(B_1, B_2, \dots, B_n)$</p> <p>EQUIJOIN: The join condition c includes one or more equality comparisons involving attributes from R1 and R2. That is, c is of the form:</p> <p>$(A_i = B_j) \text{ AND } \dots \text{ AND } (A_h = B_k); 1 \leq i, h \leq m, 1 \leq j, k \leq n$</p> <p>In the above EQUIJOIN operation:</p> <p>A_i, \dots, A_h are called the join attributes of R1</p> <p>B_j, \dots, B_k are called the join attributes of R2</p> <p>Example of using EQUIJOIN:</p> <p>Retrieve each DEPARTMENT's name and its manager's name:</p> <p>T DEPARTMENT MGRSSN = SSN EMPLOYEE</p> <p>RESULT</p> <p>DNAME, FNAME, LNAME</p> <p>(T)</p> <p>NATURAL JOIN (*):</p> <p>In an EQUIJOIN $R \bowtie R_2$, the join attribute of R2 appear redundantly in the</p>	10	CO2	L3

	<p>result relation R. In a NATURAL JOIN, the redundant join attributes of R2 are eliminated from R. The equality condition is implied and need not be specified. $R \bowtie R$ *(join attributes of R1),(join attributes of R2) R Example: Retrieve each EMPLOYEE's name and the name of the DEPARTMENT he/she works for: T EMPLOYEE *(DNO),(DNUMBER) DEPARTMENT RESULT FNAME,LNAME,DNAME (T) If the join attributes have the same names in both relations, they need not be specified and we can write $R \bowtie R$ * R2. Example: Retrieve each EMPLOYEE's name and the name of his/her SUPERVISOR:</p>			
6 a.	<p>Demonstrate the usage of the following SQL commands by writing and executing appropriate queries on a sample database: i) INSERT, ii) DELETE, iii) UPDATE, iv) ALTER, v) SELECT</p> <p>insert INSERT INTO target [(field1[, field2[, ...]])] VALUES (value1[, value2[, ...]);</p> <p>So, to add a User record for user Jim Jones, we would issue the following INSERT query: INSERT INTO User (FirstName, LastName, UserID, Dept, EmpNo, PCType) 6 VALUES ("Jim", "Jones", "Jjones","Finance", 9, "DellDimR450");</p> <p>Obviously populating a database by issuing such a series of SQL commands is both tedious and prone to error, which is another reason why database applications have front-ends. Even without a specifically designed front-end, many database systems - including MS Access - allow data entry direct into tables via a spreadsheet-like interface.</p> <p>ii) Delete Now that we know how to add new records and to update existing records it only remains to learn how to delete records before we move on to look at how we search through and collate data. As you would expect SQL provides a simple command to delete complete records. The syntax of the command is: DELETE [table.*] FROM table WHERE criteria;</p> <p>iii) drop If you have already executed the original CREATE TABLE command your database will already contain a table called User, so let's get rid of that using the DROP command:</p>	6	CO3	L3

	<p>DROP TABLE User;</p> <p>iv) alter</p> <p>Once a table is created it's structure is not necessarily fixed in stone. In time requirements change and the structure of the database is likely to evolve to match your wishes. SQL can be used to change the structure of a table, so, for example, if we need to add a new field to our User table to tell us if the user has Internet access, then we can execute an SQL ALTER TABLE command as shown below:</p> <p>ALTER TABLE User ADD COLUMN Internet BOOLEAN;</p> <p>To delete a column the ADD keyword is replaced with DROP, so to delete the field we have just added the SQL is:</p> <p>ALTER TABLE User DROP COLUMN Internet;</p> <p>v) update</p> <p>the UPDATE command, with syntax:</p> <p>UPDATE table</p> <p>SET newvalue</p> <p>WHERE criteria;</p> <p>For example, let's assume that we want to move user Jim Jones from the Finance department to Marketing. Our SQL statement would then be:</p> <p>UPDATE User</p> <p>SET Dept="Marketing"</p> <p>WHERE EmpNo=9;</p>			
6b.	<p>For the Movie database given in question number 3a, Write SQL query for the following.</p> <ol style="list-style-type: none"> 1. Retrieve the Movie Title and Release Year of all movies released after 2018. 2. Retrieve the names of all directors from the database. 3. Retrieve all details of movies that belong to the "Comedy" genre. <p>1. Retrieve the Movie Title and Release Year of all movies released after 2018</p> <p>sql</p> <p>Copy</p> <p>SELECT title AS "Movie Title", release_year AS "Release Year"</p> <p>FROM movies</p> <p>WHERE release_year > 2018;</p> <p>2. Retrieve the names of all directors from the database</p> <p>sql</p> <p>Copy</p> <p>SELECT DISTINCT director_name AS "Director Name"</p> <p>FROM directors;</p> <p>3. Retrieve all details of movies that belong to the "Comedy" genre</p> <p>SELECT *</p> <p>FROM movies</p>	4	CO3	L3

	WHERE genre = 'Comedy';			
--	-------------------------	--	--	--

Faculty Signature

CCI Signature

HOD Signature