



UNIVERSITY OF COLOMBO, SRI LANKA

UNIVERSITY OF COLOMBO SCHOOL OF COMPUTING

DEGREE OF BACHELOR OF INFORMATION TECHNOLOGY (EXTERNAL)

Academic Year 2025 – 1st Year Examination – Semester 1

IT2306 – Database Systems
Multiple Choice Question Paper

(TWO HOURS)

Important Instructions :

- The duration of the paper is **2 (two) hours**.
- The medium of instruction and questions is English.
- The paper has **40 questions** and **12 pages**.
- All questions are of the **MCQ** (Multiple Choice Questions) type.
- All questions should be answered.
- Each question will have 5 (five) choices with **one / one or more** correct answers.
- All questions will carry **equal** marks.
- There will be a penalty for incorrect responses to discourage guessing.
- The mark given for a question will vary from -1 (*All the incorrect choices are marked & no correct choices are marked*) to +1 (*All the correct choices are marked & no incorrect choices are marked*). However, the minimum mark per question would be zero.
- Answers should be marked on the special answer sheet provided.
- Note that questions appear on both sides of the paper.
If a page is not printed, please inform the supervisor immediately.
- Mark the correct choices on the question paper first and then transfer them to the given answer sheet, which will be machine marked. **Please completely read and follow the instructions given on the other side of the answer sheet before you shade your correct choices.**
- All kinds of electronic devices, including calculators are **not** allowed.
- *All Rights Reserved.*

1) Which of the following statements are common limitations of file-based systems?

- (a) Data redundancy across multiple applications.
- (b) Support for multiple user views.
- (c) Use of centralized data dictionary.
- (d) Incompatible file formats between applications.
- (e) Program-data dependency.

2) Which of the following statements best describes logical data independence?

- (a) The ability to modify the internal schema without affecting the conceptual schema.
- (b) The ability to modify the conceptual schema without affecting external schemas.
- (c) The ability to store data in multiple physical locations.
- (d) The separation of concurrent transactions to ensure consistency.
- (e) The ability to change file structures without affecting storage devices.

3) Which of the following statements best describes metadata in a database context?

- (a) The actual records stored within database tables.
- (b) Data that describes the structure, properties, and characteristics of other data.
- (c) Temporary data generated during database transactions.
- (d) The relationships among users who access the database system.
- (e) The collection of indexes used to improve database query performance.

4) Which of the following statements are correct about keys in a relational database?

- (a) A primary key uniquely identifies each tuple in a relation.
- (b) A foreign key is used to enforce relationships between relations.
- (c) A composite key consists of two or more attributes.
- (d) A candidate key may contain duplicate values.
- (e) Every superkey is a minimal superkey.

5) Which of the following statements are correct with respect to the Relational Data Model?

- (a) A relation is a collection of tuples represented in tabular form.
- (b) A schema changes whenever a new tuple is inserted.
- (c) Cardinality of a relation changes when tuples are added or removed.
- (d) The degree of a relation depends on the current number of tuples stored in the relation.
- (e) Constraints are derived from business rules.

6) Consider the following relations where Dno is a foreign key referencing Dnumber.

DEPARTMENT (Dnumber, Dname)
EMPLOYEE (EmpID, Name, Dno)

Which of the following operations would violate referential integrity?

- (a) Inserting an employee with Dno = 9 when no department with Dnumber = 9 exists.
- (b) Inserting a department with Dnumber = 10.
- (c) Updating an employee's Dno from 5 to NULL.
- (d) Deleting a department whose Dnumber is referenced by existing employees.
- (e) Updating a department name.

- 7) A hospital management system maintains patient records. For each patient, the system stores an *Address* consisting of *House Number*, *Street Name*, and *City*. Each component may be queried independently by system users.

Which of the following should be used to model the *Address* in the ER diagram?

- (a) Derived attribute
- (b) Composite attribute
- (c) Multivalued attribute
- (d) Weak entity
- (e) Atomic attribute

- 8) A software company maintains information about employees and projects. An employee may work on multiple projects, and a project may involve multiple employees. The company needs to record the number of hours each employee works on each project.


Where should the attribute *HoursWorked* be placed in the ER model?

- (a) As an attribute of the EMPLOYEE entity
- (b) As an attribute of the PROJECT entity
- (c) As an attribute of the DEPARTMENT entity
- (d) As a derived attribute of the EMPLOYEE entity
- (e) As an attribute of the WORKS_ON relationship

- 9) A company maintains employee records and classifies every employee as either *Permanent*, *Contract*, or *Intern*. Company policy states that an employee must belong to one and only one of these categories. Which specialization constraint best models this requirement?


- (a) Overlapping, Partial
- (b) Disjoint, Partial
- (c) Overlapping, Total
- (d) Disjoint, Total
- (e) Overlapping, Mandatory

Consider the following Entity-Relationship Diagram (ERD) and answer Questions 10 through 12.



- 10) Which of the following statements are correct regarding the above ERD?
- (a) A book can be associated with multiple publishers.
 - (b) A publisher may exist without publishing any books.
 - (c) Every book should be associated with a publisher.
 - (d) The relationship between Book and Publisher is binary and One-to-Many.
 - (e) Book has partial participation in the relationship.
- 11) Which of the following relations are possible if the above ERD is mapped into a relational model?
- (a) Book (ISBN, Title, PublicationYear, Price)
 - (b) Publishes (PublisherID, ISBN)
 - (c) Publisher (PublisherID, PublisherName, Country)
 - (d) Book (ISBN, Title, PublicationYear, Price, PublisherID)
 - (e) Publisher (PublisherID, ISBN, PublisherName, Country)
- 12) Which of the following changes would convert the relationship in the above ERD into a Many-to- Many relationship?
- (a) A publisher can publish many books.
 - (b) A book can be published by multiple publishers.
 - (c) Every publisher must publish at least one book.
 - (d) Every book must have a publisher.
 - (e) ISBN is no longer unique.

Consider the following ER Diagram (ERD) and answer Questions 13 and 14.



- 13) Which of the following statements are true based on the above ERD?
- (a) Every Person must be a Student or an Employee or both.
 - (b) A Student may also be an Employee.
 - (c) Student and Employee are disjoint subtypes.
 - (d) A Person can exist without belonging to either Student or Employee.
 - (e) Total participation exists between Person and the specialization.

- 14) Which of the following relations could result from mapping the above ERD to a relational schema?

- (a) Person (PersonID, Name)
- (b) Student (PersonID, DegreeProgram)
- (c) Person (PersonID, Name, DegreeProgram, Salary)
- (d) Employee (PersonID, Salary)
- (e) Employee (Name, Salary)

Consider the following relational schema and answer Questions 15 and 16.

Student (stu_id, name, age, dept_id)
Department (dept_id, dept_name)

The following sequence of relational algebra operations is given to you.

Result1 $\leftarrow \sigma_{\text{dept_name} = \text{'Computer Science'}}(\text{Department})$
Result2 $\leftarrow \text{Student} \bowtie_{\text{dept_id} = \text{dept_id}}(\text{Result1})$
Result3 $\leftarrow \sigma_{\text{age} > 20}(\text{Result2})$
Result4 $\leftarrow \pi_{\text{name}}(\text{Result3})$

- 15) What is the result of the above sequence of relational algebra operations?

- (a) Names of all students in the Computer Science department.
- (b) Names of all students older than 20 years.
- (c) Names of students older than 20 years who belong to the Computer Science department.
- (d) Names and ages of students older than 20 years who belong to the Computer Science department.
- (e) Names of Computer Science departments having students older than 20 years.

- 16) The following intermediate relations have already been computed.

Result1 $\leftarrow \pi_{\text{dept_id}}(\text{Department})$
Result2 $\leftarrow \pi_{\text{dept_id}}(\text{Student})$

Which of the following relational algebra sequences retrieves the department IDs of departments that have no students?

- (a) Result3 $\leftarrow \text{Result1} - \text{Result2}$
- (b) Result3 $\leftarrow \text{Result2} - \text{Result1}$
- (c) Result3 $\leftarrow \text{Result1} \cap \text{Result2}$
- (d) Result3 $\leftarrow \text{Result1} \cup \text{Result2}$
- (e) Result3 $\leftarrow \text{Result1} \times \text{Result2}$

- 17) Consider the following SQL query.

```
SELECT E.EmpName, D.DeptName
FROM Employee E
JOIN Department D
ON E.DeptID = D.DeptID;
```

Which of the following relational algebra expressions are equivalent to the given SQL query?

- (a) $\pi_{EmpName, DeptName} (Employee \bowtie_{Employee.EmpID = Department.DeptID} Department)$
- (b) $\pi_{EmpName, DeptName} (Employee \bowtie_{Employee.DeptID = Department.DeptID} Department)$
- (c) $\pi_{EmpName, DeptName} (\sigma_{Employee.DeptID = Department.DeptID} (Employee \times Department))$
- (d) $\pi_{EmpName, DeptName} (Employee \times Department)$
- (e) $\pi_{EmpName, DeptName} (Employee \bowtie_{E.DeptID = D.DeptID} Department)$

18) Consider the following relational schema designed for a library.

Member (m_id, mem_name, joined_date, mem_type, age, email, address)
 Membership_type (type_id, type_name, max_books)
 Book (b_id, title, author, isbn, category)
 Borrows (mem_id, book_id, borrowed_date, returned_date, renew_count)

Which of the following relational algebra sequences retrieves the names of members who have borrowed at least one book in the *Fiction* category?

- (a) Result1 $\leftarrow \pi_{b_id} (\sigma_{category='Fiction'} (Book))$
 Result2 $\leftarrow Result1 \bowtie_{b_id=book_id} Borrows$
 Result3 $\leftarrow \pi_{mem_id} (Result2)$
 Result4 $\leftarrow Member \bowtie_{m_id=mem_id} Result3$
 Result5 $\leftarrow \pi_{mem_name} (\sigma_{mem_type='Fiction'} (Result4))$
- (b) Result1 $\leftarrow \sigma_{category='Fiction'} (Book)$
 Result2 $\leftarrow \pi_{b_id} (Result1)$
 Result3 $\leftarrow Result2 \bowtie_{b_id=book_id} Borrows$
 Result4 $\leftarrow Result3 \bowtie_{mem_id=m_id} Member$
 Result5 $\leftarrow \pi_{mem_name} (Result4)$
- (c) Result1 $\leftarrow \sigma_{category='Fiction'} (Book)$
 Result2 $\leftarrow \pi_{mem_id} (Borrows)$
 Result3 $\leftarrow Result1 \times Result2$
 Result4 $\leftarrow \pi_{mem_name} (Member \bowtie_{mem_id=m_id} Result3)$
- (d) Result1 $\leftarrow \pi_{b_id} (\sigma_{category='Fiction'} (Book))$
 Result2 $\leftarrow Result1 \bowtie_{b_id=book_id} Borrows$
 Result3 $\leftarrow \pi_{mem_id} (Result2)$
 Result4 $\leftarrow \pi_{mem_id} (Member)$
 Result5 $\leftarrow \pi_{mem_name} (Member \bowtie_{m_id=mem_id} (Result3 \cap Result4))$
- (e) Result1 $\leftarrow \sigma_{category='Fiction'} (Book)$
 Result2 $\leftarrow Result1 \bowtie_{b_id=book_id} Borrows$
 Result3 $\leftarrow Result2 \bowtie_{mem_id=m_id} Member$
 Result4 $\leftarrow \pi_{title, mem_name} (Result3)$

19) Which of the following statements are TRUE regarding SQL data types?

- (a) VARCHAR(n) stores variable-length strings with a maximum length of n.
- (b) CHAR(n) stores fixed-length strings and pads shorter values with blanks.
- (c) CLOB is suitable for storing images and other binary data.
- (d) BLOB is suitable for storing large text documents.
- (e) SQL Boolean data type has only two values: TRUE and FALSE.

20) Consider the following table.

```
Employee ( EmpID INT PRIMARY KEY, Name VARCHAR(50) )
```

Which SQL command correctly adds a column named Salary of type DECIMAL(10,2) to the Employee table?

- (a) UPDATE TABLE Employee ADD Salary DECIMAL(10,2);
- (b) MODIFY Employee ADD Salary DECIMAL(10,2);
- (c) MODIFY TABLE Employee ADD Salary DECIMAL(10,2);
- (d) CREATE COLUMN Salary DECIMAL(10,2) IN Employee;
- (e) ALTER TABLE Employee ADD Salary DECIMAL(10,2);

Consider the following company database schema to answer Questions 21 and 22.

```
Department (DeptID, DeptName)
```

```
Employee (EmpID, EmpName, DeptID)
```

The following referential integrity constraint is defined on the *Employee* table:

```
FOREIGN KEY (DeptID)
REFERENCES Department (DeptID)
ON DELETE CASCADE
```

Suppose the following SQL statement is executed:

```
DELETE FROM Department WHERE DeptID = 'D01';
```

21) Which of the following will occur?

- (a) The deletion is rejected because DeptID is referenced by the Employee table.
- (b) The department record is deleted, and the corresponding employee DeptID values are set to NULL.
- (c) The department record is deleted, and all employees belonging to that department D01 are automatically deleted.
- (d) Only the employees belonging to department D01 are deleted, while the department record remains.
- (e) The department record is deleted, and all affected employees are moved to a default department.

22) Consider the following SQL query executed on the Employee table.

```
SELECT DeptID, COUNT(*)
FROM Employee
WHERE Salary > 50000
GROUP BY DeptID
HAVING COUNT(*) > 5
ORDER BY DeptID;
```

Which of the following represents the logical query processing order of the SQL clauses in the above query?

- (a) SELECT → FROM → WHERE → GROUP BY → HAVING → ORDER BY
- (b) FROM → WHERE → GROUP BY → HAVING → SELECT → ORDER BY
- (c) FROM → GROUP BY → WHERE → HAVING → SELECT → ORDER BY
- (d) WHERE → FROM → GROUP BY → HAVING → SELECT → ORDER BY
- (e) FROM → WHERE → SELECT → GROUP BY → HAVING → ORDER BY

Consider the following relations from an online shopping database and answer Questions 23 through 26.

Customer

cust_id	cust_name	city
101	Nimna	Colombo
102	Kamza	Kandy
103	Safan	Jaffna
104	Sithum	Matara

Product

prod_id	prod_name	price
P01	Laptop	150000
P02	Printer	40000
P03	Mouse	2500
P04	Keyboard	5000

Orders

order_id	order_date	cust_id	prod_id	quantity
O1001	2025-01-10	101	P01	1
O1002	2025-01-15	102	P03	2
O1003	2025-01-20	103	P02	1

- 23) Which of the following SQL statements will execute without any error?

(a) INSERT INTO Orders VALUES ('O1004', '2025-02-01', 101, 'P04', 2);
 (b) INSERT INTO Orders VALUES ('O1005', '2025-02-01', 110, 'P04', 2);
 (c) INSERT INTO Orders VALUES ('O1006', '2025-02-01', 101, 'P10', 2);
 (d) INSERT INTO Orders VALUES ('O1002', '2025-02-01', 101, 'P04', 2);
 (e) INSERT INTO Orders VALUES (NULL, '2025-02-01', 101, 'P04', 2);

- 24) Consider the following SQL query:

```
SELECT cust_name, order_id
FROM Customer C
INNER JOIN Orders O
ON C.cust_id = O.cust_id;
```

What is the number of rows returned by the above query?

(a) 2 (b) 3 (c) 4
 (d) 6 (e) 12

- 25) Which SQL query correctly displays each order ID together with its total order value, renamed as totalValue?

(a) SELECT order_id, (price + quantity) AS totalValue
 FROM Orders O JOIN Product P
 ON O.prod_id = P.prod_id;
 (b) SELECT order_id, (price * quantity) AS totalValue
 FROM Orders O JOIN Product P
 ON O.prod_id = P.prod_id;
 (c) SELECT order_id, SUM(price) AS totalValue
 FROM Orders O JOIN Product P
 ON O.prod_id = P.prod_id;
 (d) SELECT order_id, COUNT(price * quantity) AS totalValue
 FROM Orders O JOIN Product P
 ON O.prod_id = P.prod_id;
 (e) SELECT order_id, SUM (price * quantity) AS totalValue
 FROM Orders O JOIN Product P
 ON O.prod_id = P.prod_id;

26) Which SQL statement returns the names of customers who have not placed any orders?

- (a)

```
SELECT cust_name
FROM Customer C INNER JOIN Orders O
ON C.cust_id= O.cust_id;
```
- (b)

```
SELECT cust_name
FROM Customer C
LEFT JOIN Orders O
ON C.cust_id= O.cust_id
WHERE O.order_id IS NULL;
```
- (c)

```
SELECT cust_name
FROM Customer
WHERE cust_id IN (SELECT cust_id FROM Orders);
```
- (d)

```
SELECT cust_name
FROM Customer
CROSS JOIN Orders;
```
- (e)

```
SELECT DISTINCT cust_name
FROM Orders;
```

27) Which SQL statement correctly creates a unique index named idx_emp_dept on the combination of EmpName and DeptID in the Employee table?

- (a)

```
CREATE UNIQUE INDEX idx_emp_dept ON Employee(EmpName, DeptID);
```
- (b)

```
CREATE INDEX UNIQUE idx_emp_dept ON Employee(EmpName, DeptID);
```
- (c)

```
CREATE UNIQUE INDEX idx_emp_dept ON Employee.EmpName,
Employee.DeptID;
```
- (d)

```
CREATE UNIQUE idx_emp_dept INDEX ON Employee(EmpName, DeptID);
```
- (e)

```
CREATE INDEX idx_emp_dept ON Employee EmpName, DeptID);
```

28) Which of the following SQL statements can be used to remove all rows from the Employee table while keeping the table structure intact?

- (a)

```
DROP TABLE Employee;
```
- (b)

```
TRUNCATE TABLE Employee;
```
- (c)

```
DELETE TABLE Employee;
```
- (d)

```
DELETE FROM Employee;
```
- (e)

```
DELETE *FROM Employee;
```

29) Consider the relation: Employee (EmpID, EmpName, Salary)

Which of the following queries returns employees earning the highest salary?

- (a)

```
SELECT * FROM Employee WHERE Salary IN (SELECT MAX(Salary) FROM
Employee GROUP BY EmpID);
```
- (b)

```
SELECT * FROM Employee WHERE Salary = (SELECT MAX(Salary) FROM
Employee);
```
- (c)

```
SELECT * FROM Employee ORDER BY Salary DESC LIMIT 1;
```
- (d)

```
SELECT * FROM Employee E WHERE NOT EXISTS (SELECT * FROM Employee
X WHERE X.Salary > E.Salary);
```
- (e)

```
SELECT * FROM Employee WHERE Salary = (SELECT Salary FROM Employee
ORDER BY Salary DESC);
```

- 30) Which of the following statements about views are FALSE?
- (a) Views can be queried in a similar manner to base tables.
 - (b) A view can be created based on another view.
 - (c) Views always store a physical copy of the data contained in the underlying tables.
 - (d) Every view is automatically updatable.
 - (e) Views can be used to hide sensitive columns from users.
- 31) Which of the following best describes an inline view in SQL?
- (a) A permanent view stored in the database catalog.
 - (b) A view that can only be used inside a trigger.
 - (c) A view that automatically updates the underlying tables.
 - (d) A temporary table created using the CREATE TABLE statement.
 - (e) A subquery that appears in the FROM clause and behaves like a temporary table during query execution.
- 32) User A grants the privilege to read data from the Employee table to User B and allows User B to grant the same privilege to other users. User B subsequently grants the privilege to User C. Later, User A revokes User B's privilege and specifies that any privileges that depend on User B's grant should also be revoked.
- Which of the following explains the outcome?
- (a) Only User B loses the SELECT privilege on the Employee table.
 - (b) Only User C loses the SELECT privilege on the Employee table.
 - (c) User C retains the SELECT privilege because it was granted directly by User B.
 - (d) Both User B and User C lose the SELECT privilege on the Employee table.
 - (e) Neither User B nor User C loses the SELECT privilege on the Employee table.
- 33) Which of the following statements about polyinstantiation in a multilevel secure database are true?
- (a) Polyinstantiation means that multiple users share the same database account.
 - (b) Polyinstantiation helps prevent inference of sensitive information in multilevel secure databases.
 - (c) Different tuples with the same apparent key may exist at different security classification levels.
 - (d) Polyinstantiation requires every relation to have multiple primary keys.
 - (e) Polyinstantiation allows users with different security clearances to access all versions of a tuple simultaneously.
- 34) Consider the SQL command: DROP VIEW Emp_Payroll;
Which of the following statements are correct regarding the effect of executing the above command?
- (a) The view Emp_Payroll and the underlying data are permanently deleted.
 - (b) The definition of the view Emp_Payroll is removed from the database.
 - (c) All records in the underlying base table(s) are permanently deleted.
 - (d) Queries that reference Emp_Payroll will fail unless the view is recreated.
 - (e) Data stored in the underlying base table(s) remains unaffected.

- 35) Consider the following view definition.
- ```
CREATE VIEW DeptSummary AS
SELECT DeptID, COUNT(*) AS NumEmployees
FROM Employee
GROUP BY DeptID;
```

Which of the following reasons explain why the view DeptSummary is generally not updatable?

- (a) The GROUP BY clause combines multiple base table rows into a single result row.
- (b) The view contains an aggregate function (COUNT (\*)).
- (c) The view contains more than one attribute.
- (d) The use of an attribute alias (NumEmployees) prevents updates.
- (e) The view does not contain a WHERE clause.

- 36) Consider the relation: Employee (EmpID, EmpName, Salary, DeptID)  
The Sales Manager wants a view that contains only employees working in the department 'D01'.

Which of the following SQL statements correctly create the view?

- (a) 

```
CREATE VIEW Sales_Emp AS
SELECT *
FROM Employee
WHERE DeptID = 'D01';
```
- (b) 

```
CREATE VIEW Sales_Emp AS
SELECT *
FROM Employee
GROUP BY DeptID='D01';
```
- (c) 

```
CREATE VIEW Sales_Emp
SELECT EmpID, EmpName, Salary, DeptID
FROM Employee
WHERE DeptID = 'D01';
```
- (d) 

```
CREATE VIEW Sales_Emp AS
SELECT *
WHERE DeptID = 'D01'
FROM Employee;
```
- (e) 

```
CREATE VIEW Sales_Emp AS
SELECT EmpID, EmpName, Salary, DeptID
FROM Employee
WHERE DeptID = 'D01';
```

- 37) Which of the following statements correctly distinguishes Second Normal Form (2NF) from Third Normal Form (3NF)?

- (a) 2NF removes transitive dependencies, whereas 3NF removes partial dependencies.
- (b) 2NF removes partial dependencies, whereas 3NF removes transitive dependencies.
- (c) 2NF and 3NF both remove only partial dependencies.
- (d) 2NF applies only to relations with a single-attribute primary key.
- (e) 3NF allows transitive dependencies as long as the relation is in 2NF.

Consider the sports platform at *Nexus-Global* Company and answer Questions 38 through 40.  
The company currently stores tournament information in the following unnormalized (UNF) relation:

**Tournament\_Master**

(MatchID, PlayerID, MatchDate, PlayerName, TeamName, TeamLogo, GameTitle, PrizePool, SponsorName, SponsorCategory)

The business rules related to the Tournament\_Master relation are listed below.

- Each MatchID is associated with exactly one MatchDate, GameTitle, and PrizePool.
- Each PlayerID corresponds to one PlayerName, TeamName, and TeamLogo.
- Each TeamName is associated with exactly one SponsorName.
- Each SponsorName belongs to one SponsorCategory.

38) Which of the following functional dependencies (FDs) are implied by the business rules?

- (a) MatchID  $\rightarrow$  {MatchDate, GameTitle, PrizePool}
- (b) PlayerID  $\rightarrow$  {PlayerName, TeamName, TeamLogo}
- (c) TeamName  $\rightarrow$  SponsorName
- (d) SponsorCategory  $\rightarrow$  SponsorName
- (e) SponsorName  $\rightarrow$  SponsorCategory

39) Which of the following attributes form the primary key of the Tournament\_Master relation?

- (a) MatchID
- (b) PlayerID
- (c) TeamName
- (d) {MatchID, PlayerID}
- (e) {MatchID, PlayerID, TeamName}

40) Which normal form is achieved by the resulting relations after decomposing the Tournament\_Master relation to remove all partial and transitive dependencies?

- (a) The resulting relations do not satisfy First Normal Form (1NF).
- (b) The resulting relations satisfy First Normal Form (1NF) only.
- (c) The resulting relations satisfy Second Normal Form (2NF) but not Third Normal Form (3NF).
- (d) The resulting relations satisfy Third Normal Form (3NF).
- (e) The normal form cannot be determined from the given information.

\*\*\*\*\*