**QP Code:CS2T06-2** **(Pages: 4)**

Reg No : ..... Name : .....

**MAR ATHANASIUS COLLEGE OF ENGINEERING, KOTHAMANGALAM**

(Govt. Aided Autonomous Institution)

Fourth Semester B.Tech. Degree Examination, May, 2026

Course Code: B24CS2T06

Course Name: DATABASE MANAGEMENT SYSTEMS

Max. Marks: 100 Duration:3 Hours

**PART - A: 30 Marks (10 Qns x 3 Marks)**

**Answer all questions**

1. Write briefly about any three relational database integrity constraints.
2. What is meant by a recursive relationship type? Give an example of recursive relationship type.
3. Distinguish between physical records and logical records in the context of file organization in DBMS. Explain with a suitable example why this distinction is important.
4. Compare SELECT and PROJECT operations on relational databases with examples.
5. With an example define primary key, candidate key and super key.
6. List any three aggregate functions in SQL.
7. State Armstrong's three basic axioms. Why are they sound and complete.
8. Write an algorithm to compute the attribute closure of a set of attributes (X) under a set of functional dependencies (F).
9. Define concurrency control in the context of transaction processing. Why is it necessary in a multi-user database system?
10. "If every transaction in a schedule follows the two-phase locking protocol, the schedule is guaranteed to be serializable", justify the statement.

**PART - B: 70 Marks (5 Qns x 14 Marks)**

**Answer any One full question from each module.**

**Module 1**

11. a) Classify and explain the different types of end users in a Database Management System with examples. 7 marks

    b) Classify the following data items into structured, semi-structured, or unstructured: 7 marks

    a) Patient medical history in JSON format

    b) MRI scan image

c) Pharmacy medicine stock table

d) Doctor's voice note about a diagnosis

OR

12. a) Compare database languages DDL, DML, DCL, and TCL with their functions and examples. Which language category would a database administrator use most frequently? 7 marks

b) A small online bookstore currently uses a single centralized computer for its entire inventory database. All tasks—application processing, database queries, and user interface—run on this one machine. 7 marks

(a) Identify three major limitations of this centralized architecture, especially during peak sales periods.

(b) The bookstore plans to upgrade to a client-server architecture. Explain how this architecture differs from the current centralized system.

(c) Describe two specific advantages the bookstore would gain by switching to a client-server architecture.

**Module 2**

13. a) Consider a large, sorted file with 1,200,000 records, where each record is 120 bytes long and disk blocks are 2 KB (2048 bytes). A sparse primary index is built on the ordering key, with each index entry being 12 bytes. Compute step-by-step and construct the multi-level sparse index structure: (8 Marks)

(a) Data blocking factor (records per block)

(b) Total number of data blocks

(c) Index blocking factor (index entries per block)

(d) Number of entries and blocks in the first-level index

(e) Number of levels required in the multi-level index

(f) Total number of index blocks across all levels

b) Explain dynamic hashing. Describe extendible and linear hashing. (6 Marks)

OR

14. a) A file is sorted on the search key 'employee_id' and contains 50,000 records occupying 2,000 disk blocks. (5 Marks)

(a) If you build a primary index on employee_id, will it be dense or sparse? Why?

(b) Approximately how many index entries will exist?

Explain the advantage of sparse indexing in this case.

b) A company maintains a customer database with 20,000 records (record size = 200 bytes) on disk blocks of 2000 bytes. A dense index is built on the customer ID field, where each index entry takes 20 bytes (including key and pointer).Determine: (9 Marks)

(a) How many customer records fit in one data block?

(b) Total number of data blocks required

(c) How many entries will the dense index contain?

(d) How many blocks are needed to store the entire dense index?

**Module 3**

15. a) Describe explicit cursors in PL/SQL, including declaration, opening, fetching, and closing. Apply it by writing a complete PL/SQL block to fetch and print employee details where salary > 50000, using the schema Employees (emp_id INT, name VARCHAR(50), salary DECIMAL(10, 2), dept_id INT). 7 marks

b) In an inventory management system, create a row-level trigger that automatically updates stock quantity in a Products table after an INSERT into Sales. Explain the trigger's execution for each row and discuss a scenario where this ensures data integrity. 7 marks

Schema:

- Products (product_id INT, name VARCHAR(50), stock INT)
- Sales (sale_id INT, product_id INT, quantity INT)

OR

16. a) Employees( emp_id, first_name, last_name, email, phone, department_id, job_title, salary, hire_date) 7 marks

Departments(dept_id, dept_name, location, manager_id)

For the above schema write the queries for the following

(a) Find all employees whose first name starts with 'A'. Display first name, last name, and email.

(b) Find all employees whose last name contains 'me' (case-insensitive) anywhere in the name. Display first name, last name, and email.

(c) Using implicit join syntax display employee name, job title, department name, and location for all employees.

(d) Using implicit join, find all employees working in departments located in 'Mumbai'. Show employee name, department name, and salary. Sort by salary in descending order.

b) Explain the need for SQL joins when retrieving data from multiple tables. Describe the working of an INNER JOIN and OUTER join with an example scenario. 7 marks

**Module 4**

17. a) A library stores data in relation: (8 Marks)

LIBRARY(BookID, Title, AuthorName, AuthorCountry, Publisher, PubCity ) with the following functional dependencies.

- BookID → Title, AuthorName, Publisher
- AuthorName → AuthorCountry
- Publisher → PubCity

Normalize into 3NF.

b) A relation R (A, B, C, D), with FDs: AB→C, C→D is decomposed to $R_1$ (A, B, C), and $R_2$ (C, D). Determine whether the decomposition is lossless and dependency preserving. (6 Marks)

OR

18. What is Query Optimization? Explain the various techniques used by a DBMS to optimize a query. Discuss in detail the difference between Heuristic-based (Rule-based) optimization and Cost-based optimization with suitable examples.

**Module 5**

19. a) What is the two-phase locking (2PL) protocol? How does it guarantee serializability? How strict 2PL differs from basic 2PL? (8 Marks)

b) R1(x), R2(x), R1(z), R3(x), R3(y), W1(x), C1, W3(y), C3, R2(y), W2(z), W2(y), C2; Check whether the given schedule is strict? Justify your answer. (6 Marks)

OR

20. a) Explain BASE properties and interpret how they differ from ACID properties in distributed databases. (7 Marks)

b) A banking core system requires: (7 Marks)

(a) Strict consistency

(b) Transactions

(c) Complex joins

Explain why NoSQL may not be suitable and recommend an appropriate database solution.

**********