BACHELOR OF TECHNOLOGY (CBCS - 2023)
B. Tech. Sem-III INFORMATION TECHNOLOGY : WINTER: 2025
SUBJECT: DATABASE MANAGEMENT SYSTEMS

Day : Monday
Date : 15/12/2025

W-29244-2025

Time : 10:00 AM-01:00 PM
Max. Marks : 60

NB :
1. Assume suitable data, if necessary.
2. Draw neat labelled diagrams WHEREVER necessary.
3. Figures to the right indicate FULL marks for the question.
4. All questions are COMPULSORY.

Q. 1 Consider following Schema (10)

STUDENT(SID, Name, Program, Year)
COURSE(CID, CName, Department)
ENROLLMENT(SID, CID, Grade)
INSTRUCTOR(IID, IName, Department)
TEACHES(IID, CID, Semester)

Solve the following relational algebra statements for the above schema

a. List the names of all students who are enrolled in the "Operating Systems" course.
b. Find the IDs and names of instructors who teach at least one course in the 'Mathematics' department.
c. Determine the names of students who are enrolled in more than one course in the current semester.
d. Retrieve the names of students and the courses they are enrolled in, along with the respective instructor names.
e. List all students who are enrolled in every course taught by a specific instructor (e.g., "Dr. Brown").
f. Find all course IDs (CID) that are not taught in the 'Fall 2024' semester.
g. Identify instructors who do not teach any course currently, using appropriate set operations.

OR

Q. 1 Analyze how logical and physical data independence contribute to the flexibility and maintainability of a database system, using a real-world example. (10)

Q. 2 Inspect the impact of violating referential integrity in a database. Use an example of student-course enrollment database to illustrate the problem. (10)

OR

Q. 2 Investigate a decomposition of a relation schema and determine whether it is dependency preserving. Justify your steps. (10)

Q. 3 Analyze the result sets produced by different join types (INNER, LEFT, RIGHT, FULL OUTER) on the Employees and Departments tables and explain their practical implications in reporting. (10)

OR

Q. 3 Analyze the differences between procedures and functions in PL/SQL in terms of usage, (10)

return types, and side effects. Provide example use cases for each.

Q. 4 Examine the impact of join strategies on the optimization process. Specifically, discuss how the choice between hash join, nested loop join, and sort-merge join can affect query execution time for complex queries. (10)

OR

Q. 4 Scrutinize the performance trade-offs of using hash indexing in comparison to B+ tree indexing for a query that searches for employees by Name in a large employee database. (10)

Q. 5 Given a scenario where two transactions are executed concurrently, analyze the differences between locking-based concurrency control and timestamp-based concurrency control. Discuss their advantages and disadvantages in terms of deadlock prevention and transaction throughput. (10)

OR

Q. 5 Explore the role of concurrency control in ensuring the isolation property of transactions. Discuss the impact of different isolation levels (Read Uncommitted, Read Committed, Repeatable Read, Serializable) on data consistency and transaction performance. (10)

Q. 6 Inspect the differences between Discretionary Access Control (DAC), Mandatory Access Control (MAC), and Role-Based Access Control (RBAC) models. Discuss the scenarios where each model is most suitable. (10)

OR

Q. 6 Given a scenario in which large amounts of unstructured data need to be processed from various sources (such as audio, video, and text), analyze the potential challenges in handling this data using Hadoop and NoSQL systems. (10)

* * * * * *