Exam. Code: 0919
Sub. Code: 33473

2125

B.E. (Computer Science and Engineering)
Seventh Semester
CS-702: Advance Database Systems

Time allowed: 3 Hours Max. Marks: 50

NOTE: Attempt five questions in all, including Question No. I which is compulsory and selecting two questions from each Unit.

x-x-x

I. Attempt the following:

a) Define strict two-phase locking.
b) Explain BCNF with an example.
c) What is multi version concurrency control?
d) Explain Shadow Paging with example.
e) Define Data Marts with the use of an example. (5x2)

UNIT - I

II. a) Explain heuristic-based query optimization with considerable heuristics. How the heuristic rules are applied? Explain with the use of SQL query example?

b) Explain multiple granularity level locking with suitable diagram and example. Also explains the types of intention locks used in DBMS for a transaction at various levels. (2x5)

III. a) A university maintains the database containing the following attributes for each course offered: ***CourseID, CourseName, InstructorID, InstructorName, Department, RoomNo, TimeSlot, StudentID, StudentName, InstructorPhone***

i) Identify and explain all the functional dependencies.
ii) Explain each step showing the normalization of the relation into 3NF, BCNF and 4NF.

b) Explain the Object-Oriented Data Model with examples of objects, classes, and inheritance. Also describe the major enhancements introduced in SQL3 to support object-relational features. (2x5)

IV. a) Apply all the normal forms upon the following database system and describe the update anomalies for each.

Borrowing*(Borrow_ID, Member_ID, Member_Name, Bood_ID, Book_Title, Book_Price, Borrow_Date, Member_City)*

Functional Dependencies (FDs):

Borrow-ID → Member-ID Member-ID → Member-Name Member-ID → Member-City
Borrow-ID → Book-ID Book-ID → Book-Title Book-ID → Book-Price
Borrow-ID → Borrow-Date

b) A banking system has three interleaved transactions:

T1: R(A), W(B), R(C), W(A)
T2: W(A), R(B), W(C)
T3: R(C), W(B)

Contd.....P/2

Answer the following questions:

i) Construct the conflict graph and find out if the schedule is conflict serializable.

ii) Explain the outcome under Strict 2PL, Timestamp Ordering, and Multiversion concurrency control and find out which transactions are commit or abort. Provide justification. (2x5)

UNIT - II

V. a) A multinational company maintains a distributed database across three sites in India – Delhi, Mumbai, and Bangalore

i) Propose a fragmentation strategy with justification.

ii) Explain how distributed concurrency control is maintained across these sites using either two-phase commit or distributed locking.

b) A cloud database fails three times in one day of three individual events such as a server crash occurred during a long update transaction, a media failure corrupted a segment storing index pages and a power failure happened immediately after a check point. The system uses Write-Ahead Logging and mixed of Deferred with immediate update, and shadow paging for some tables. Explain the following:

i) How transaction recovery is performed for the update transaction using Undo or Redo rule?

ii) Describe the media recovery process including when full backup, differential backup and log replay are used.

iii) Compare recovery under Shadow Paging and Log-based recovery for this scenario. (2x5)

VI. a) What is data Mining? Explain the data mining process in detail with each individual step including cleaning, integration, selection, transformation, pattern discovery, and evaluation. Also differentiate between OLTP, OLAP, and Data Warehouse.

b) Compare various concurrency control techniques on the basis of conflict serializability, deadlock avoidance, recoverability, and cascade lessness, with suitable examples. (2x5)

VII. a) Explain the architectural differences among Microsoft SQL Server, Oracle, IBM DB2, and MySQL along with their features.

b) What is data Cube? How it used in data warehouse? Explain with example. (2x5)

x-x-x