

Reg. No. :

Question Paper Code : 20416

B.E./B.Tech. DEGREE EXAMINATIONS, APRIL/MAY 2022.

Third/Fourth/Fifth/Eighth Semester

Computer Science and Engineering

CS 8492 – DATABASE MANAGEMENT SYSTEMS

(Common to : Computer and Communication Engineering/
Mechanical and Automation Engineering/Mechatronics Engineering/
Computer Science and Business Systems/Information Technology)

(Regulations 2017)

Time : Three hours

Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. List out the different types of data models.
2. What are the different types of integrity constraints used in relational databases?
3. State the types of attributes in ER Model.
4. Define join dependencies.
5. List out the timestamp based deadlock prevention schemes.
6. What is the use of save points in recovery?
7. When it is preferable to use a dense index rather than a sparse index?
8. What are the steps involved in query processing?
9. Specify the types of fragmentation in distributed databases.
10. Give the significance of XQuery.

PART B — (5 × 13 = 65 marks)

11. (a) Illustrate the database system architecture with a neat diagram.

Or

- (b) Consider the relational database :

employee(empname, street, city)

works(empname, companynname, salary)

company(companynname, city)

manages(empname, managername)

Write each of the following queries in Relational Algebra and SQL.

- (i) Find the names, street addresses, and cities of all employees who work for Indian Bank Corporation and earn more than 10 lakhs per annum. (2 + 2)
- (ii) Insert the tuple ('Siva', Gandhi Street, Chennai) in the employee database. (2 + 1)
- (iii) Delete all employees who work for Indian Bank Corporation. (2 + 1)
- (iv) Give all managers in this database a 15 percent salary raise. (2 + 1)
12. (a) (i) Construct an E-R diagram for a car insurance company whose customers own one or more cars each. Each car has associated with it zero to any number of recorded accidents. (9)
- (ii) Convert the designed ER into a relational design. (4)

Or

- (b) (i) Consider a relation schema $R = (A, B, C, D, E)$.

$F = \{A \rightarrow BC, CD \rightarrow E, B \rightarrow D, E \rightarrow A\}$. Show the decomposition of R into ABC and ADE is lossless join and dependency preserving if the above set F of functional dependencies holds. (6)

- (ii) Explain the fourth normal form with an example. (7)

13. (a) Explain the two phase locking protocols for concurrency control.

Or

- (b) Define Transaction. Represent the transaction states and explain the various desirable properties of a transaction. (2 + 5 + 6)

14. (a) (i) Explain the RAID levels with the nature of applications supported by each level. (8)
(ii) Illustrate how the B+ tree structures are used for indexing files. (5)

Or

- (b) Let relations $r_1(A, B, C)$ and $r_2(C, D, E)$ have the following properties : r_1 has 20,000 tuples, r_2 has 45,000 tuples, 25 tuples of r_1 fit on one block, and 30 tuples of r_2 fit on one block. Estimate the number of block accesses required, using each of the following join strategies for $r_1 \bowtie r_2$:
- (i) Nested-loop join. (4)
(ii) Block nested-loop join. (4)
(iii) Merge join. (5)
15. (a) Explain the reference architecture and transaction processing of distributed database system. (7 + 6)

Or

- (b) Write short note on following :
- (i) XML databases (7)
(ii) ODMG Model of ODL and OQL. (6)

PART C — (1 × 15 = 15 marks)

16. (a) (i) Consider the following relation schema LOTS which describes parcels of land for sale in various countries of a state. Suppose that there are two candidate keys: $Property_id\#$ and $\{Country_name, Lot\# \}$; that is, lot numbers are unique only within each country, but $property_id\#$ numbers are unique across countries for the entire state.
- LOTS($Property_id\#, Country_name, Lot\#, Area, Price, Tax_rate$)
Based on the given primary key, is the relation in 1NF, 2NF, 3NF and BCNF, 4NF? Why and why not? How would you successfully normalize it completely? (10)
- (ii) Why 3NF is preferable than BCNF? (5)

Or

- (b) (i) Consider the following two transactions :

T1 : read(A); read(B); if A=0 then B:=B+2; write(B).

T2 : read(B); read(A); if B=0 then A:=A+2; write(A).

Add lock and unlock instructions to the transactions T1 and T2, so that they observe the two-phase locking protocol. Can the execution of these transactions result in deadlock? (8)

- (ii) Consider the following log corresponding to a particular schedule at the point of a system crash for four transactions T₁, T₃, T₅ and T₇. Suppose that we use the general immediate update protocol with check pointing. Describe the recovery process from the system crash. Specify which transactions are ROLLEDBACK, which operations in the log are REDONE and which (if any) are UNDONE, and whether any cascading ROLLBACK takes place. (7)

<T₁, Start>

<T₁, D, 20, 25>

<T₁, Commit>

<Checkpoint{T₁}>

<T₃, Start>

<T₃, B, 12, 18>

<T₅, Start>

<T₅, D, 25, 15>

<T₇, Start>

<T₇, C, 30, 40>

<T₇, A, 30, 20>

<T₇, commit>

<T₅, D, 15, 25> ← System Crash