

Code No.: DS743PE

R20

H.T.No.

8 R

CMR ENGINEERING COLLEGE: : HYDERABAD

UGC AUTONOMOUS

IV–B.TECH–I–Semester End Examinations (Supply) – April – 2025

DISTRIBUTED DATABASES

(CSD)

[Time: 3 Hours]

[Max. Marks: 70]

Note: This question paper contains two parts A and B.

Part A is compulsory which carries 20 marks. Answer all questions in Part A.

Part B consists of 5 Units. Answer any one full question from each unit. Each question carries 10 marks and may have a, b, c as sub questions.

PART-A

(20 Marks)

1. a) What is data allocation in the context of DDBMS? [2M]
- b) List the major problems in DDBMS. [2M]
- c) What is the purpose of data localization in query processing? [2M]
- d) What is query decomposition? [2M]
- e) List the four properties of a transaction. [2M]
- f) Identify two types of transactions commonly used in database systems. [2M]
- g) Define the terms: Reliability and Availability. [2M]
- h) Why fault tolerance is important in distributed systems? [2M]
- i) List two key features of the Object-Oriented Data Model (OODM) [2M]
- j) Define an object in the context of object-oriented databases. [2M]

PART-B

(50 Marks)

2. a) Discuss the design issues of Distribute Databases. [5M]
- b) Explain the challenges in distributed systems. [5M]

OR

3. Demonstrate how data fragmentation can improve performance in distributed databases with suitable example. [10M]
- 4.a) Compare and contrast the steps involved in centralized vs. distributed query optimization. [5M]
- b) Analyze the factors that influence the choice of algorithms in distributed query optimization. [5M]

OR

5. Illustrate generic layering scheme for query processing in DDBMS with diagram and explain the functions of each layer. [10M]
6. Analyze the impact of deadlock on distributed transaction processing and explain how deadlock detection mechanisms work. [10M]

OR

- 7.a) Discuss in detail about the types of the transaction models. [5M]
- b) Illustrate the algorithm for basic Time Stamp Ordering algorithm. [5M]
8. Examine the impact of site failures on data consistency and availability in a distributed database. [10M]

OR

9. Analyze the challenges of load balancing in a parallel database system and suggest possible solutions. [10M]
10. Demonstrate with a scenario where object query processing would be essential in a distributed system. [10M]

OR

11. Apply the concept of persistence in an object-oriented database, showing how it allows objects to maintain state. [10M]
