

ADVANCE DATABASE MANAGEMENT SYSTEMS

Time : 3 Hours Total Mark : 100

Note :-Attempt all questions.

1. Attempt any two parts of the following: (10x2=20)

(a) What are the problems that can arise during concurrent execution of two or more transactions ?

Discuss methods to prevent or avoid these problems.

(b) What are serializable schedules of transactions ? Differentiate between conflict and view serializability

(c) (i) Discuss the rules to be followed while preparing a serializable schedule. Why should we prefer serializable schedules instead of serializable schedules ?

(ii) What are cascadeless schedules ? Discuss

2. Attempt any two parts of the following: - (10x2=20)

(a) How do we enforce serializability by locks ? Discuss whether preparing a serializable schedule is preferable over enforcement of serializability by locks.

(b) What is the importance of selection of granularity of data items in lock based protocols ?

What is the effect of granule size over the performance of transaction processing ? Explain in detail.

(c) Discuss the time stamp based protocols to maintain serializability in concurrent execution.

3. Attempt any two parts of the following: (10x2=20)

(a) What is the purpose of fragmentation in a distributed database ? Describe the various types of fragmentation methods.

(b) How is transaction management in distributed database different from transaction management in a standalone database ? Describe distributed transaction management.

(c) Describe the following in brief:

(i) Replication techniques      (ii)      Two phase locking algorithm.

4. Attempt any two parts of the following: (10x2=20)

(a) Describe the process of recovery in a distributed database system. What are the various issues in recovery ?

(b) Describe log based recovery in concurrent transactions. How is it different from recovery in message passing systems ? Describe recovery line.

(c) Discuss the concepts of orphan and inconsistent messages. Illustrate with example.

5. Attempt any two parts of the following: (10x2=20)

(a) Describe cost estimation of query in a distributed database. How does replication in a distributed database effects update query ?

(b) Describe query optimization in a distributed database. How is it different from query optimization in a standalone database ?

(c) Describe following in brief:

(i) Distributed deadlock detection

(ii) Eager replication technique