Printed pages: 02 **Sub Code: ECS076**

Paper Id: | 1 | 0 | 1 | 5 | **Roll No.** | | | | | | | | | | |

B.Tech.
(SEM VII) THEORY EXAMINATION 2017-18
DISTRIBUTED DATABASE

Time: 3 Hours ***Total Marks: 100***

Note: 1. Attempt all Sections. If require any missing data; then choose suitably.

SECTION A

1. Attempt *all* questions in brief. 2 x10 = 20

a. Define the states of transaction?
b. Why testing of serializability is used?
c. Define granularity of a lock.
d. Explain Rigorous 2 phase locking protocol.
e. Explain concurrency control in distributed database.
f. How the transactions are managed in distributed database.
g. What are checkpoints?
h. Explain the concept of inconsistent messages.
i. What is distributed deadlock detection?
j. Explain the objectives of distributed query processing.

SECTION B

2. Attempt any *three* of the following: 10 x 3 = 30

a. Discuss Recoverable schedule & also explain cascading rollback.
b. Describe the architecture of locking scheduler in brief?
c. Describe 3 phase commit protocol? How 3PC is different from 2PC?
d. Explain the following in detail: (i) Orphan messages (ii) Problem of livelocks
e. Explain Edge chasing algorithm for distributed deadlock detection

SECTION C

3. Attempt any *one* part of the following: 10 x 1 = 10

(a) (i) Differentiate between conflict & view serializability in detail.
(ii) Differentiate between distributed database & centralized database.
(b) What are schedules? Describe the concepts in recoverable & cascadeless schedules.

4. Attempt any *one* part of the following: 10 x 1 10

(a) (i) Differentiate between fine granularity & coarse granularity.
(ii) Explain the working of 2 phase locking protocol in brief
(b) Discuss the working of time stamp based protocols? Also explain how a unique global tike stamp is generated in distributed system.

5. Attempt any *one* part of the following: 10 x 1 = 10

(a) (i) Explain moss concurrency protocol.
(ii) Describe the correctness rules that must be considered during data fragmentation.
(b) Discuss the locking techniques for concurrency control in detail.

6. Attempt any *one* part of the following: 10 x 1 = 10

(a) (i) Explain the type of failure in distributed database.
(ii) Explain issues of recovery in distributed database.

(b) Generate an algorithm for synchronous check pointing in a Distributed Database System.

7. Attempt any *one* part of the following: 10 x 1 = 10

(a) (i) Differentiate between multiway joins & semi joins.
(ii) Differentiate between Eager & Lazy replication

(b) Explain Ho-Ramamoorthi algorithm for deadlock detection. What are the phantom deadlocks? Does this algorithm detect the phantom deadlock?