

CS 54200: Distributed Database Systems
Professor Bharat Bhargava
Homework 1
Due date: January 30, 2020 11:59 PM (EST)

Note: *Submit the written homework via Blackboard. Answer each question briefly.*

Max. Marks: 100

Q1: Write a one-page summary of ideas for three papers identifying the problem and the solution (1-page per paper, 3 pages in total). You may select papers from the [reading list](#), or choose papers outside the list, that you want to read, related to *distributed databases* (such as implementations, concurrency control, reliability, communication network, big data, security/privacy, mobile) or *cloud computing*. Three papers that can give you a quick background are: (30 points)

- a) [The Raid Distributed Database System](#), Bharat Bhargava and John Riedl, IEEE Trans on Software Engineering, 15(6), June 1989.
- b) [Concurrency Control in Database Systems](#) Bharat Bhargava, IEEE Trans on Knowledge and Data Engineering, 11(1), Jan.-Feb. 1999
- c) [Building Distributed Database Systems](#), Bharat Bhargava.

Q2: Give examples of transaction executions that lead to: (10 points)

- a) lost update
- b) dirty read.

Q3: What makes the problems in a distributed database system different (or complex) as compared to the problems in a single site system? List at least four differences. (20 points)

Q4: Is there any sense in running a: (10 points)

- a) **'Read only'** transaction?
- b) **'Write only'** transaction?

Q5: Discuss the role of the following issues in distributed database system architecture: (15 points)

- a) Autonomy
- b) Heterogeneity
- c) Scalability.

Q6: What are advantages and disadvantages of a replicated database compared with a nonreplicated database? (15 points)