CS54200: Distributed Database Systems
Prof. Bharat Bhargava

Homework 3
Due date: March 30, 2017

Note*: Written homework assignments are due at the beginning of class on the due date. Answer each question briefly.*

Q1: Design a central site termination protocol. This requires that the operational sites agree on a coordinator. Suggest a mechanism that only one coordinator is elected. (Hint: Rank the sites and elect the site with the highest ranking as the coordinator)

Q2: Give two examples of the application of commit and termination protocols in a distributed system. (Hint: Mail facility is one such example. If you use this example, explain how would you make changes in the current mail facility that exists on our system across ITAP and CS machines.)

Q3: Imagine the situation in 2PC protocol when one participant sends a "vote-abort" message while another one sends "global-commit". Is this possible? Why?

Q4: In the centralized 3PC termination protocol, the first step involves sending the coordinator's state to all participants. The participants move to new states according to the coordinator's state. It is possible to design the termination protocol such that the coordinator, instead of sending its own state information to the participants, asks the participants to send their state information to the coordinator. Modify the termination protocol to function in this manner.

Q5:

(a) What are the problems in a commit protocol if a failure of a participating site occurs? Give an example to illustrate the problem.
(b) How can the simple termination protocol be made robust? Give the summary of messages that are needed to abort, and commit the transaction.
(c) Define the concurrency set of a local state in a commit protocol.

Q6:

(a) What are session numbers and nominal session numbers? What role they play in recovery?
(b) What are control transactions? What is their role?

Q7: Explain the various ideas of determining a 'unique majority partition' that can continue to process transactions during a network partition. Consider cases where multiple partitions and merging may take place. Suggest what useful processing can be done in minority partitions so that the merge process is quick and efficient. How can the size of the partition be determined at the time of partition and how can we determine if the mutual consistency has been violated at the time of merge?