Universität Freiburg
Institut für Informatik
Prof. Dr. P. Fischer/Prof. Dr. F. Kuhn

Georges-Köhler Allee, Geb. 51
D-79110 Freiburg
peter.fischer@informatik.uni-freiburg.de

Exercises
Distributed Systemes: Part 2
Summer Term 2015
21.7.2015

6. Exercise sheet: Distributed Concurrency Control and Replication

Exercise 1
Consider the following local schedules:

- $S_1: \quad R_1A \quad W_1A \quad R_2A \quad W_2A$
  $S_2: \quad R_2B \quad W_2B \quad R_1B \quad W_1B$
- $S_1: \quad R_1A \quad W_2A$
  $S_2: \quad R_3B \quad W_1B \quad R_2C \quad W_3C$
- $S_1: \quad R_1A \quad R_3A \quad R_3B \quad W_3A \quad W_3B \quad R_2B$
  $S_2: \quad R_4D \quad W_4D \quad R_1D \quad R_2C \quad R_4C \quad W_4C$
- $S_1: \quad W_1A \quad c_1 \quad R_3A \quad R_3B \quad c_3 \quad W_2B \quad c_2$
  $S_2: \quad W_2C \quad c_2 \quad R_4C \quad R_4D \quad c_4 \quad W_1D \quad c_1$

(1) Verify whether or not the schedules are serializable.
(2) Demonstrate that by applying Distributed 2PL/Timestamp Protocol prevents non-serializable schedules.
(3) Check whether or not the schedules are rigourous and commit-deferred.
(4) Demonstrate that by applying a Ticket-based concurrency control prevents non-serializable schedules.

Exercise 2
Keeping consistency in replicated data is a key issue, for which several approaches exist

a) Compare the combinations of update primary copy/update anywhere and eager/lazy propagation in terms of availability, consistency and cost for read/write operations
b) What kind of consistency problems could occur with a read quorum $\frac{2}{3}$N+1 and a write quorum of N/3+1?

Exercise 3
Eventual consistency provides high availability and scalability, but limits consistency

a) Provide examples of consistency problems/anomalies that could occur!
b) In current cloud storage systems, *Latest write wins* is a popular approach to resolve concurrent updates. Explain the problems that may occur when using physical/wall-clock timestamps!
c) Describe an approach that uses logical clocks to handle such concurrent updates

Exercise 4
Different consistency models provide different tradeoffs between availability and consistency

a) Explain why preventing lost updates can lead to unavailability
b) How can you guarantee Read Committed, but stay available?