

Excercise Sheet 11

Solution

Lecture Distributed Systems

Winter Term 2025/26

Exercise 1: Two-Phase Commit

A participant could wait in its INIT status for a VOTE-REQUEST message from the coordinator. If this message is not received after a certain time, the participant simply decides to cancel the transaction locally and sends a VOTE-ABORT message to the coordinator.

Similarly, the coordinator can block in the WAIT state, where it waits for the other participants to vote. If not all votes can be determined within a certain time, the coordinator should also decide to cancel the transaction and send a GLOBAL-ABORT to all participants.

The third status in which it is possible to block is the READY status of the participants. Here you could block because a participant is waiting for the global voting result sent by the coordinator. If this message is not received within a certain time, the participant cannot simply decide to cancel the transaction. Instead, it must determine which message the coordinator has sent. The easiest way to solve this problem is to block each participant until the coordinator is up and running again.

Exercise 2: Causal Consistency

- a) P_2 : Reading x and writing y creates a causal dependency. I.e., $W_1(x)_1$ is causally before $W_2(y)_2$. But to explain the behavior at P_3 , the global order would have to be $W_2(y)_2, R_3(y)_2, R_3(x)_0, W_1(x)_1$, which contradicts this ($W_1(x)_1$ is after $W_2(y)_2$).
- b) The causal consistency may be sufficient. The problem is that reactions to changes in stock values should be consistent. Changes to independent stock values can be displayed in different order.

Exercise 5: Replication

- a) The three most important points that speak in favor of generating replicas are reliability, data safety, and access speed.
- b) Types of replicas:
 - Permanent: permanently available replicas - initial set of replicas that make up a distributed data store (e.g. mirroring of web pages, replicated DNS servers).
 - Server-initiated: initiated on demand by the server - copies of a data store that exist to enhance performance and which are generated by the initiative of the owner of the file storage (e.g. web hosting services, content delivery networks).
 - Client-initiated or client caches: created by the client - local storage facility used by the client to temporarily store a copy of the previously requested data.

Exercise 6: Active Replication

No. For example, take read operations that take place on unmodified data, or commutative write operations. In principle, these situations allow different sequences to exist on different replicas. However, it can be difficult, if not impossible, to determine whether, for example, two write operations are commutative.

Exercise 8: Any Questions?