Chalmers Tekniska Högskola
Data- och Informationsteknik

CTH TDA297
GU DIT290

Distribuerade system fk

Tentamen 2014-03-12

Dag, Tid, Sal: March 12th 2014, 08:30-12:30, M building

Kursansvarig: Philippas Tsigas (Tel: 772 5409)

Hjälpmedel: Inga

Totalt Poängtal: 60

Betygsgränser:

CTH: 3:a 30 p, 4:a 38 p, 5:a 48 p

GU: Godkänd 30p, Väl godkänd 48 p

Instructions

- Please answer in English, if possible.
  If you have very big difficulty with that, though, you may answer in Swedish.
- **Do not forget to write your personal number and if you are a GU or CTH student and at which "linje".**
- Please start answering each assignment on a new page; number the pages and use only one side of each sheet of paper.
- Please write in a tidy manner and explain (briefly) your answers.
- Students must **not** write their personal number on the answer sheets since the exam is anonymous; they shall write that **only** on the name slip area that they will seal.

LYCKA TILL !!!!

1. *10 marks*

   a) Give the definitions of Linearizability and Sequential Consistency.

   b) A correctness property $\mathcal{P}$ is compositional if, whenever each object in the system satisfies $\mathcal{P}$, the system as a whole satisfies $\mathcal{P}$. Is Sequential Consistency compositional? If yes explain why, if no give an example.

2. *10 marks*

   Describe the differences between the three-phase commit protocol and the two-phase commit protocol. Draw the three-phase commit protocol as a state machine where you also describe the behavior of the protocol when time-outs are triggered and processes are recovered after crashing.

3. *5 marks*

   What are the properties that a causal broadcast must satisfy?

   Compare the causal broadcast property with the following property: "if a process delivers messages m_1 and m_2, and $m_1 \rightarrow m_2$, then the process must deliver m_1 before m_2."

4. *10 marks*

   a) Describe an algorithm for leader election on a ring topology. All processes have unique ids and the system is asynchronous. Any process can initiate the leader election protocol at any time.

   b) Compute the time complexity and communication complexity of your algorithm.

   c) Is it possible to design a symmetric algorithm for leader election? If yes, provide such an algorithm, if no provide a proof.

5. *15 marks*

   a) Eric wants to build a replicated storage system. In his system there is only one client. The client performs just one storage operation (read or write) at a time, waiting for each operation to complete before starting the next. Eric wants availability even in the case of one server failure. Because of that he decides to store the data on two servers. Eric wants to ensure that the data on the two servers stay identical all the time. In order to guarantee that he is thinking making the writes at the two servers atomic using three-phase commit, to ensure both-or-nothing behavior. In the design he has in mind, the client acts as the transaction coordinator. The client would execute a write as described in the three-phase commit protocol. The system uses the timeout recovery scheme explained in the lectures. Eric thinks about this design for a while, and eventually realizes that three-phase commit is fundamentally not suited to providing *availability* via replication. Please explain why. Eric wants to formalize the consistency properties of the replication system. Is it sequential consistent? Is it a linearizable one? Please explain your answer.

   b) Eric decides to change his design and uses now the gossip architecture that lazily synchronizes the two servers. The single client contacts any server that is available and gets the value that this server has, updates are also performed on the first availiable server first and then lazily propagated on the next server. Each client request uses a unique id.

What is the availability that he can achieve? Eric wants to formalize the consistency properties of the replication system in case where there are no failures. Is it sequential consistent? Is it a linearizable one? Please explain your answer.

c) Eric decides to change his design and uses now nine replicas and quorums to ensure strong consistency and availability. What are the constraints on the sizes of read and write quorums? What is the availability that he can achieve? Give an example of quorum processing involving a write, two reads, then a write. Eric wants to formalize the consistency properties of the replication system. Is it sequential consistent? Is it a linearizable one? Please explain why?

6. *10 marks*

The algorithm by Choy and Singh uses two Doorways. The Asynchronous and the Synchronous Doorway. Describe both Doorways in pseudo-code and informally.

If we use only the Asynchronous Doorway together with the coloring, is this algorithm going to give us a solution to the resource allocation problem (e.g. is it going to guarantee mutual exclusion and no-starvation)? Please provide a proof sketch or a counterexample. What is the time complexity of this algorithm?