

Distribuerade system fk
Tentamen 2018-03-14

Dag, Tid, Sal: March 14th 2018, 08:30-12:30, SB-MU building

Kursansvarig: Philippos 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. (13 points) Suppose you have a protocol for a synchronous message-passing ring that is anonymous (all processes run the same code) and uniform (this code is the same for rings of different sizes), i.e. symmetric. Suppose also that the processes are given inputs marking some, but not all, of them as leaders. Give an algorithm for determining if the size of the ring is odd or even, or show that no such algorithm is possible.

Clarification: Assume a bidirectional, oriented ring and a deterministic algorithm.

Hint: You might want to consider executions of different ring-sizes together.

2. (4+1+1 points) Data replication is a very important area in distributed systems. Replicating data may boost performance, and also be used for masking failures. There are two common schemes for data replication: *active* and *passive*.
 1. Describe both the active and the passive replication scheme, including the differences between them.
 2. Briefly compare the two schemes from a performance point of view.
 3. Which replication scheme is more appropriate to use for replicating sensitive data in a system that can experience arbitrary failures?
3. (11 points) Suppose that we modify the problem of synchronous agreement with crash failures so that instead of crashing a process forever, the adversary may jam some or all of its outgoing messages for a single round. The adversary has limited batteries on its jamming equipment, and can only cause f such one-round faults. There is no restriction on when these one-round jamming faults occur: the adversary might jam f processes for one round, one process for f rounds, or anything in between, so long as the sum over all rounds of the number of processes jammed in each round is at most f . For the purposes of agreement and validity, assume that a process is non-faulty if it is never jammed. Can we modify the round based algorithm that we discussed to work in this model? If yes, as a function of f and n , how many rounds does it take to reach agreement in the worst case in this model, under the usual assumptions that processes are deterministic and the algorithm must satisfy agreement, termination, and validity? If no, provide a proof.
4. (10 points) Suppose that you have a bidirectional but not necessarily complete asynchronous message-passing network represented by a graph $G = (V, E)$ where each node in V represents a process and each edge in E connects two processes that can send messages to each other. Suppose further that each process is assigned a weight 1 or 2. Starting at some initiator process, we'd like to construct a shortest-path tree, where each process points to one of its neighbors as its parent, and following the parent pointers always gives a path of minimum total weight to the initiator. Give a protocol that solves this problem, show that it works and describe its complexity.
5. (10 points) Moa is building a shared object that is replicated on 2 servers, server S and server N. To access the object, all client from the south accesses the local copy at server S and all clients from the north access the local copy at server N. The two servers start from the same initial state. There is no communication and synchronization between the two servers. If there is no mobility, i.e. no client from the south can move to the north or vice versa, what is the consistency (linearizability, sequential consistency, other) that the scheme provides? Provide a proof sketch. If there is mobility what is the consistency that this replication scheme provides.

6. (10 points) Each statement is either true or false. A correct answer gives 1 point, a wrong answer gives -1 point, no answer gives 0 points. Overall you cannot get less than 0 points for this question.

1. Leader election can be solved with a consensus algorithm.
A. True B. False
2. Atomic Broadcast is weaker than Consensus.
A. True B. False
3. In an asynchronous system, where only one processor might crash, consensus is solvable.
A. True B. False
4. Sequential Consistency is composable.
A. True B. False
5. The Gossip architecture for replication was designed to provide highly available service.
A. True B. False
6. The 3 Phase Commit protocol was introduced to improve the latency of the 2 Phase Commit protocol in executions where no faults take place.
A. True B. False
7. There exist a consensus solution able to handle up to $f < n/3$ Byzantine failures.
A. True B. False
8. 3PC was introduced to handle undetected message losses.
A. True B. False
9. Any transactional service is linearizable if each request is executed at most once.
A. True B. False
10. A system can be linearizable but not sequential consistent.
A. True B. False