Chalmers Techniska Högskola
Data- och Informationsteknik

CTH TDA297
GU DIT290

Distribuerade system fk

Tentamen 2014-08-22

Dag, Tid, Sal: August 22th 2014, 14:00-18:00, V 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. *15 marks*

   (a) Define the specification of the i) Reliable Broadcast, ii) FIFO Broadcast and iii) Causal Broadcast.

   (b) Describe three algorithms that implement respectively the three different types of Broadcast mentioned above in an asynchronous system, with process crash failures.

   (c) What is the time and communication complexity of your FIFO Broadcast Algorithm? Provide a complexity proof.

2. *10 marks*

   Describe an algorithm that computes a spanning tree of a network $G(V, E)$. How a node of the network can use the existence of such a spanning tree in order to broadcast information to all nodes of the network?

3. *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.

4. *10 marks*

   Describe the two generals problem. Can you find a solution to the problem? If yes describe your solution and provide a proof of its time and communication complexity. If not provide a proof that the problem is not solvable.

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?

Chalmers Techniska Högskola
Datavetenskap

CTH TDA297
GU DIT290

Distribuerade system fk

Tentamen 2013-03-16

Dag, Tid, Sal: March 16th 2013, 14:00-18:00, V 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) Give an example pseudocode for a simple program that would operate correctly with sequential consistency but incorrectly with linearizability. (Briefly explain your answer.)

c) Describe the advantages and disadvantages of using each for an application, paying particular attention to any assumptions you must make about the users' behaviour. As two example applications you can consider:

- a web site allowing registered users to upload videos and allowing anybody to download and watch them. (Comments on videos are not supported.)
- a multi-user game, the players move figures around a common scene. The state of the game is replicated at the players' workstation and a server. The figures may through projectiles at one another, and a hit debilitates the unfortunate recipient for a limited time. (Application II is from Chapter 18th of the course Book.)

2. *10 marks*

a) What are FIFO ordering, causal ordering and total ordering guarantees of replica managers? How are they related to each other?

b) Compare the causal ordering property with the following property: "if a replica manager processes message m_1 and m_2, and $m_1 \rightarrow m_2$, then the replica manager must process m_1 before m_2."

3. *10 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 two-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 two-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 two-phase commit is fundamentally not suited to providing *availability* via replication. Please explain why.

b) Eric decides to change his design and use 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.

4. *10 marks*

   A distributed conference application provides a shared whiteboard. Each member of the conference has a replica of the whiteboard that is managed by a specific member of a closed process group.

   Describe an approach that uses the algorithm by Choy and Singh to achieve mutually exclusive access to the whiteboard, prior to propagation of the updates to the whole group.

   If the number of members of the group is n, what is the time complexity of the scheme, provide a brief proof?

5. *5 marks*

   Describe an algorithm that solves consensus in a synchronous system where up to f of the n processes exhibit crash failures, assume the existence of a reliable multicast protocol. What is is the time complexity of the algorithm.

6. *10 marks*

   The flooding algorithm is a straight forward reliable broadcast algorithm. The initiator sends to all its neighbors a message of kind broadcast. When a process receives message broadcast for the first time, it sends to all other adjacent processes further broadcast messages.

   a) What are the properties of reliable broadcast.

   b) What is the Time and Message Complexity of the algorithm in fault-free asynchronous executions?

   c) What is the Time and Message Complexity of the algorithm in an asynchronous execution in the presence of crash faults? (No partitions.)

   Please provide brief complexity analysis.

7. *5 marks*

   Consider it as given that there is no Byzantine agreement protocol for a system with 3 processes, from which 1 of them faulty. Using this proof prove that there is no Byzantine agreement protocol for a system with $n \geq 3$ processes, from which f are faulty, if $n \leq 3f$.

Chalmers Techniska Högskola
Datavetenskap

CTH TDA297
GU DIT290

Distribuerade system fk

Tentamen 2011-08-23

Dag, Tid, Sal: August 23rd 2011, Fm, M

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*

   Construct a solution to reliable, totally ordered multicast in a synchronous system, using a reliable multicast and a solution to the consensus problem.

2. *5 marks*

   (a) Can Byzantine agreement be reached for 8 generals, with 3 of them faulty?

   (b) Can Byzantine agreement be reached for 8 generals, with 3 of them faulty, if the generals digitally sign their messages?

3. *8 marks*

   A three-phase commit protocol has the following parts

   Phase 1: is the same as the two-phase commit.

   Phase 2: the coordinator collects the votes and makes a decision; if it is No, it aborts and informs participants hat voted Yes, if the decision is Yes, it sends a precommit request to all participants. Participants that voted Yes wait for a precommit or doAbord request. They acknowledge precommit requests and carry out doAbort requests.

   Phase 3: the coordinator collects the acknowledgements When all are received, it commits and sends a do commit to he participants. Participants wait for a doCommit request. When it arrives they Commit.

   Explain how this protocol avoids delay to participants during their "uncertain" period due to the failure of the coordinator or other participants. Assume that communication does not fail.

4. *10 marks*

   A processor P is part of a network $G(V, E)$. P believes that processor Q is also connected to the network. Describe a protocol that P together with the other processors of the network can use in order to find Q or to find that Q is not part of the network. Prove the time complexity of the algorithm.

5. *15 marks*

   Give a solution to the dinning philosophers problem. Prove the time complexity of the algorithm.

6. *12 marks :*

   A quorum consensus method with three replicas is used for replication. The parameters of the replication system are described in the following table. Calculate: i) the read latency, ii) the write latency, iii) the blocking probability of a read and iv) the blocking probability of a write.

Chalmers Techniska Högskola
Datavetenskap

CTH TDA297
GU DIT290

Distribuerade system fk

Tentamen 2011-03-14

Dag, Tid, Sal: March 14th 2011, 14:00-18:00, V

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*

   Two processes, A and B, both share a concurrent stack. Consider the following history that your system can give to concurrent accesses to the stack objects by the two processes A and B:

   A invokes a push(5), A successfully finishes the push(5) operation, B invokes a push(6) operation, B successfully finishes the push(6) opeartion, A invokes a pop operation, A successfully finishes the pop operation that returns 5.

   (a) Is the execution linearizable? Is the execution sequential consistent? Please provide an explanation for that.

   (b) Give the definitions of linearizability and sequential consistency. Explain the difference between linearizability and sequential consistency.

2. *10 marks*

   Assume that your system provides a reliable, totally ordered (atomic) multicast. Design a replication scheme based on the state machine approach that quarantees that all replicas go exactly through the same sate transitions. Implement a highly availiable concurrent Queue on top of this replication scheme.

3. *15 marks*

   (a) Which type of faults are called Byzantine faults?

   (b) Prove that it is impossible to reach agreement in a system with three processes if one of them is Byzantine faulty.

   (c) How can the above proof be generalised for a system with n processes?

   (d) Is it possible to reach agreement in a system with three processes if one of them is Byzantine faulty by using authentication (unforgeable signatures)? If your answer is yes, describe an algorithm. If your answer is no give a proof to support it.

4. *5 marks*

   Define the specification of the i) Reliable Broadcast, ii) FIFO Broadcast and iii) Causal Broadcast.

5. *10 marks*

   Describe an algorithm that computes a spanning tree of a network $G(V, E)$. How nodes of the network can use the existence of such a spanning tree in order to perform routing? Is this a good solution to be used on wireless sensor networks with battery constraints?

6. *10 marks*

   For solving the dinning philosophers problem the following solution has been proposed: All philosophers except from philoshopher P_0 seek their right fork first. Philoshopher P_0 seek her left fork first.

Does this solution solve the dinning philosophers problem? If your answer is yes provide a proof and an analysis of the time complexity of the algorithm.

If your answer is no provide a counterexample and a fix to the above solution that makes it work. Please also provide an analysis of the time complexity of your solution.

Chalmers Techniska Högskola
Datavetenskap

CTH TDA297
GU DIT290

Distribuerade system fk

Tentamen 2012-03-05

Dag, Tid, Sal: March 5th 2010, 14:00-18:00, V 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*

   Give an example of an execution that is sequential consistent but not linearizable.

   Show that every linearizable history is sequential consistent.

2. *10 marks*

   Give a specification of a state machine replication abstraction and an underlying algorithm to implement it using a total order broadcast abstraction.

3. *5 marks*

   Is it possible to implement a total order broadcast deterministically in asynchronous systems with message loses?

   Provide an algorithm or an impossibility proof.

4. *10 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."

5. *10 marks*

   The flooding algorithm is a a straight forward broadcasting algorithm. The initiator sends to all its neighbors a message of kind broadcast. When a process receives message broadcast for the first time, it sends to all other adjacent processes further broadcast messages.

   What is the Time and Message Complexity of the algorithm? Please provide a complexity analysis.

6. *10 marks*

   The algorithm by Choy and Singh uses two Doorways. The Asynchronous and the Synchronous Doorway.

   Describe them in pseudo-code and informally.

   Use both Doorways to construct a solution to the Dinning Philosophers problem for a system of 3 philosophers.

7. *5 marks*

   Show that Byzantine agreement can be reached for three generals, if the generals digitally sign their messages.