

Exam in Distributed Systems

Justin Pearson

25 Aug 13-18, 2008

Cover Sheet

This sheet should be handed in together with the exam.

Each problem must be solved on a separate sheet. Write your name on each sheet. Indicate below which questions you have answered.

Problem no.	Solution provided	Max	Your points
1		6	
2		8	
3		5	
4		5	
5		8	
6		7	
7		9	
8		6	
9		9	
10		5	
Total:		54/68	

Name :

Pers.no. :

Exam Rubric

A mark of 60% is required for a pass, a 4 and 5 are distributed evenly. 85% is required for a VG (similar scheme of a 3-4-5). Answer can be in English or Swedish.

Each full answer should be started on a separate sheet. Please write your name and personal number on each sheet.

Students taking from F4Sy that is students on the 3 point version of the course only do questions 1-6.

I will not be able to come to the exam. If you are unsure about how to answer a question. If in doubt, make some reasonable assumptions, state your assumptions and answer the exam.

Hjälpmedel:

Pen, pencil, ruler, rubber (eraser for people from the Antipodes or from across the pond) calculator.

General Comments

I like examples and pictures. If you give interesting examples and draw useful pictures to illustrate any points that make. If you do all this you are more likely to get full points for the question.

Always show your workings when ever you do a calculation. If you do not show your workings then even if the answer is correct you will get 0 points.

1. General Questions on Distributed Systems

- (a) Give an example of a distributed system. Explain why it is distributed and what advantage over a non-distributed implementation of the same system. **(4 points)**
- (b) Why is it difficult to synchronise things in distributed systems. **(1 points)**
- (c) Explain why it might be a problem to detect failure in a distributed system. **(1 points)**.

2. Communication and Naming

- (a) Explain, with examples, the client server model. **(2 points)**
- (b) What are Client and Server Stubs and how are they used in remote procedure calls? Explain in detail how a remote procedure call is executed. **(4 points)**
- (c) When calling a remote procedure or method where a reference parameter is passed, what problems have to be solved and how are they solved? **(2 points)**

3. You are to design a distributed database system for a online ticket booking system for an airline. There will be more than one copy of the database for reliability and fault tolerance. Explain some of the problems (in distributed systems) you would have solve during the development of this system. In particular what problems have to be solved the are special for the domain (an online ticket booking system). **(5 points)**


4. Clock Synchronisation

- (a) What is a leap second and why is it used? **(1 points)**
- (b) Explain the Berkley algorithm for synchronising clocks. **(2 points)**
- (c) Suppose that you have two clocks both with drift $\rho = 10^{-7} \text{seconds/seconds}$. Suppose that you want the clocks synchronised within 0.01 of a second. How often should the clocks be resynchronised to achieve this? **(2 points)** (Obs. Show your workings, no workings no points for a correct or incorrect answer)

5. Logical Clocks

- (a) Consider the following three processes p_1, p_2 and p_3 with the fol-

following pattern of communication:



Label the each event with a normal Lamport timestamp (1 points) and a vector timestamp (1 points)

- (b) What does it mean for two events to be concurrent and what is the relation of the Lamport timestamps of the two events. (2 points).
 - (c) Is it possible for two events to have the same Lamport (non vector) timestamp? If it possible give an example, if it is not possible argue why it not possible. (2 points).
 - (d) Why are vector timestamps better than scaler timestamps? (2 points)
6. Transactions.
- (a) With examples define what a transaction is. (2 points)
 - (b) Explain what problems can happen there is no concurrency control where multiple transactions are being executed at the same time. (3 points) (Obs. Give examples).
 - (c) Define what it means for an interleaving of two transactions to be serially equivalent? Note that you must define what it means for two transactions to be interleaved(2 points)
7. Two and Three Phase Commit
- Explain the two phase commit algorithm. You should give the algorithms that the server and the clients implement and explain what happens. (3 points)
 - Suppose the server crashes during the two-phase commit algorithm. Explain if it possible and how the clients can come to an agreement. (2 points)
 - Explain the three phase commit algorithm and how it is an improvement over the two-phase algorithm. (4 points)
8. Consistent Cuts.
- With examples define what a consistent cut is. (3 points)

- Describe the snapshot algorithm and explain how it records consistent cuts. (3 points)

Stop here if you are a TF student.

9. Shared Memory

- What is strict consistency and why is not possible to achieve in a distributed system? (3 points)
- Define causal-consistency. You must give examples. (2 points)
- Is the following data store causally consistent? Explain your answer. Note that depending on your explanation there is more than one possible correct answer.(2 points)

A	W(x)a	W(x)b	
B	R(x)a	W(x)c	
C		R(x)c	R(x)b

- Explain with examples what release consistency is. (2 points)

10. Peer to peer and anonymous routing

- Explain with a node in a Gnutella network finds the location of a file which is stored in another node. (2 points)
- Describe some problems with Gnutella networks. (1 points)
- Outline how onion routing works. (2 points)