Exam in Distributed Systems

Justin Pearson 1DT641 1TT835

2007-06-13

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		16	
2		8	
3		5	
4		8	
5		8	
6		7	
7		14	
	Total:	52/66	

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. Answer can be in English or Swedish.

Each full answer should be started on a separate sheet. Please write you 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.

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) Describe the difference between asynchronous and synchronous communication in a network. (**1 points)**

   (b) Give two examples of algorithms or protocols that are easier to implement when the message delay is bounded and delivery is reliable. (**4 points)**

   (c) Explain why it might be a problem to detect failure in a distributed system. (**1 points)**

   (d) Give two examples of where replication could be used in a distributed system, explain in each case why replication is a good thing. (**2 points)**

   (e) What is middleware in a distributed system and why is it used? (**2 points)**

   (f) Explain the different types of transparency that are desirable in a distributed system. (**2 points)**

   (g) Give a breif overview of the differences between Corba and Java RMI. (**4 points)**

2. Communication

   (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. Clock Synchronisation

   (a) What is a leap second and why is it used? (**1 points)**

   (b) Explain Cristian's method for synchronising clocks. (**2 points)**

   (c) Suppose that you have two clocks both with drift $\rho = 10^{-12} seconds/seconds$. 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)

4. Logical Clocks

   (a) Consider the following three processes p_1,p_2 and p_3 with the following pattern of communication:

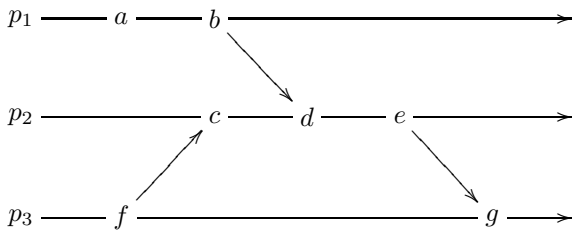

   Label the each event with a normal Lamport timestamp (**2 points)** and a vector timestamp (**2 points)**

(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)**.

5. You are to design a distributed system which allows users to book tickets for concerts. The systems must be robust, it must allow many users to book tickets online at the same time. Using some of the ideas that you have learnt in this course outline how you would design such a system. (**8 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? (**2 points)**

Stop here if you are a TF student

7. Shared Memory and Two Phase Commit

   (a) What is strict consistency and why is not possible to achieve in a distributed system? (**3 points)**

   (b) Define causal-consistency. You must give examples. (**2 points)**

   (c) Is the following data store causally consistent? Explain your answer. (**1 points)**

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

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

   (e) Describe in detail the two phase commit algorithm. (**6 points)**