Name______________ Lastname______________ Student number

Distributed Systems
academic year 2011-2012

Rules

- This is a closed books exam.
- The operation of any electronic device is prohibited (e.g. no calculator, phone or PDA).
- Answer the questions being *precise*, *complete*, and *formal*.
- Write as *clearly* as possible, both in terms of handwriting and wording.

Questions

1. Illustrate the five phases executed when performing a client request in a passive replication architecture, that is, with a primary back-up.
2. Consider the flooding algorithm in a graph (the simple version, without the counting of hops). The idea is that a source transmits a message, and each node forwards the message to all its neighbours, if it has not seen it yet. What is the time complexity of this algorithm in the synchronous and in the asynchronous model? What is the time complexity (in the worst case) of broadcasting in the resulting Spanning Tree in each case?

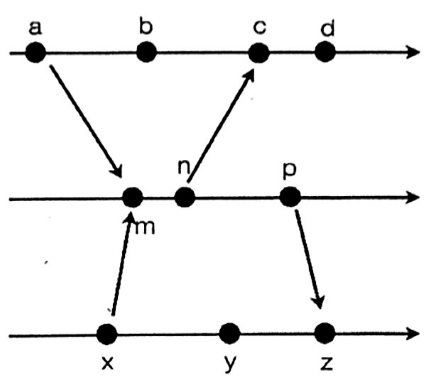

Figure 1: Events on three processes.

3. Consider three processes in a distributed system with the events depicted in Figure 1.
   (a) List all happened before relations among the 10 events.
   (b) List with which events is y parallel to, if any.
   (c) Provide the Vector clocks values for all the events.
   (d) Give a consistent snapshot of the system containing the event d.

4. Consider a synchronous system with reliable FIFO channels where the processors try to reach consensus on a binary value.
   (a) How many byzantine faulty processors can the system tolerate in general so that consensus is correctly reached for the non-faulty processors?
   (b) How many in the asynchronous case?
   (c) Consider the case of five processors $\{P_0 \ldots P_4\}$. Provide an example of dealing with one and with two bizantine failing processors and describe how the other processors may detect or not the faulty situation.
5. The Network Time Protocol can be utilized to synchronize the time on computers across a network. A NTP time server is utilized to obtain the correct time from a time source and adjust the local time in each participating computer. Consider the case of using an atomic clock for the NTP time server. The clock defines the second as 9,192,631,770 periods of the cesium-133 atom, and it diverts from real time of about one second every 15 million years. Now, two hosts $H1$ and $H2$ connect directly to the server reliably: the channel of $H1$ has a latency of 2 seconds (plus or minus 0.5 seconds), while the channel of $H2$ has a latency of 1.5 seconds (plus or minus 0.25 seconds).
   (a) What is the atomic clock resolution? What is its drift rate?
   (b) If $H1$ and $H2$ synchronize in the same instant, what will be their minimal and maximal skews?
   (c) What drift rates of the clocks of $H1, H2$ are tolerable, if one can accept at most a skew of 5 seconds one day after a synchronization?