

Examination Distributed Algorithms (IN4150)

5 april 2011, 9-12 AM

Notes:

1. The number of exercises is 4, and the number of pages is 2.
 2. The solutions to the exercises can be either in Dutch or in English.
 3. Try to give **short, concise, and precise answers**.
 4. The maximum number of points to be obtained for each part of each exercise is indicated between parentheses. The final grade is computed as 12 plus the total number of points obtained, divided by 10 and rounded to the nearest integer.
-
1. (a) (5) Give the definition of the *happened before* relation on events in a distributed system.
(b) (4) Give the definition of vector clocks. In particular, show how processes assign vector timestamps to events.
(c) (4) What is the meaning of the elements in the vector timestamp of an event in some process in relation to the events in the same and in the other processes?
(d) (5) How can causally ordered broadcast be implemented with vector clocks?
(e) (4) Show an example of how this algorithm indeed causes a message that arrives too early to be delivered at the correct moment.
 2. (a) (5) Explain why the minimum-weight spanning tree (MST) of a weighted graph with different weights is unique.
(b) (4) Give definitions of the concepts of *fragment*, the *level* of a fragment, and the *minimum-weight outgoing edge* of a fragment in the MST algorithm of Gallager, Humblet, and Spira (GHS).
(c) (5) Explain the concepts of *merging* two fragments and of *absorbing* one fragment into another in this algorithm.
(d) (4) Let G be a weighted graph that is a tree (with no particular structure) with any number of nodes. Show an assignment of weights to the edges of G and an execution of the GHS algorithm such that the final MST of G is of level 1. Explain your answer.
(e) (4) Let G be a complete binary tree with 7 nodes (i.e., all non-leaf nodes have two child nodes). What is the maximal value for the level of the MST of G ? Show an assignment of the weights 1, 2, ..., 7 to the edges of G such that the MST has this level.

3. (a) (4) Formulate the Byzantine agreement problem. In particular, state the conditions for *agreement* and *validity*.
- (b) (8) Give in words or in pseudocode the algorithm for randomized Byzantine agreement.
- (c) (4) Argue why this algorithm is also suitable for asynchronous systems.
- (d) (6) Given a system with 11 processes, two of which are faulty. Show an assignment of the initial values of 0 and 1 to the non-faulty processes and an execution of the algorithm of (b) with that assignment in which:
- all non-faulty processes will pick a random value at the end of round 1, and
 - all non-faulty processes will decide in round 2 of the algorithm.
4. (a) (6) Describe the structure and the contents of the routing cache (routing table) maintained by each of the nodes in the Freenet peer-to-peer system.
- (b) (6) Explain the algorithm of Freenet for inserting a new file.
- (c) (5) Argue why in Freenet the values of the file keys in the routing caches of single peers tend to become clustered.
- (d) (5) Give an explicit numerical example (with values in the routing caches) of an insert operation in Freenet that shows the clustering effect of (c).