Parallel and Distributed Systems Group
Department of Software Technology
Faculty of Electrical Engineering, Mathematics, and Computer Science
Delft University of Technology

Examination Distributed Algorithms (IN4150)

3 april 2009, 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) (4) Give the definition of causal message ordering of point-to-point messages.

   (b) (6) Give in words or in pseudocode the algorithm of Schiper-Eggli-Sandoz for causal message ordering of point-to-point messages.

   (c) (6) Consider the distributed system in the figure below with processes P_1, P_2, P_3, and P_4 and messages m_1, m_2, m_3, n_1, and n_2. Suppose that:

   - P_1 first sends n_1 and then m_1;
   - P_2 receives m_1 before first sending n_2 and then m_2;
   - P_3 receives m_2 before sending m_3;
   - P_4 first receives m_3, then n_2, and then n_1.

   Give the buffer and the timestamp accompanying each of the five messages when the Schiper-Eggli-Sandoz algorithm is employed.

   (d) (6) Show exactly how the conditions in P_4 for delivering or for postponing the delivery of the three messages it receives are or are not satisfied.

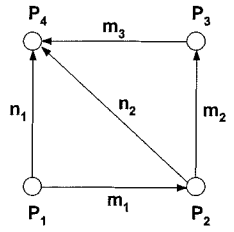

2. (a) (6) Give in words or in pseudo-code the *candidate process* of Afek's and Gafni's election algorithm in a *synchronous* complete network.

   (b) (4) Give in words or in pseudo-code the *ordinary process* of this algorithm.

   (c) (4) Argue that in the general case with n processes, the worst-case time complexity of the algorithm is $\log_2 n$.

   (d) (8) Assume that there are 16 processes and that the ids in the system are the numbers $0, 1, \ldots, 15$. Show an execution of the algorithm in which after 1, 2, 3, and 4 rounds, exactly 8, 4, 2, and 1 processes are still active candidate processes, respectively.

3. (a) (4) Formulate the Byzantine agreement problem. In particular, state the conditions for *agreement* and *validity*.
   Assume in the rest of this exercise that no authentication is used, and that n is the total number of generals and f the number of traitors.

   (b) (8) Give in words or in pseudo-code the algorithm OM(f) for Byzantine agreement.

   (c) (5) Assume that $n = 7$ and $f = 2$, and that the commander does not exhibit failures (is loyal). List in a *systematic way* all the messages a loyal lieutenant receives in every round of the algorithm, and show in detail how he deduces his final decision from these messages.

   (d) (5) The same question as in (c) when the commander does exhibit failures (is a traitor).

4. (a) (4) Explain the structure of the identities of nodes and files in the Pastry P2P system.

   (b) (8) Explain the structure of the routing table used by Pastry nodes for searching for a file.

   (c) (4) What is the size of this routing table (number of rows and columns) in terms of the parameters of a Pastry system?

   (d) (6) Explain how a new node that joins a Pastry system fills its routing table.