Massachusetts Institute of Technology | Handout 8
6.852: Distributed Algorithms
Prof. Nancy Lynch | September 22, 2005

Problem Set 2, Part a

Due: Thursday, October 6, 2005

Reading:

Chapter 6 of *Distributed Algorithms*
Aguilera, Toueg paper, listed in Handout 3
(Optional) Keidar, Rajsbaum paper, Sections 4.1-4.4 (skim)
(Optional) Attiya-Welch, Sections 5.1 and 5.2

Reading for next week: Chapter 7, skipping Section 7.3.4
Chapter 8

Problems:

1. Consider the following variant of the coordinated attack problem from Section 5.1 (the one that is supposed to tolerate message losses). Assume that the network is a complete graph of $n = 4$ nodes. The termination and validity requirements are the same as those in Section 5.1. However, the agreement requirement is weakened to say: "If any process decides 1, then at least three processes decide 1." Is this problem solvable or unsolvable? Prove.

2. Modify the *FloodSet* algorithm of Section 6.2.1 by adding a local *early decision* test condition, in order to obtain the following additional *early decision* time bound property:
   If the execution has only $f' \leq f$ failures, then all nonfaulty processes decide (but don't halt) by the end of round $f' + 2$.
   Prove that your algorithm works, that is, that it solves the stopping agreement problem for stopping failures, and that it has the additional time bound property.
   Note: We require that all processes that decide (even if they later fail), decide on the same value.

3. In class, we sketched the proofs of Theorems 6.31 and 6.32, which describe the $f = 1$ and $f = 2$ cases, respectively, of the lower bound of $f+1$ rounds for the agreement problem with stopping failures. Each of these proofs involved constructing a chain of executions spanning from one in which everyone starts with 0 and no one fails, to one in which everyone starts with 1 and no one fails.

   (a) How many executions appear in the chain constructed in Theorem 6.31?

   (b) How many executions appear in the chain constructed in Theorem 6.32:

   (c) If we extend this construction to arbitrary f, how many executions would appear in the resulting chain?

4. Exercise 6.20.

5. Consider a different kind of process failure model for synchronous systems: a "transient failure" model. In this model, a process may fail at a particular round, which means that it sends an arbitrary subset of the messages it is supposed to send (perhaps all of them), and does not perform its state transition. A process that exhibits a transient failure at a round r continues as if nothing is wrong at the following round $r + 1$. Permanent failure of a process at round r is modeled by transient failure at all rounds greater than or equal to r.

For this problem, we assume that, at each round, at most one process exhibits a transient failure. We do not assume an overall bound on the number of processes that ever exhibit a transient failure during an execution.

Now consider the agreement problem in this transient failure model: each process that does not fail permanently should eventually decide, subject to the usual agreement condition for stopping agreement, and the strong validity condition (every process' decision is some process' initial value).

Is this problem solvable in the given model? If so, describe an algorithm and sketch a proof that it works. If not, try to prove impossibility (carefully), using techniques like the ones in the Aguilera-Toueg paper.

6. Exercise 6.33.