

Exercise Session 4

Consensus

October 25, 2010

Problem 1

Consider all our fail-stop consensus algorithms (Consensus Algorithm I and Consensus Algorithm II). Explain why none of those algorithms would be correct if the failure detector turns out not to be perfect.

See Exercise 5.4, page 265 in “Introduction to Reliable and Secure Distributed Programming” book.

Problem 2

Explain why any fail-noisy consensus algorithm (one that uses a $\diamond P$ failure detector) actually solves uniform consensus.

See Exercise 5.5, page 265 in “Introduction to Reliable and Secure Distributed Programming” book.

Problem 3

Explain why any fail-noisy consensus algorithm (one that uses a $\diamond P$ failure detector) requires a majority of the correct processes. More precisely, provide a “bad run” in the case where there isn't a majority correct.

See Exercise 5.6, page 265 in “Introduction to Reliable and Secure Distributed Programming” book.

Problem 4

Improve Consensus Algorithm I (Algorithm 5.2 “Hierarchical Consensus”) to save one communication step. Consensus Algorithm I requires N communication steps for all correct processes to decide. By a slight modification, it can run in $N - 1$ steps: suggest such a modification.

See Exercise 5.2, page 265 in “Introduction to Reliable and Secure Distributed Programming” book.