Homework 3 (Failure Detection, Consensus, Peer-to-Peer, Routing, Distributed Object) - 100 Points

CS425/ECE 428 Distributed Systems, Fall 2009, Instructor: Klara Nahrstedt

Out: Thursday, October 15, **Due Date**: Thursday, October 29

Instructions: (1) Please, hand in hardcopy solutions that are typed (you may use your favorite word processor). We will not accept handwritten solutions. Figures and equations (if any) may be drawn by hand. (2) Please, start each problem on a fresh sheet and type your name at the top of each sheet. (3) Homework will be due at the **beginning of class** on the day of the deadline.

Problem 1: Failure Detection (20 Points)

Centralized heart-beating and ring heart-beating may not detect simultaneous multiple failures of processes, while all-to-all heart-beating is too expensive (in terms of messages sent per time unit). Suppose, in an **asynchronous system** that initially has ***N*** processes, you are given that at most ***N/4*** processes may crash.

(a) Design an efficient failure detector algorithm for this system. Your failure detector should satisfy completeness. Either give pseudo-code for your algorithm, or explain it using a figure.
(b) How many failure detector messages are sent by your algorithm if no failures occur?
(c) Calculate the best-case and worst-case detection times for your failure detector (hint: These are likely to occur with 1 and N/4 simultaneous failures respectively.)

Problem 2: Distributed Graph Algorithms – Routing (20 Points)

Consider a subnet that consists of routers A, B, C, D, and E (with links having symmetrical costs, where a lower cost is more desirable).

A	B	C	D	E
Age	Age	Age	Age	Age
Seq	Seq	Seq	Seq	Seq
B 5	A 5	B 8	B 2	A 3
E 3	C 8	D 8	C 8	C 8
	D 2	E 8	E 3	D 3

(a) Suppose the subnet uses **link state routing**. If router A receives the above link state advertisement (LSA) packets from all the other routers, **derive the network topology**. Note that an entire LSA packet contains a single age field that describes how long the information will be valid. Each entry in the LSA packet then specifies a neighbor and the link cost to the neighbor.

(b) Suppose the subnet uses **distance vector routing** instead, then consider the converged state, where all routes are optimal. Now, if the **link E − D goes down**, then list all the distance vector entries throughout the system that will change due to this link going down. For each changed entry, state the (old and new) next hop and the (old and new) cost of the entry.

Problem 3: Consensus (20 Points)

Explain briefly why the impossibility of consensus proof (proofs of Lemmas 2 and 3 in the FLP paper) would break if the system were synchronous. Specifically, give at least one statement in the proof that may not hold in a synchronous system.

Problem 4: Unstructured Peer-to-Peer Systems (20 Points)

Consider a Gnutella p2p system with 1023 total nodes that has become structured as a binary tree with height =9 (from a given root node.) What are the maximum and minimum number of nodes that can receive a Query message that is initiated with a TTL=7 from a leaf node of the tree?

Problem 5: Distributed Objects (20 Points)

Consider distributed objects in the distributed systems.

(a) Specify where (component(s)) and how (protocol) does the translation between local and remote procedure/object references happen when the distributed system is using remote procedure call (RPC)?

(b) Specify **where** (component(s)) and **how** (protocol) does the translation between local and remote object references happen when the distributed system is using remote method invocation (RMI)?