Problem Set 5, Part b

Due: Wednesday, May 3, 2006
Problem sets will be collected in class. Please hand in each problem on a separate page, with your name on it.

Reading

Location services Awerbuch, Peleg: Concurrent online tracking of mobile users
Li, Janotti, et al.: GRID location service paper
Abraham, Dolev, Malkhi: LLS paper
Clustering Mittal, Demirbas, Arora: Local clustering in large-scale wireless networks
Middleware services Malpani, Welch, Vaidya: Leader election for mobile ad-hoc networks
Walter, Welch, Vaidya: Mutual exclusion for ad-hoc mobile networks
Walter, Cao, Mohanty: k-exclusion paper

Reading for next week

Middleware services Malpani, Chen, Vaidya: Token circulation
Chen, Welch: Self-stabilizing dynamic mutual exclusion
Dolev, Schiller, Welch: Random walk for self-stabilizing group communication
Virtual objects Dolev, Gilbert, et al.: GeoQuorums: implementing atomic memory in ad hoc networks
Compulsory protocols Hatzis, Pentaris, et al.: Fundamental control algorithms in mobile networks
Chatzigiannakis, et al.: Efficiency of distributed communication and control algorithms
Chatzigiannakis, et al.: An efficient communication strategy for ad-hoc mobile networks
Chatzigiannakis, et al.: An efficient routing protocol for hierarchical ad-hoc mobile networks

Problems

1. Consider the algorithm in the Awerbuch-Peleg paper.
   (a) Draw an interesting network graph, and define a path that a mobile user ξ travels in the graph. Define the contents of the routing directories and the values of the forwarding pointers, for each node, after each MOVE step.
   (b) Discuss some issues involved in extending the Awerbuch-Peleg algorithm to the case where FIND and MOVE operations can proceed concurrently.
2. The behavior of the GRID location service in the Li, Janotti, et al. paper is analyzed for a static network.
   (a) What can you say about the behavior of the algorithm in dynamic cases, where nodes join and leave the network, fail and recover, and move?
   (b) Describe some ideas for improving the algorithm so it has better behavior in dynamic cases. What kinds of claims might you hope to prove about the behavior in such cases?
3. Describe (via pictures, or text), some interesting executions of the leader election algorithm of Malpani, Welch and Vaidya. In particular, illustrate the behavior in the following scenarios:

(a) A split of a component into two.

(b) A merge of two components that were previously separated.

(c) A link loss that does not cause a disconnection of components.

(d) An addition of a new node to the system (or, recovery of a previously failed node). This is not directly described in the given algorithm, but they allude to this extension in the Discussion section.

It's fine to use the first algorithm—assuming sequential changes— but of course you may tackle the concurrent one if you feel ambitious.

4. The k-exclusion algorithm of Walter, Cao and Mohanty is based on the mutual exclusion algorithm of Walter, Welch and Vaidya. List as many differences as you can between the two algorithms, and discuss why each difference was introduced.