# EECS 461 Homework #3
# Due Date: Friday, January 30, 3:30pm

All questions weighted equally. You are free to discuss this assignment with other students and the TA, but you must each individually write up and submit your own solutions.

1. Peterson and Davie, Chap. 6, Problem 2

2. Peterson and Davie, Chap. 6, Problem 3

3. Peterson and Davie, Chap. 6, Problem 10

4. Write a pair of RPC stubs to turn your code for homework #1 into a "priority queue server". In other words, the client calls "queue::Add()" and "queue::Remove()", but the actual code to implement the queue runs on a different machine. For extra credit, demonstrate that your code works by running it on two separate machines. What would happen if there were two clients sharing the same server (e.g., one client added an object to the queue and the other removed it)?

5. The book claims that the initial sequence numbers (ISN) of a TCP connection are chosen "randomly" by each end point. This is not quite true. Find out what TCP actually does to choose ISNs and briefly describe the algorithm. You might look this up in the TCP specification itself, as described in "Request for Comments 793" (RFC-793). To find an on-line copy, try your favorite search engine or go to http://www.ietf.org/, and follow links.

6. Consider the checksum algorithm used in the UDP and TCP protocols.

(a) If exactly one bit of the packet (including the header) is corrupted, does the checksum necessarily detect this error? If not give a counterexample of where it fails. If so explain why.

(b) If exactly two bits of the packet are corrupted, does the checksum necessarily detect this error? If not give a counterexample of where it fails. If so explain why.

(c) A "pseudoheader" consisting of the IP source and destination addresses, the packet length, and the IP protocol type is prepended to the UDP/TCP packet before computing the checksum. Why is this extra information included? (Give a brief justification for each field.)