HW #7

ELEN E4710 - Intro to Network Engineering　　Due 4/24/2003
Spring 2003　　Prof. Rubenstein

Homework must be turned in at the beginning of class on the due date indicated above. CVN students have one additional day. Late assignments will not be accepted.

1. A sender wishes to send its data to a set of n receivers, $r_1, r_2, \cdots, r_n$ as follows: The sender sends the data to r_1, and r_i sends the data to r_{i+1} for $1 \leq i < n$. r_n then sends acknowledgments to the sender. Packets are lost at each hop (including the acknoweldgment) via a Bernoulli process with probability p. Assume the sender never times out early (i.e., the sender does not retransmit a packet while a previous copy is still heading toward a receiver using the Selective-Repeat protocol or an acknowledgment for a previous transmission is on its way back to the sender).

   (a) Compute the expected number of times the sender transmits a particular packet (i.e., packet i for any i) for it to receive an acknowledgment for the packet.

   (b) Compute the expected number of times receiver r_i receives a copy of the packet.

   (c) If the window size used by the protocol is w, and it takes time τ for a packet to be received and acknowledged when no loss or corruptions take place, give an upper bound on the maximum rate of the protocol. Explain why your result is an upper bound (hint: *at most* w packets can be in transmission at any given time).

2. Suppose a sender and a receiver communicate over a lossy channel (data packets and ACKs can both be lost), where transmissions are not reordered. The size of the window is w, and packets are assigned sequence numbers modulo m (i.e., the ith packet is assigned the sequence number $i \pmod{m}$). Give examples where the protocol would fail to perform correctly if:

   (a) Go-Back-N is used and $m = w$.

   (b) Selective Repeat is used and $m = 2w - 1$.

3. A sender communicates with two receivers r_1 and r_2 using the go-back-N protocol with a window size of w. The sender can communicate with the receivers individually, so it can send packets to only r_1 or to only r_2 when desired. Packets sent from the sender to r_1 are never lost, and packets sent from the sender to r_2 are lost with probability p_1. Receiver r_2 can also get lost packets by recovering them individually (i.e., via selective repeat) from r_1. Transmissions from r_1 to r_2 are lost with probability $p_2 > p_1$, acknowledgments (from r_1 to the sender or from r_2 to anyone) are never lost.

   (a) Show that if $p_1 < 1 - w(1 - p_2)$ then the expected number of transmissions of packets to r_2 is minimized by never asking for transmissions from r_1, but always getting them from the sender directly via Go-Back-N. (Hint: consider the scenario where all packets are received in a window except the first. This is where using Go-Back-N is the most expensive.)

   (b) Consider now the case where there are m packets in the window that have not been received. What is the minimum value of m (as a function of p_1, p_2, and w) for which r_2 tries to recover packets from r_1 to minimize the expected number of transmissions? Solve this by figuring the expected number of packets received (that have not already been received) per packet sent, (i.e., $E[R]/E[S]$ where R is a r.v. that counts the number of non-duplicates received, and S is a r.v. that counts the number of packets sent).