HW #7

ELEN E4710 - Intro to Network Engineering — Due 11/26/2003
Fall 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 transmit a set of packets to n receivers. A protocol is used to execute this transmission by reliably delivering a set of k packets to all receivers before proceeding to the next set of k packets. The protocol proceeds over a series of rounds. Each round, the sender sends **all** k packets to **all** n receivers. If any receiver does not get all k packets, it sends a NAK. If, in the $i-1$st round, any receiver sends a NAK, the protocol proceeds to the ith round.

   Assume that each packet transmission is lost at each receiver inedependent from all other packet transmissions and with respect to all other receivers with probability p.

   (a) What is the probability that, after i rounds, a specific receiver j has not yet received a particular packet ℓ?

   (b) What is the probability that, after i rounds, all receivers have received all k packets?

   (c) What is the expected number of rounds needed to deliver all k packets to all n receivers (you may express in summand form)?

2. 2 versions of a reliable transfer protocol are designed, where in each version the sender is simultaneously transmitting W packets. For both protocols, when a packet is transmitted by the sender, the expected amount of time taken for the sender to learn the fate of that transmission (received or lost) is T. The sender waits for this information before trying to resend the packet (when the previous transmission was not received)) or send a new packet: this ensures that, from the sender's perspective, W packets are always in transmission.

   The packet loss process is Bernoulli with probability of loss, p. The time it takes for a sender to learn the fate of a particular packet is independent of the loss process. Determine the rate at which the receiver receives new packets (i.e., don't count duplicates twice) for the following two protocol versions:

   (a) A selective-repeat variant where there is no constraint on the size of the window, but the sender ensures that W packets are always in transmission. For instance, if $W = 5$, and the sender has been notified that packets $1, 2, 4, 7, 9, 12$ have been received, then packets $3, 5, 6, 8, 10$ are actively being transmitted.

   (Hint: how can you get the rate if you know (a) the expected time a packet spends in the system and (b) the expected number of packets in the system? Now how do you compute (a) and (b)?)

   (b) A Go-Back N where W consecutive packets are sent simultaneously, and the sender learns the fate of the transmission of all W packets simultaneously. (i.e., after all W packets have had enough time to reach the receiver, the receiver responds with the smallest sequence number i of a packet not yet received, and the sender immediately sends packets $i, i+1, \cdots, i+W-1$). The receiver drops any packets received out-of-order.

   (Hint: you may find it useful to think about the following system that would exhibit the same rate. When the sender sends W packets numbered $i, i+1, \cdots, i+w-1$, if all packets up to $i+k-1$ are received and the $i+k$th packet is dropped, then only k packets wind up being "accepted" by the receiver - it's as if only those k packets were injected that round. So what is the expected number of packets in the system now?)

3. 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) Go-Back-N is used and $m = w - 1$.

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

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

4. A sender wishes to transmit a series of packets to a pair of receivers, r_1 and r_2. An alternating bit protocol is used to perform the delivery where the sender communicates only with r_1, and r_1 communicates with both the sender and r_2. The sender and r_2 implement the normal sender and receiver versions of the Alternating Bit Protocol, where the sender sends its packets to r_1 and r_2 sends its ACKs to r_1. Three versions of r_1's protocol are presented below. For each version, indicate whether or not the protocol functions correctly (by correct, we mean that r_1 and r_2 always pass up packets to their respective applications in order, without passing up duplicates, regardless of which transmissions are lost). When the version is correct, explain why. When the version is faulty, demonstrate a faulty scenario.

   (a) Upon receiving a packet with sequence number i from the sender, r_1 passes the packet to the application layer whenever the previous packet passed up to the application layer had sequence number $\hat{i}$. r_1 forwards every packet it receives from the sender to r_2 (even those it does not pass up to the application layer), and forwards every ACK received from r_2 to the sender.

   (b) r_1 passes packets up to the application layer as described above, and only forwards to r_2 these packets that it has passes up to its application layer As above, r_1 passes all ACKs from r_2 to the sender.

   (c) r_1 passes packets up to the application layer as described in the above two versions, and forwards all packets from the sender to r_2, but only forwards ACKs for sequence number i from r_2 to the sender when the last packet r_1 received was a packet with sequence number i.