Exercise 4 — Retransmission Protocols + LANs

Due: Friday, Mar. 6, 1998

This assignment is a combination of problems about retransmission protocols used at the data link layer and medium access control protocols (used primarily in Local Area Networks).

To prepare for the discussion of the medium access control problem you should go back to chapter 3 in Shay and read §3.4 and then move ahead to chapter 6 and read §6.1 through 6.2.

In addition, I will put two papers I would like you to read on reserve in the library. I have been able to find both of these papers online in various electronic forms, so those who prefer can read them using an appropriate viewer.

The first paper, *Ethernet: Distributed Packet Switching for Local Computer Networks*, is the report by the inventors of the Ethernet on the "experimental" Ethernet built at Xerox PARC (Palo Alto Research Center). The second paper, *Measured Capacity of an Ethernet*, discusses an experimental evaluation of the performance of an Ethernet and reviews analytical work on Ethernet performance. Complete references to the papers can be found in the "Topics and readings" sub-section of the course home page.

1. In class, I mentioned that a few members of the class had brought it to my attention that there were some problems with the way negative acknowledgments or "naks" were handled by the pseudo-code for a Go Back N implementation provided in the text. That got me thinking about how to handle naks correctly in Go Back N. Unfortunately, I can't say I have all the details figured out, but I did at least come up with the following problem.

   One first step the sender might be expected to take upon receiving a nak from the receiver is to determine which packet caused the receiver to send the nak. While this may seem obvious, actually, it is not immediately clear that it is even possible.

   The problem arises from the fact that it is typical when implementing Go Back N to take advantage of the fact that if the sender's window size is W, then sequence numbers can be computed modulo W+1. In class, we saw that while using such a range of sequence numbers was enough to enable the receiver to distinguish the packet it should accept next from all others, it was not enough to enable the receiver to know whether a packet it rejects is a repeat copies of a previously accepted packet or an out-of-order first arrival. Luckily, the receiver didn't need to make such distinctions.

   When a receiver sends a nak, it can only identify the packet being rejected using its sequence number modulo W+1. In addition, it will include the sequence number of the packet it hopes to receive next in the nak. When the nak reaches the sender, it may be very desirable for the sender to be able to tell whether the packet was rejected because it was a repeat copy of a previously accepted packet or an out-of-order first arrival. In the former case, the sender will not need to retransmit the packet that caused the nak. In the later case, retransmission of the packet will be required.

   What I would like you to determine is whether or not the sender can distinguish between these two kinds of naks. If you believe they can be distinguished, give an argument similar to those presented in class explaining how the information available to the sender about the state of the system makes it possible to work with sequence numbers modulo W+1. If you believe they can not in general be distinguished, describe two example sequences of events that would appear identical from the sender's perspective such that in one sequence the final event is the receipt of a nak for a duplicate transmission and the second ends with the receipt of a nak for an out of sequence transmission.

You should assume that the receiver will transmit one nak for each packet that it can not accept based on its sequence number (i.e. naks are not sent for damaged packets).

2. Assume that node A is sending packets to node B using GO-BACK-N over a link that fails to deliver packets with probability p_1 and fails to deliver acks with probability p_2. Also assume that node A has sufficient buffer space to use a window of packets whose total transmission time exceeds the round trip propagation time on the link (i.e. $W * TR_F > TR_F + TR_A + 2 * P$ using the alphabet soup introduced in class). Suppose that A has transmitted packets with sequence numbers 1 through N where N is small enough that A would not yet expect to have received ACKs for any of the packets sent (i.e. $N * TR_F < TR_R + TR_A + 2 * P$). At this point, the algorithm described in the text would just send packet $N + 1$. It is not clear that this is the best choice. Given the flexibility we gave the sender in our specification of the Go Back N protocol, it could instead repeat some earlier packet. Obviously (I hope), A's optimal strategy would be to send the packet that has the highest probability of being received and accepted by B. So,

   (a) Determine the probability that packet number $N + 1$ will be received and accepted by B if transmitted at this point (By "received and accepted" I mean that it gets there intact and that B doesn't discard it as a packet out of sequence.).

   (b) Determine the probability that a repeated copy of packet number M (where $1 <= M <= N$) will be received and accepted by B if transmitted at this point.

   (c) Either argue that the value of the probability you determined for part a will always exceed that for part b (in which case A should never repeat packets before receiving a NAK or timing out) or give as simple a formula as you can expressing the relationship that must exist between M and N to justify sending packet M rather than $N + 1$.

3. Do problem 16 at the end of Chapter 3 in Shay.

4. I give up. There will only be three problems this week.