Midterm Examination

April 2-6, 2001

Name:

Started at: **Completed at:**

This is an open-book examination. You may consult the text, your notes, or any other inanimate source of information while completing this examination. There are four questions. You should stop all work on the exam at most two and one half hours after you begin.

1. (a) Calculate the CRC error-detection bits when G=$x^3 + x^2 + 1$ and M=10010110011.

   (b) Assuming that one is only willing to use two bits for error detection based on the CRC, there are only four possible generator polynomials: x^2, $x^2 + 1$, $x^2 + x$, and $x^2 + x + 1$. Which of these would be the best to use? Justify your answer.

2. Suppose one plans to send packets from some point A to some point Z using the stop-and-wait protocol to ensure reliable delivery. If A and Z are close enough together, you would just use a direct transmission line from A to Z. If A and Z are sufficiently far apart, you might instead consider introducing an intermediate repeater. In this problem, I would like you to analyze the performance of these two approaches.

   To make this interesting, I want you to assume that the repeater is fairly sophisticated. In particular, I want you to assume that the repeater itself uses the stop-and-wait protocol to talk to A and Z.

   In your answer, you should use the following identifiers to refer to the parameters that determine the behavior of the system.

   D - Time required to transmit a data frame.

   R - Time required to transmit an acknowledgment.

   P - Time required for a signal to propagate from A to Z. You should assume that the time for a signal to propagate from A or Z to the repeater is $P/2$.

   L_1^D - Probability that a data frame is lost on a single, long line from A to Z.

   L_2^D - Probability that a data frame is lost on one of the "short" lines from A or Z to the repeater.

   L_1^R - Probability that an acknowledgment is lost on a single, long line from A to Z.

   L_2^R - Probability that an acknowledgment is lost on one of the "short" lines from A or Z to the repeater.

   L_1 - Probability that a data frame or its acknowledgment is lost on a single, long line from A to Z. Note that $L_1 = 1 - (1 - L_1^D)(1 - L_1^R)$.

   L_2 - Probability that a data frame or its acknowledgment is lost on one of the "short" lines from A or Z to the repeater. Note that $L_2 = 1 - (1 - L_2^D)(1 - L_2^R)$.

   You should assume that $L_1 > L_2$. You may also assume that the processing time at the nodes and the repeater are negligible. Unlike the homework problem on stop-and-wait, you need only consider data flow in one direction when answering this question. Finally, assume that the timeouts used are as tight as possible.

(a) Assuming that there are no transmission errors, how long would it take for a packet to get from A to Z using each of these two methods.

(b) Assuming that there are no transmission errors and that A has a steady stream of packets to send to Z, what would be the utilization of the link from A to Z with and without the repeater.

(c) Now, assuming that L_1 and L_2 are greater than 0, give a formula for the utilization of the link using each of the two approaches. Can you say which approach results in higher utilization?

3. We have seen that Ethernet and ring networks naturally support the broadcast of a single message to all machines attached to a network. This is not true for point-to-point networks. Our discussion of point-to-point networks emphasized routing schemes, such as shortest path first, designed to transfer a message from a single source in a point-to-point network to a single destination.

For this question, I would like you to compare two approaches to broadcasting messages in a point-to-point network.

- The first approach is for the source node to simply send a copy of its messages to every other machine in the network using the standard routing scheme. That is, if there were 20 machines in the network, a machine that wanted to send a broadcast message would actually send 19 messages which would each be separately routed to the remaining machines in the network.
- The second approach is to use a flooding scheme like that usually used to distribute link-state information to the routers in a point-to-point network.

In her paper on the broadcast of routing information, Perlman provides the following brief description of the use of flooding to broadcast link state packets:

> "Flooding" means that a node broadcasts a Link State Packet on all its links except for the one on which it was received. "Intelligent flooding" means that the node recognizes duplicates, and does not flood a packet unless it is a new packet.
>
> With intelligent flooding, assuming no need for retransmissions, each packet will traverse each link at most twice, once in each direction."

Consider a network with the topology shown in the figure below and assume that node A wants to broadcast a message. Further assume that every message transmitted is delivered intact but that the transmission time on each link is unpredictable.

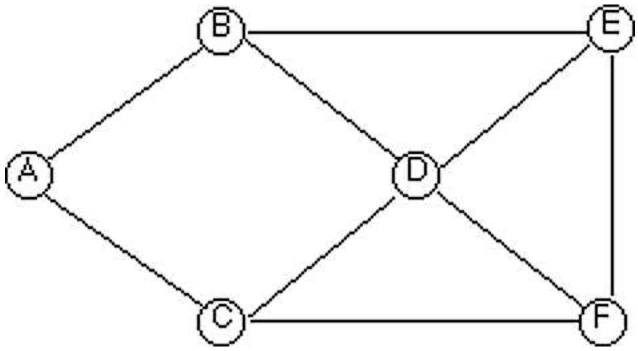

(a) Assume that A's message is broadcast using intelligent flooding. Briefly describe the sequence and timing of transmissions and receptions that would yield the smallest number of packet transmissions sufficient to deliver a broadcast message from A in this case. The best way to do this is probably to list the order in which nodes receive their first copies of the packet being broadcast indicating to which nodes each such node will forward the packet. So, an answer might look like:

A sends to B and C,
B receives message and forwards to D and E
E receives message and forwards to ...
C recieves message and forwards to ...

What number of transmissions occurs in this case?

(b) Still assuming intelligent flooding, what would be the largest number of packets that could be transmitted in the process of broadcasting a single message from A? Again, briefly describe the sequence and timing of transmissions that would yield this number of transmissions.

(c) How many packets would be sent if a message were broadcast from A by sending messages directly to each machine other than A using standard point-to-point routing for each copy?

(d) Describe a network topology that would ensure that the number of packets transmitted when flooding is used is always minimal (i.e. bounded by the number of nodes in the network). How many transmissions would occur if a message were broadcast by sending multiple point-to-point messages through a network of N nodes with this topology.

4. Suppose that A, B, C and D are computers attached to an Ethernet and that A, B, and C all become ready to send data while D is already transmitting. Describe a scenario in which A, B, and C contend for the use of the network after D's transmission is complete such that:

- initial transmission attempts occur in the order A, B, and C, but
- successful transmissions are in the order C, B, and A, and
- there are at least 4 collisions.

Your description should include a diagram showing the relative physical locations of A, B, C and D and a time-line showing the sequence of transmissions, attempts, collisions and backoff choices (indicate the range of possible backoff choices each time a choice is made).