CS 164 Fall 1995 — Final Examination
3 Hours, Closed Book
Aids allowed: Calculator, one sheet of notes

1. One of the jobs of a data link protocol is to recover from the loss of a data frame. A data link protocol is *only* used between two nodes that are *directly* connected by a physical link. Although bit errors may occur, for every bit sent from one end, a bit arrives at the other end. How could an entire frame get lost on the link?
2. When virtual circuits are used internally by the communications subnet, instead of datagrams, a shorter header can be used on each packet. Explain how this is possible.
3. When centralized adaptive routing is used in the subnet, what problem could arise if the routing updates were distributed to each node along the *old* routes?
4. A dialup modem that transmits 28,800 bits/sec over a voice grade line that can carry only a limited range of frequencies (roughly 0−3000Hz) obviously cannot be transmitting a separate symbol for each bit. Briefly describe two techniques that modems use to squeeze many bits of information into each transmitted symbol.
5. In Go-Back-N with window size 8, consider the following sequence of events in space-time between the client and server:

   a. What kind of connection is used to make this scenario happen? Connectionless or connection oriented? Why?

   b. Using the protocol implemented in our project 5 and the final project, and assuming that the line is noiseless, can this scenario cause the protocol to fail? Give your answer in a way that explains how the protocol survived or failed.

   c. Calculate the maximum throughput (in bits/sec.) that can be achieved by this protocol over a noiseless channel, assuming a data rate of 10Mbit/sec, frame size of 10,000 bits, and a round-trip propagation delay of 10 msec. You may ignore the processing delays at the two ends, and assume that the length of an acknowledgement frame is negligible.
6. In the unfaithful husbands problem, the wives use a counting algorithm to determine whether or not their own husbands are unfaithful. The basis of the algorithm was that each wife knows the status (faithful/unfaithful) of all husbands *except her own.* Therefore, since all of the unfaithful husbands are known to every wife (except, possibly, for her own husband) the

total number of unfaithful husbands known to any given wife can differ from the corresponding totals for every other other wife by at most one. Thus, a given wife knows that her own husband is faithful if she can determine (through the counting algorithm) that at least one other wife has a smaller count. Otherwise, she knows that her husband is unfaithful, since everyone knows that there is at least one unfaithful husband.

a. During the times of Queen Henrietta I and Queen Henrietta II, the wives used the same linear time counting algorithm. Queen Henrietta I called the wives together into a public meeting to tell them to get rid of the unfaithful husbands, whereas Queen Henrietta II informed them by letter through a reliable mail system. Briefly explain why the algorithm worked in Queen Henrietta I's time but not in Queen Henrietta II's time.

b. During the time of Queen Margaret, the wives were permitted to fire shots into the air, in addition to shooting their husbands as soon as they were found to be unfaithful. This allowed the wives to develop a new algorithm that worked in constant time. Briefly explain how the new algorithm works and state how many days are required.
Do not try to prove that it is correct.

7. In a concentrator, data from several input lines are sent on a single output line. Since the line speed (in bits/sec.) of the output line is less than the sum of the line speeds for all of the input lines, the utilization of the input lines must be low enough so that, *on average,* the output line can carry all the data. However, the *peak* input rate could exceed the output rate, so the concentrator needs to buffer its input. Assume that the arrival of frames from all input lines can be modelled as a Poisson process at rate λ frames/sec., and that the transmission time for frame on the (single) output line is exponentially distributed with a mean of $1/\mu$ seconds.

a. If the concentrator had infinite space for storing frames, then its buffer could be represented exactly by an M/M/1 queue. Find the average queue length, assuming $\lambda = 30$ frames/sec., an average frame length of 1800 bits, and a single 64,000 bit/sec ISDN channel as its outgoing line.

b. What happens to the *combined* buffer requirements if we replace 24 such concentrators with a single new concentrator that handles 24 times as much input traffic and uses the full capacity of a T1 channel (== 24 ISDN channels) as its outgoing line?

c. Now assume that the concentrator memory is limited, so that no more than B frames can be stored in the buffer at any time. Now it is possible for input frames to be *blocked* because they arrive when the buffer is already full. One common way to approximate this blocking probability is to use the usual infinite storage M/M/1 model to estimate the distribution of the number of frames in the queue, and then simply say that a new frame will be blocked if it arrives when the number of frames already present in the (infinite) queue is B or more. Write a formula for the blocking probability based on this method.

d. The exact value of the blocking probability for this system can be found from the solution to an M/M/1 queue with *finite* storage (i.e., no more than B in the system). Draw the state diagram for this case.

e. Find p_k, the probability that there are k in the system, in terms of ρ and B.
[HINT: $\sum_{j=0}^{N} x^j = (1 - x^{N+1})/(1 - x)$.]

f. What is the PASTA principle? Briefly explain how we can use it to conclude that the blocking probability in the method in parts (c) and (d) is simply p_B.

g. Assuming $\rho = 0.5$ and $B = 2$, estimate the blocking probability using both methods. Explain why the first method will always give a larger value.

END OF EXAMINATION