Final Study Guide

The class midterm examination will be held in DCL 1320 from 8:00 to 11:00 a.m. on Wednesday May 9[th]. The exam will begin and end promptly. Please arrive before 8:00 to allow everyone to settle into their seats before the test begins. No extensions will be granted to those who are late, nor will any non-emergency excuse for absence be accepted.

You may not consult any materials during the exam: no textbooks, no crib sheets, no calculator, etc.

The final will contain five to six parameterized problems and a set of short-answer questions. About half the total points on the exam will be for each type of problem. Each parameterized problem will consist of multiple parts, and all parameterized problems will carry approximately equal weight overall, but may break down unevenly amongst the parts. The short-answer questions require you to explain or comment on a topic relevant to the course in twenty-five words or less.

The problems and questions on the final will all be variants of some of the problems and questions that are either on this study guide or the problem sets.

On the final, you must show all work and reasoning, writing both work and solution legibly, and should box all answers. If the course staff cannot read a solution, no credit will be given. All short-answer questions must be stated in twenty-five words or less; longer answers will be graded by looking at only the first 25 words of the answer. Be concise, but do not spend your time counting words.

Parameterized Problems

1. Workstations as Switches

You are entrusted with the purchase of a workstation to serve as a switch between two high-speed local area networks (LAN's). One of the networks is an UltraRING, a 1 Gbps token ring LAN with 80 bytes of total overhead (headers and trailers) required for each frame. The other network is an OmniNet, a 1.3 Gbps Ethernet-like LAN with 100 bytes of overhead required for each frame. All packets sent on either network have exactly 1,000 bytes of data. After some research, you narrow the options to the two architectures described in the table below.

Name	Admiral J9000	SPQ
CPU handles	100,000 packets/second	60,000 packets/second
Number of I/O buses	3	1
I/O bus bandwidth	480 Mbps	1 Gbps
Memory bus bandwidth	2 Gbps	1.4 Gbps
Price	$10,000	$8,000

    (a) Pick one. Justify your decision, showing all work.
    (b) Draw a block diagram of the workstation architecture that you have chosen in part (a), labeling all components with appropriate names and data rates.
    (c) At the maximum sustainable bandwidth (i.e., with no packets dropped), what is the transmission rate - the total number of bits per second, including headers and trailers - sent over each network link?

2. Medium Access Control

This question concerns medium access control on a microwave network using carrier sense multiple access with collision detection (CSMA/CD, the algorithm used with Ethernet). The network consists of four hosts distributed as shown in the figure below. The microwaves are broadcast, and the signal travels directly along a line of sight from sender to all receivers. Assume that the signals propagate at the speed of light in a vacuum, 3×10^8 m/sec.

a. If a transmitter sends at 1 Mbps, how long must packets be to guarantee collision detection by the transmitter?
b. Divide time into slots the length of the maximum round-trip propagation delay in the network. One packet may be transmitted each time slot. Assume that each of the four hosts attempts to transmit with probability p in each time slot. What is the probability of a successful transmission in any given slot if
    a. $p = \frac{1}{4}$?
    b. $p = \frac{1}{2}$?
    c. $p = \frac{3}{4}$?
c. Using the minimum transmission length from part (a) and the probability of successful transmission from part (b)(ii) (for $p = \frac{1}{2}$), calculate the average throughput of the network if each packet requires 20 bytes of header/trailer and
    a. 10 bytes of data, and
    b. 50 bytes of data.

3. Bit- and Byte-Stuffing

Consider a data stream of 8-bit ASCII characters with values 0 to 255. Assume that the probability of a byte assuming each possible is equal (i.e., is exactly 1/255).

a. Using the bit-stuffing protocol discussed in class, what is the average number of bits that must be stuffed (inserted) per byte in the stream?
b. Answer the same question posed in part (a) for a byte-stuffing protocol in which the DLE character (value 16) must be escaped by stuffing a second DLE byte.

Now assume that the data stream contains only values in the range 32 to 127, again with a uniform probability distribution amongst the possible values.

c. Recalculate your answer to part (a) with the new probability distribution.
d. Recalculate your answer to part (b) with the new probability distribution.

4. Token Ring

Consider a token ring with latency of 500 μsec. Answer for both a single active host and for "many" active hosts. For the latter, assume that there are sufficiently many hosts transmitting that the time spent advancing the token can be ignored. Assume a packet size of 1500 B and the token is released after every packet.

    a. Assuming that the delayed token release strategy is used, what is the effective throughput rate that can be achieved if the ring has a bandwidth of 3 Mbps?

    b. Assuming that the immediate token release strategy is used, what is the effective throughput rate that can be achieved if the ring has a bandwidth of 3 Mbps?

    c. With delayed release, what is the effective throughout rate that can be achieved if the ring has a bandwidth of 100 Mbps?

    d. With immediate release, what is the effective throughout rate that can be achieved if the ring has a bandwidth of 100 Mbps?

5. Switch Fabrics

Banyan and Batcher networks are two types of self-routing fabrics often used to construct large switches from simpler components. A single stage of a $n \times n$ Banyan network consists of $n/2$ switches of dimension 2×2. An $n \times n$ Batcher network can be made from two Batcher networks of size $n/2 \times n/2$ plus a merge network with $n/2 \log_2 n$ switches.

    a. For n = 32, how many stages are required to route packets from the inputs to the outputs of a Banyan Network?

    b. How many 2×2 switches are required for the network in part a)?

    c. Write down a recurrence relation T(n) for the number of switches in a Batcher network of size $n \times n$.

    d. Give the number of switches required for n = 16.

6. IP Fragmentation

Consider two hosts, A and B, each on a separate shared Ethernet with MTU=1500 bytes. In addition to these LAN's, the route connecting host A to host B through the Internet contains an additional hop over a point-to-point link between a router on A's Ethernet and a second router on B's Ethernet. The point-to-point link has MTU=1000 bytes. Recall that MTU is the maximum amount of data that can be sent in a frame at the physical layer and thus includes all TCP and IP headers (each of which occupies 20 bytes). Also recall that IP fragmentation breaks data along 8 byte boundaries.

    a. An application on host A passes 2400 bytes of data to TCP. Following the approach used to draw Figure 4.4 (page 255 of P&D) but including the TCP header, sketch the packets that cross each link in the route. How many bits are delivered to the network layer protocol at host B?

    b. If the probability that any IP datagram crossing any link arrives intact (without error) is given by p, calculate the probability that the entire 2400 bytes sent in part (a) arrives without the need for retransmission.

    c. Calculate the average amount of data, including TCP and IP headers, and including all transmissions and retransmissions, that must be sent by host A in order to successfully deliver the 2400 bytes to an application on host B given $p = \frac{3}{4}$, where p is as defined in part (b). Assume that host B will buffer any data that arrives at the TCP level.

    d. Most IP datagram reassembly algorithms have a timer to avoid having a lost fragment tie up reassembly buffers forever. Suppose a datagram is fragmented into four fragments. The first three fragments arrive, but the last one is delayed. Eventually the timer goes off and the three fragments in the receiver's memory are discarded. A little later, the last fragment stumbles in. What happens to this last fragment at the receiver?

7. Comparison of Multiplexing Strategies

Consider a 15 Mb/s link that is shared by 4 flows. The mean data rate of the first flow is 6 Mbps, and the mean data rate of the other flows are 2 Mbps each. For ease of analysis we will make some assumptions so that the simple M/M/1 queueing model can be applied. Suppose that the packets of all flows are exponentially distributed in length with a mean length 12,000 bits, and that the packet arrival times for each of the flows form a Poisson arrival process. Suppose there is a transmit buffer at the front end of this link. Throughout this problem, the delay of a packet is taken to be the time from when the packet arrives to the buffer until the last bit of the packet is transmitted. (That is, "time in queue plus time in service" is counted, but not the propagation delay.) For each of the four scenarios below, find (i) the mean delay for flow one, (ii) the mean delay for flow two, three or four, and (iii) the mean delay, averaged overall packets of all flows. The last of these can be computed as

0.5×(delay for flow 1) + 0.166×(delay for flow 2) + 0.166×(delay for flow 3) + 0.166×(delay for flow 4)

or, by using Little's law, as

(Mean amount of data queued or in transmission)/(sum of arrival rates).

(Hint: Think of the packets as customers to a queueing system. Express mean arrival rates and service rates in packets/second, rather than in bits/second. If an arrival rate exceeds a service rate, take the mean delay to be infinite.)

Scenario A: A Time Division Multiplexing (TDM) scheme is used with a small frame size and *equal allocation of transmission rates*. Thus, each flow effectively sees the same service it would from a 3.75 Mbps link that is not shared with other flows.

Scenario B: Same as Scenario A, but instead of equal allocations, the allocations are proportional to the arrival rates. Thus, flow one is allocated 7.5 Mbps, while the other three flows are allocated 2.5 Mbps each.

Scenario C: Same as Scenario A, but find and use the allocation of rates for the four flows which minimizes the mean delay, averaged over all flows.

Scenario D: Statistical multiplexing is used in which the link serves packets in first come first served order (Note: the combined flows again form a Poisson flow, and all the substreams have the same mean delay.)

8. Token Bucket Filtering

Consider a traffic stream with rate indicated in the figure below. Note for example that 100 bits were sent between time 10 and time 20.


Suppose this stream is passed into a token bucket filter with token generation rate r (where one token is needed *per bit)* and the token buffer size is *B*.

    a.   Suppose *r=10* kbps (i.e., 10 kilo *tokens* per second). What is the minimum size of *B* required so that the filter lets the stream pass with no loss or delay? (Hint: As shown in class, this is the same as the maximum queue size if the stream is fed into a queue served at constant rate *r.)*

    b.   Repeat for r = *5kbps.*

    c.   Find the minimum *B* needed for arbitrary r > 0 and sketch your answer.

9. Distance-Vector Routing

Consider the following network configuration where the routers calculate shortest routes using the Distance Vector Routing Protocol.

Initially, the router tables for routes to node A look like the following:

B		C		D		E	
Cost	Next Hop	Cost	Next Hop	Cost	Next Hop	Cost	Next Hop
1	A	2	B	2	B	3	D

Now, assume node A goes down.

Given the following sequence of routing update messages, fill in the table for the routing entries for reaching A at each event, where the notation B → C indicates that node B sent a routing update to node C.

	B		C		D		E	
Event	Cost	Next Hop	Cost	Next Hop	Cost	Next Hop	Cost	Next Hop
	1	A	2	B	2	B	3	D
Node A goes down								
C → B								
B → D								
D → E								
E → D								
D → B								
B → C								
C → B								
B → D								

Analyze this carefully and try to see what's happening compared to what should happen ideally. So what, according to you, is the problem here? Why is it a problem?
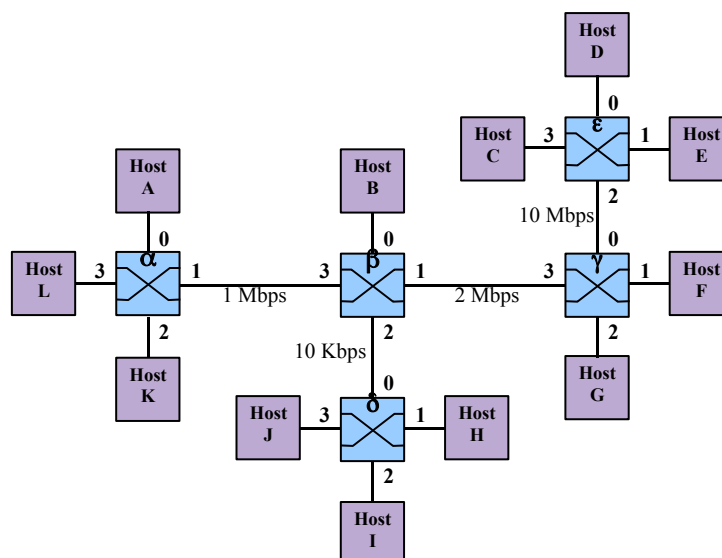
10. TCP Slow Start
Although slow start with congestion avoidance is an effective technique for coping with congestion, it can result in long recovery times in high-speed networks.
a. Assume a roundtrip time delay of 100 ms (about what might occur across a continent) and a link with an available bandwidth of 500 Mbps and a segment size of 576 octets. Determine the window size needed to keep the pipe full and the time it will take to reach that window size after a time out using the Jacobson Algorithm.
b. Repeat for a segment size of 16 Kbytes.

11. Max-Min Fairness

Consider the network shown in the figure. The links are labeled with capacity. This problem deals with circuit-switched forwarding and fairness.

Host D

Host C — 3 ε 1 — Host E

0 (Host D to ε)

10 Mbps

Host A

Host B

Host L — 3 α 1 ——— 3 β 1 ——— 3 γ 1 — Host F

0 (A to α)   1 Mbps   0 (B to β)   2 Mbps   0 (to γ)

2   2

Host K   10 Kbps

Host G

0

Host J — 3 δ 1 — Host H

2

Host I

a. Give the VCI assignment for circuits set up between the following hosts.

   (L → E), (A → I), (I → E), (F → K), (G → B), (C → F)

   Assume for a given link and a given direction, whenever a new virtual circuit is set up through the link, the smallest unused (on the given link and direction) virtual circuit identifier is selected.

Switch α				
IN		OUT		Circuit
PORT	VCI	PORT	VCI	
0	3	0	1	
0	0	1	1	
0	1	0	2	

Switch β				
IN		OUT		Circuit
PORT	VCI	PORT	VCI	
0	3	0	1	
1	3	0	2	
0	2	1	1	
0	1	0	3	
1	1	0	0	

Switch γ				
IN		OUT		Circuit
PORT	VCI	PORT	VCI	
0	3	0	0	
1	3	1	0	
0	1	0	3	
0	2	1	3	
0	0	0	1	

Switch δ				
IN		OUT		Circuit
PORT	VCI	PORT	VCI	
0	0	0	2	
0	2	0	0	

Switch ε				
IN		OUT		Circuit
PORT	VCI	PORT	VCI	
0	2	0	1	
1	2	1	1	
0	3	0	2	

b. Given the following bandwidth requests for the connections, determine a fair allocation of resources for each connection. For this question, use the definition of max-min fairness discussed in class.

(L → E) = 1 Mbps; (A → I) = 5 Kbps; (I → E) = 7 Kbps; (F → K) = 1 Mbps; (G → B) = 1 Mbps; (C → F) = 10 Mbps

12. Fair Queueing

A router implements fair queueing of four incoming flows over a single outgoing channel. In this problem, you will determine the order of packets sent out from the router and calculate statistics from your results. Assume for simplicity that the problem starts at time $t = 0$ and that sending a packet of length N requires N time units. Packets arrive on the four flows, named A, B, C, and D, as follows:

A packets of length 10 arrive at times 0, 5, 10, 15
B packets of length 8 arrive at 25, 45, 65, 85, 105
C packets of length 5 arrive at 10, 20, 30, 40, 50, 60, 70, 80
D packets of length 4 arrive at 1, 2, 30, 31, 32, 33, 80, 81, 82, 83

All service counters start at 0 at time 0.

a. For each packet sent on the outgoing link, record the start time, the service counter for each flow at that time, and the name of the packet *(e.g.,* write *B3* for flow B's third packet).
b. Calculate the percentage of the link used by each flow up to time 100. Was the link fairly distributed? Why or why not?
c. Calculate the average packet delay for each flow-the average difference between the arrival time of the packets and the time that they begin to be sent on the outgoing link. Explain the differences in the averages in terms of the burstiness of the arrivals.
d. Repeat parts (a), (b), and (c) using round-robin scheduling. Skip queues that are empty on their turn. Explain the differences in utilization and delay in comparison with fair queueing.

13. Sliding Window Algorithms

This question considers a sliding window implementation across a full-duplex point-to-point link. The link has a bandwidth of 327 kbps in each direction and a one-way propagation delay of 100 milliseconds. All packets sent across the link are 1,024 bytes long, including all headers and trailers.

a. How much data is required to fill the pipe for a round-trip delay on the network?
b. What send window size (SWS) is necessary to fully utilize the network?
c. For RWS= $\lfloor$ SWS/2 $\rfloor$, construct an example demonstrating that SWS+2 sequence numbers (e.g, from 0 to SWS+1, where SWS is your answer to part (b)) are not enough to guarantee correct operation of the sliding window algorithm.
d. Given a go-back-n algorithm (with RWS=1), assume that data frames are received with probability $p = 0.9$ and that acknowledgements (ACK's) are always received (probability of 1). Further assume that the retransmission timeout used by the sender is a negligible amount of time longer than the round-trip time, implying that a packet is retransmitted as soon as the sender could possibly detect its loss on the previous transmission. Calculate the average transmission rate (in bits per second) achieved for long streams of data.

14. Delay-Bandwidth Product for Links in Series

Consider three nodes in series. Node A is connected to node B via a 75 Mbps fiber optic link, 1200 km in length. Node B is connected to node C via a 2 Mbps link, 5 km in length. The links are full duplex. The rate of transmission errors on the links, the time to switch a packet at node B, and the time to transmit an ACK are all negligible. A large file is to be sent from node A to node C, and there is no other traffic on the links. Packets are 1 KB, including headers.

a. Ignoring reliability and packet headers, what is the maximum throughput that can be achieved (in Mbps)? Explain.
b. What is the round trip time from A to C?
c. What is the roundtrip bandwidth delay product for the path from A to C? (Specify the units you use).
d. Suppose an end-to-end sliding window protocol is used with SWS=RWS. What size of SWS is optimal?
e. Why wouldn't you want SWS to be many times larger than the value you suggested in part d?

15. TCP over a Raw Wireless Link

In this problem, you are basically asked to predict the throughput for a TCP connection with a wireless link in the path between the TCP source and TCP destination. Packets are 800 bytes. The wireless link corrupts packets, with no attempt at retransmission or forward error correction (hence the adjective "raw") with each bit being corrupted with probability 3×10^{-6}. TCP Reno is used. The TCP-Reno protocol enters slow-start only after a timer expiry, but for this problem ignore time-outs (which is a good assumption if the round trip times are not highly variable), so assume that the TCP connection is always in the congestion avoidance phase. Upon receiving three duplicate ACK's, the source goes back to resend the requested packet and also cuts the congestion window in half. Assume that the socket buffers and resulting advertised congestion windows are large enough so as not to impede throughput. Suppose the data rate of the wireless link is 10 megabits per second, and that the data rate on the wireline portion of the path is much greater than that. Suppose the round trip time plus packet transmission time plus ACK transmission time is 20 ms.
  a. Find the packet error probability.
  b. What is the round-trip bandwidth-delay product, expressed in packets? That is, under error-free operation without limitations by the congestion window, what is the typical number of unacknowledged packets?
  c. Ignoring the possible impeding effects of the congestion window (i.e. assuming it is always very large), what is the end-to-end throughput, expressed as a fraction of the throughput possible on an error-free link. A fraction p of the packet transmissions are in error. To simplify your analysis, assume that beginning with a retransmitted packet, the next faulty transmission occurs in the 1/p th packet (rather than randomly). Round 1=p to the nearest integer for this purpose.
  d. Find the end-to-end throughput, now incorporating the effect of the congestion window.

16. Poisson Processes

A communication line capable of transmitting at a rate of 150Kbits/sec will be used to accommodate 10 sessions. Each session generates Poisson traffic at a rate 250 packets/min. Packet lengths are exponentially distributed with a mean of 1000 bits.
  a. For each session, find the average number of packets in queue, the average number of packets in the system, the average delay per packet when the line is allocated to the sessions by using
      i. 10 equal-capacity time-division multiplexed channels.
      ii. Statistical multiplexing
  b. Repeat part (a) for the case where the five streams transmit at a rate of 350 packets/min while the other five transmit at a rate of 100 packets/min.
      i. time-division multiplexed channels
      ii. Statistical multiplexing

17. Queueing Theory

Consider a M/M/1 queue, which represents a network interface. Packets arrive, possibly wait in the queue before being serviced (transmitted), get transmitted and leave the system. Suppose that the Poisson arrival rate of packets for execution is 4 packets/sec and the average (exponentially distributed) service time is .20 sec.
  a. What is the service rate (in packets/sec)?
  b. What is the throughput (in packets/sec) for the given arrival rate and service time?
  c. What is the utilization?
  d. What is the probability that an arriving packet will have to wait before receiving service?
  e. What is N_{avg}, the average number of packets in the system?
  f. What is Q_{avg}, the average number of packets waiting?
  g. What is T_{avg}, the average amount of time a packet spends in the system?
  h. What is W_{avg}, the average amount of time a packet spends waiting to be serviced?
  i. Suppose that the service rate and arrival rate are both doubled, how do T_{avg} and N_{avg} change?

Short-Answer Questions

1. Explain the exposed terminal problem and how it is solved.
2. Name the OSI layer or layers in which medium access control (MAC) is addressed and state whether MAC is typically handled in hardware, in software, or in both in the Internet architecture.
3. Explain the effect of layering on end-to-end bandwidth.
4. Name the OSI layer or layers in which framing is addressed and state whether framing is typically handled in hardware, in software, or in both in the Internet architecture.
5. Show that the final parity check in a horizontal and vertical parity check code, if taken as the modulo 2 sum of all data bits, is equal to the modulo 2 sum of the horizontal parity checks and also equal to the modulo 2 sum of the vertical parity checks.
6. Describe the benefits of error correction over error detection.
7. Give two arguments against IP reassembly in routers.
8. List three components that contribute to end-to-end latency.
9. The sequence number field in the TCP header is 32 bits long, which is big enough to cover over 4 billion bytes of data. Even if this many bytes were never transferred over a single connection, why might the sequence number still wrap around from $2^{32} - 1$ to 0?
10. Explain the circumstances that give rise to the count-to-infinity problem.
11. Assuming SWS=3, RWS=1, and independent timeouts per packet, construct a minimal timeline such that timeouts for packets in the send window are neither monotonically increasing nor monotonically decreasing.
12. Name the OSI layer or layers in which error detection is addressed and state whether error detection is typically handled in hardware, in software, or in both in the Internet architecture.
13. For TCP, why does the maximum packet lifetime, T, have to be large enough to ensure that not only the packet, but also its acknowledgements, have vanished?
14. Caching is an important mechanism whereby frequently used information is replicated in order to provide fast access at different physical locations. Name three instances of caching discussed in the course that arise in the context of standard Internet operation.
15. At what OSI layer do Internet routers typically operate?
16. Why are differentiated services approaches to supporting QoS considered more scalable than integrated services approaches?
17. What Hamming distance is necessary for n-bit error detection? n-bit error correction?
18. If SWS=RWS=5 in a sliding window protocol, if packet numbers do not wrap around, if packets do not arrive out of order, and if the next frame expected (NFE) is currently 17, why can't the receiver next receive a packet with sequence number 10?
19. Explain how a receiver detects the end of a frame with sentinel-based framing.
20. Why will UDP require a checksum with IPv6?
21. Why might a host using an IEEE 802.11 access point (AP) continue scanning for other access points?
22. Explain the meaning of the error, "address already in use," sometimes returned from calls to bind.
23. Name the four components that uniquely specify a TCP connection and the length of each component in bits.
24. What does TCP use in addition to an estimate of RTT to calculate timeouts for adaptive retransmission?
25. Suppose a dynamic routing algorithm is employed to try to make routing tables correspond to least cost paths. What types of routing metrics are prone to producing load oscillations?
26. Explain in words (no equations) what the memoryless property of a random, exponentially distributed lifetime is.
27. What does UDP provide in addition to those services contained in IP?
28. What is the difference between congestion avoidance and congestion control?
29. How does TCP guarantee that new connections do not receive segments from previous incarnations of the connection?
30. Describe the responsibilities of hosts and routers using DECbit to avoid congestion.
31. Give an argument why the leaky bucket algorithm should allow just one packet per tick, independent of how large the packet is.
32. What is the purpose of the protocol field in the IPv4 header?
33. List five services demanded by many applications but not provided by IP (nor typically provided by user-level code).
34. Explain the fundamental conflict between tolerating burstiness and controlling network congestion.

35. Why does TCP begin by multiplicatively increasing its congestion window? What is "slow" about this approach?

36. Having ARP table entries time out after 10 -15 minutes is an attempt at a reasonable compromise. Describe the problems that can occur if the timeout value is too small or too large.

37. Give an example of scheduling discipline that is not work-conserving.

38. An approach to building special purpose hardware for massive high-speed switching fabrics is to use a Batcher sorting network followed by a self routing Banyan network. Why is the Batcher network included?

39. How does IP limit messages to 64 kB in the common case? Why does IPv6 provide for longer messages?

40. Why doesn't the adaptive time out mechanism of TCP update EstimatedRTT in case an ACK is received for a segment that was retransmitted?

41. If all the links in the Internet were to provide the reliable-delivery service, would the TCP reliable-delivery service be completely redundant? Why or why not?

42. With Go-Back-N, is it possible for the sender to receive an ACK for a packet that falls outside of its current window?

43. Why does TCP use a 32-bit sequence number space instead of calculating a tighter bound based on RTT and link speed? Assume that complexity is minimal and that saving two bytes of header space {for example) is worthwhile.

44. Explain what "fair" means for flows traversing a router.

45. Ethernet frames must be at least 64 bytes long to ensure that the transmitter is still going in the event of a collision at the far end of the cable. Fast Ethernet has the same 64 byte minimum frame size, but can get the bits out ten times faster. How is it possible to maintain the same minimum frame size?

46. Explain the relationship between physical distance and end-to-end latency in a TCP connection.

47. Under what circumstances does TCP Vegas increase its window size?

48. Recall that with IP tunneling, we said that an IP datagram is carried inside of another IP datagram. How does the IP router at the end of the multicast tunnel know that the outer datagram contains an inner IP datagram (as opposed to simply being a normal IP datagram that should be forwarded along)?

49. TCP waits until it has received three duplicate ACK before performing a fast retransmit. Why do you think the TCP designers chose not to perform a fast retransmit after the first duplicate ACK for a segment is received?

50. What does a host do when it receives an ARP from an unknown host to a second unknown host?

51. Why do internet routers stop at the IP layer rather than passing data up to TCP or UDP internally?

52. What traditional network class was under the most pressure before CIDR? Why?

53. Explain how fair queueing prevents flows from "saving up credit."

54. Suppose the throughput for a particular TCP connection is limited primarily due to the fact that one of the links it traverses is heavily congested. The congested link is shared by several TCP connections. How does the propagation delay for the TCP connection affect the throughput it receives?

55. In RED gateways, explain why MaxThreshold is actually less than the actual size of the available buffer pool.

56. Suppose packets on a wireless link consist of N data bits and H header bits each, where H is fixed. Suppose bits are received in error with probability P, independently of each other, and that N is adjusted to maximize the throughput of data in bits per second. If P gets larger, does the optimal value of N get larger or smaller? Why?

57. Data link protocols almost always put the CRC in a trailer, rather than in a header. Why?

58. What is the maximum bandwidth attainable on a TCP connection with RTT=100 milliseconds? Explain how TCP options can be used to raise this limit.

59. As we have seen many times in class, a sliding window abstraction can be used to bound transmission rates. Why would anyone propose a rate-based mechanism, given that buffer (window) space is intrinsically available from the end hosts?

60. Why is an ARP query sent within a broadcast frame? Why is an ARP response sent within a frame with a specific destination LAN address?

61. List the contents and explain the purpose of each segment transmitted when a TCP connection closes in a typical way.

62. Give two good reasons to allow branching – that is, the ability to support multiple protocols above and below any given protocol-in protocol graphs.

63. Intuitively, what is the goal of TCP Vegas? Why is this goal desirable?

64. Draw a diagram explaining the function of a self-healing ring topology.

65. Name an advantage of switched media over shared media. Name a disadvantage.

66. How does CIDR solve the problems of inefficient address allocation and limited number of networks associated with the traditional class model?

67. How does fast retransmission improve TCP's overall utilization of network resources?

68. Compare the problem solved by Nagle's algorithm to silly window syndrome and describe the similarities between the two problems.

69. IP hosts that are not designated routers are required to drop packets misaddressed to them, even if they would otherwise be able to forward them correctly. In the absence of this requirement, what would happen if a packet addressed to IP address A were inadvertently broadcast at the link layer?

70. Assuming that all routers and hosts are working properly and that all software in both is free of all errors, is there any chance, however small, that a packet will be delivered to the wrong destination?

71. The original Internet mechanism for looking up names used central hosts.txt table, which was distributed to all hosts every few days. Describe two reasons why this mechanism is no longer used.

72. Describe two advantages of using encapsulation (tunneling) for distributed internet applications such as virtual private networking.

73. A friend comes to you and asserts that network programming is too hard; complaining that after select indicates data available on a connection, read returns no data. Explain your friend's problem.

74. How does a RED gateway act to avoid congestion?

75. Under what circumstances may coarse-grained timeouts still occur in TCP even when the fast retransmit mechanism is being used?

76. List the contents and explain the purpose of each segment transmitted during a TCP connection setup.

77. Due to the use of CIDR, it is possible that the destination address on an incoming packet will match several entries in a routing table. In such a case, which routing entry or entries will be used for forwarding the packet?

78. A class B network on the Internet has a subnet mask of 255.255.240.0. What is the maximum number of hosts per subnet?

79. Most IP datagram reassembly algorithms have a timer to avoid having a lost fragment tie up reassembly buffers forever. Suppose a datagram is fragmented into four fragments. The first three fragments arrive, but the last one is delayed. Eventually the timer goes off and the three fragments in the receiver's memory are discarded. A little later, the last fragment stumbles in. What should be done with it?

80. Give a potential disadvantage when Nagle's algorithm is used on a badly congested network.

81. Explain Karn's algorithm. Why do we need it?

82. How can ARP be used to redirect traffic for a given host? Give pros and cons of such a technique.

83. Explain the hidden terminal problem and how it is solved.

84. Give the major reason why TCP does not behave well in wireless environments.

85. In queueing theory, explain why the arrival rate, λ, must be less than the service rate, μ.

86. Name and describe the type of multiplexing traditionally employed in data networks.

87. Define the Hamming distance of an encoding.

88. In the Perlman distributed spanning tree algorithm, why does the root bridge periodically send messages even after the tree is determined?

89. What is the difference between a cumulative and a selective acknowledgement?

90. In a sliding window protocol, explain why would you ever use an RWS that is not equal to the SWS.

91. Why does Ethernet used fixed time slots during backoff? What could go wrong if the fixed slots were not used?

92. Cell switching methods essentially always use virtual circuit routing rather than datagram routing. Give a specific argument why this is so.

93. Define the process of routing.

94. Describe the problem solved by sliding window based ARQ protocols.

95. Describe the problem solved by reliable transmission.

96. Describe two ways in which the topology of the Internet has evolved over the last ten years.

97. Draw a protocol graph for the Internet, including at least the following: ATM, Ethernet, FDDI, FTP, HTTP, IP, TCP, TFTP, and UDP.

98. Explain the main bottleneck for sending short and long messages.

99. Explain the main drawback of the stop-and-wait ARQ algorithm.

100. What delay is relevant and what bandwidth is relevant for computing the delay-bandwidth product of two links in series?

101. Why does Ethernet use binary exponential backoff during contention resolution?

102. Why does reverse path broadcast assign parent routers to each LAN?

103. Why does TCP use a 32-bit sequence number space instead of calculating a tighter bound based on RTT and link speed? Assume that complexity is minimal and that saving two bytes of header space {for example) is worthwhile.

104. With the selective repeat protocol, is it possible for the sender to receive an ACK for a packet that falls outside of its current window?

105. What is the role of the NAV in IEEE 802.11?

106. In Token Ring networks, why would we want to emulate a shared medium with point-to-point links?

107. How do sniffers work? Will they work on all networks?

108. Describe the pros and cons of using MTU discovery on a path prior to data transmission.

109. How can NAT be used to load balance a company's servers?

110. What is the role of the root servers in DNS?

111. How can an attacker poison a DNS cache and what is the impact?

112. What is the role of a BGP border router?

113. Why does BGP us path-vector routing?

114. Why would **connect()** be used with a UDP socket?

115. Why can 4B/5B encoding be transmitted using NRZI?

116. Why is it ineffective to use and ACK for broadcast and multicast communication in wireless networks?

117. Why does a node in 802.11 suspend is collision counter when the medium is busy?

118. Given an example of when the use of 802.11 can lead to unfairness?

119. Why is buffering needed in switches?

120. Why are metrics like link utilization and delay difficult to use effectively in routing?