Midterm Study Guide

The class midterm examination will be held in DCL 1320 from 7:00 to 9:00 p.m. on Wednesday March 7th. The exam will begin and end promptly. Please arrive before 7:00 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 after Tuesday February 29th.

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

The midterm will contain three to four 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 midterm will all be variants of some of the problems and questions that are either on this study guide or were in problem sets 1 and 2.

On the midterm, 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. Channel Rates and Shared Media

You are entrusted with the design of a network to interconnect a set of geographically distributed hosts within your corporation. After some research, you narrow the options to two choices, a fiber-based token ring or a copper-based switched network. The pertinent statistics appear in the table below.

Type	fiber-based token ring	copper-based switched network
signal bandwidth	10 MHz	1 MHz
signal-to-noise ratio at transmitter	20 dB	20 dB
attenuation rate	1 dB/km	2 dB/km

The longest link in the network in either case is 10 km.

(a) What link bandwidth is possible according to Shannon's Law
  i. for the fiber network?
  ii. for the copper network?

(b) Assuming that hosts in the copper network can all transmit at their link rate (the values found in part (a)) simultaneously, roughly how many hosts are necessary for the networks to provide equal aggregate bandwidth (the sum of bandwidth for all hosts)?

(c) Using the copper-based network with a 32-point QAM encoding, what modulation rate (baud) is necessary to obtain the bandwidth found in part (a)?

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.

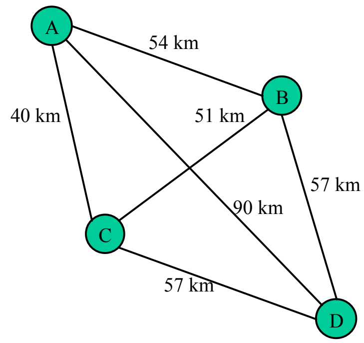

(a) If a transmitter sends at 1 Mbps, how long must packets be to guarantee collision detection by the transmitter?

(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

  i. $p = 1/4$?

  ii. $p = 1/2$?

  iii. $p = 3/4$?

(c) Using the minimum transmission length from part (a) and the probability of successful transmission from part (b)(ii) (for $p = 1/2$), calculate the average throughput of the network if each packet requires 20 bytes of header/trailer and

  i. 10 bytes of data, and

  ii. 50 bytes of data.

3. 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?

4. Forwarding Tables

Consider the network shown in the figure below. The links are labeled with relative costs. The three parts of this problem deal with datagram forwarding, circuit-switched forwarding, and source-routed forwarding, respectively

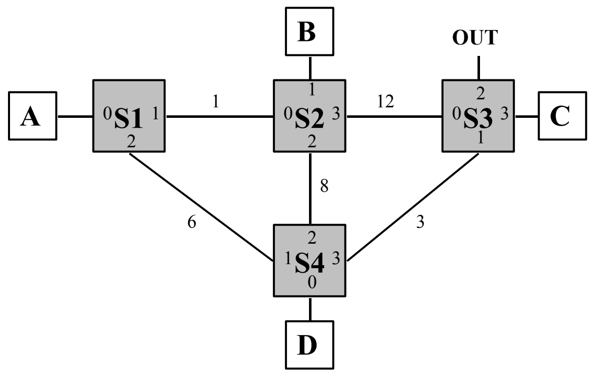

(a) Give the datagram routing table at switch S2, assuming least-cost paths are used. Your table should consist of one row for each possible destination (including the default destination, OUT) consisting of the destination ID, output port, and distance

(b) Suppose virtual circuit forwarding is used for the network shown above with the routing tables show below.

S1:	$port_{in}$	VCI_{in}	$port_{out}$	VCI_{out}
	0	0	2	0
	0	1	1	0
	2	0	1	1

S3:	$port_{in}$	VCI_{in}	$port_{out}$	VCI_{out}
	0	0	3	0
	0	1	2	1
	3	0	0	0

S2:	$port_{in}$	VCI_{in}	$port_{out}$	VCI_{out}
	0	0	2	0
	0	1	3	1
	2	0	3	0
	2	2	1	0
	3	0	1	1

S4:	$port_{in}$	VCI_{in}	$port_{out}$	VCI_{out}
	0	0	2	0
	0	2	1	0
	1	0	2	2
	2	0	0	0

When setting up a new virtual circuit on a given output port, a switch should assign the smallest unused virtual circuit identifier for that port. Indicate how the routing tables change after the following two (cumulative) events: (i) The circuit beginning with (port,VCI)=(0,0) at switch S1 is torn down, and (ii) subsequently, a new circuit is set up from host D to host B using a least-cost path.

(c) Now assume the use of source routing for the network. Indicate the sequence of absolute port identifiers to be found in a packet header for a packet sent by host B destined for host C along the least-cost path. (Assume that the sequence of port identifiers in the header is transmitted in the order written, from left to right.)

5. Error Detection with Cyclic Redundancy Checks

Use the CRC-8 generator polynomial x^8+x^2+x+1 for both parts of this problem.

(a) Calculate the CRC value of the bit sequence **0011 1100 0011**.

(b) Recall that error detection with a CRC works by appending a CRC value to the message to make it a multiple of the generator polynomial. Find a 12-bit burst error polynomial $E(x) = x^{11} + \ldots + 1$ that cannot be detected by a CRC check (with CRC-8).

6. Spanning Tree Algorithm for Intelligent Bridges

Suppose the Perlman spanning tree algorithm and the bridge learning algorithm for forwarding are used for the network shown below.

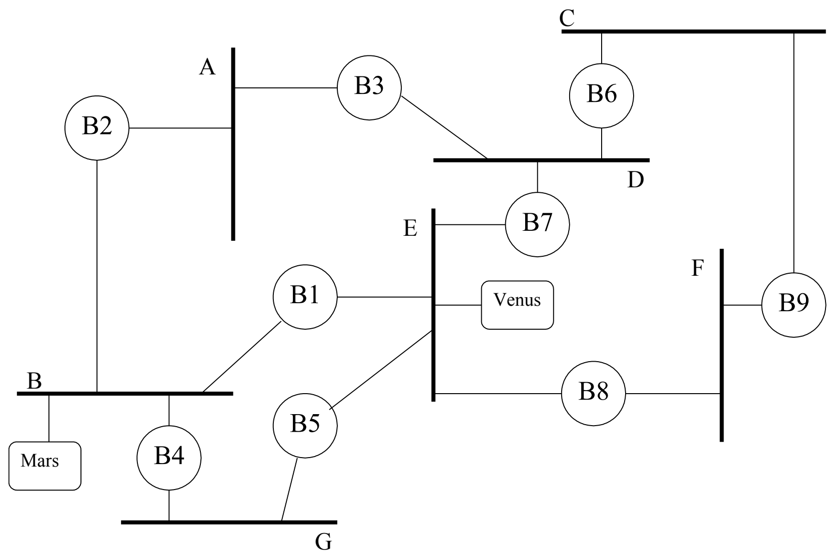

(a) Indicate which bridge is root, which ports are root ports (i.e. the preferred port for reaching the root bridge), which bridge is the designated bridge for each LAN, and which ports are designated ports (i.e. the ports that connect some LAN to its given designated bridge). Hint: bridges that are not designated bridges for any LAN, and ports that are not either root ports or designated ports do not play a role in the routing of packets. The remaining bridges together with the LANs form a spanning tree.

(b) Suppose after the configuration is complete, host Mars attaches to LAN B and host Venus attaches to LAN E. Suppose Mars sends a message to Venus, then Venus sends a message to Mars, then Mars sends a second message to Venus. For each of the three messages, indicate which LANs the message is heard on.

7. Bit- and Byte-Stuffing

Consider a data stream of 8-bit ASCII characters with values 0 to 127. Assume that the probability of a byte assuming each possible is equal (i.e., is exactly 1/128).

(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.

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

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

8. Token Ring

Consider a token ring with a data rate of 75 Mbps, a ring latency of 100 μsec, and 1000 bit packets.

a. Assuming only one host wants to transmit and the delayed token release scheme is used, what is the maximum effective throughput rate that can be achieved? What is the efficiency?
b. Now assume N hosts want to transmit on the token ring and the token holding time (THT) is 250 μsec. What is the token rotation time? What is the maximum effective throughput rate that can be achieved? What is the efficiency?
c. Under the assumptions of part b, and using the immediate release scheme, what is the throughput rate that can be achieved?

9. Link-State Routing

Show how the link-state algorithm builds the routing table for node A in the following network. Use the same format as in P&D (Table 4.9).

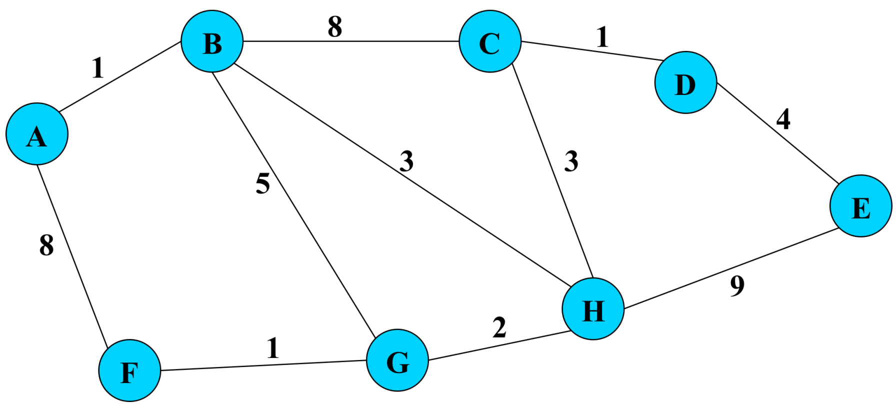

10. Distance-Vector Routing

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

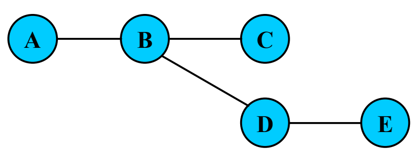

Initially, the router tables for routes to node A look like the following:

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?

Short-Answer Questions

11. Explain the exposed terminal problem and how it is solved.
12. 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.
13. Give two good reasons to allow branching - that is, the ability to support multiple protocols above and below any given protocol - in protocol graphs.
14. Explain one advantage of abstracting networked communication into multiple layers.
15. Explain how a receiver detects the end of a frame with length-based framing.
16. 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.
17. Explain two methods of solving the problem of communicating between machines with mixed endianness on a network.
18. Explain how a receiver detects the end of a frame with sentinel-based framing.
19. What is head-of-line blocking, and when does it occur?
20. Define the Hamming distance of an encoding.
21. Explain the effect of layering on end-to-end bandwidth.
22. Name the OSI layer or layers in which encoding is addressed and state whether encoding is typically handled in hardware, in software, or in both in the Internet architecture.
23. Explain a drawback of forwarding packets with source routing.
24. What delay is relevant and what bandwidth is relevant for computing the delay-bandwidth product of two links in series?
25. Describe a problem associated with communicating between heterogeneous architectures (e.g., a mixture of Sun and Intel hosts) on a network.
26. What purpose do the four addresses in an IEEE 802.11 packet serve?
27. 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.
28. Explain how a receiver detects the end of a frame with clock-based framing (e.g., SONET).
29. Explain what frequency-hopped spread spectrum modulation is, and a motivation for using it.
30. 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?
31. Name the four components that uniquely specify a TCP connection and state the length of each component in bits.
32. Consider a frame consisting of two characters of four bits each. Assume that the probability of error is 10^{-3}, independent for each bit. What is the probability that the frame is received correctly? Add a parity bit to each character. Now what is the probability?
33. State an advantage of direct memory access (DMA) over programmed input/output (PIO).
34. Name the OSI layer or layers in which error detection is addressed in the Internet architecture and state whether error detection is typically handled in hardware, in software, or in both.
35. Name and explain two effects that complicate the process of signal transmission.
36. Explain the main drawback of the use of multiple logical channels for reliable transmission.
37. To provide more reliability than a single parity bit can give, an error detecting coding scheme uses one parity bit for checking all the odd numbered bits and a second parity bit for all the even numbered bits. What is the hamming distance of this code?
38. What does 4B/5B encoding accomplish, besides expanding the number of bits by 25%?
39. Describe the problem solved by error detection.
40. Explain the hidden terminal problem and how it is solved.
41. Name and describe the type of multiplexing traditionally employed in data networks.
42. The data rate of a QAM system using M-ary symbols can be doubled by increasing M and holding the bandwidth and baud rate constant. How must larger must M be, and why?
43. Draw a protocol graph for the Internet, including at least the following: ATM, Ethernet, FDDI, FTP, HTTP, IP, TCP, TFTP, and UDP.
44. In the Perlman distributed spanning tree algorithm, why does the root bridge periodically send messages even after the tree is determined?

45. Name the OSI layer or layers in which reliable transmission is addressed and state whether reliable transmission is typically handled in hardware, in software, or in both in the Internet architecture.
46. What Hamming distance is necessary for n-bit error detection? n-bit error correction?
47. Why are many new local area networks built using multi-mode fiber, despite the fact that single-mode fiber provides higher capacities?
48. Explain a drawback of datagram-based forwarding.
49. State the recursive definition of a network.
50. Explain a drawback of virtual-circuit-based forwarding.
51. Why is byte stuffing necessary with some sentinel-based framing schemes?
52. Explain why a CSMA/CD type protocol cannot be used in a wireless environment.
53. For a small data packet, which is more relevant, bandwidth or latency? Explain.
54. Explain the benefits gained by framing.
55. Under what circumstances will error detection using CRC fail?
56. Describe the benefits of error correction over error detection.
57. In Ethernet, how does a sender detect a collision?
58. Why does Ethernet use binary exponential backoff during contention resolution?
59. Describe the role of the receiver in Ethernet. How is this different from the role of the receiver in IEEE 802.11?
60. Why does Ethernet have a minimum packet size? How is it determined?
61. What is the role of a monitor in a token ring network?
62. What do "learning" bridges actually learn? What do the use this information for?
63. What are the limitations of bridges?
64. How do sniffers work? Will they work on all networks?
65. Why does Ethernet used fixed time slots during backoff? What could go wrong if the fixed slots were not used?
66. In Token Ring networks, why would we want to emulate a shared medium with point-to-point links?
67. What is the role of the NAV in IEEE 802.11?
68. Explain why network administrators use VLANs?
69. 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?
70. Describe "label swapping", and how it is used when setting up virtual circuits.
71. 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 frame size?
72. Why is the port number needed for TCP and UDP communication?
73. Why is a second socket needed in traditional client-server communication over TCP?
74. Why would **`connect()`** be used with a UDP socket?
75. Explain the dangers of not checking the return value of **read()** or **write()**.
76. Why is determining and handling byte order left up to the programmer and not handled by the operating system?
77. How can a simple non-threaded server handle multiple clients without starving one client while waiting for another?
78. What are the limitations of NRZ and NRZI encoding?
79. Explain the efficiency of Manchester encoding.
80. Why can 4B/5B encoding be transmitted using NRZI?
81. What is the relationship between bit rate and baud rate?
82. The FDDI token ring protocol guarantees that each node gets to send a fixed amount of synchronous data and any extra time can be used for asynchronous data. What is the maximum TRT? How frequently can this happen?
83. What is the effect of a wireless signal traveling over multiple paths to a receiver?
84. Why can't we really represent a wireless "link" as a node with a circular radius?
85. Explain the difference between transmission range and interference range in a wireless network.
86. Why is it ineffective to use and ACK for broadcast and multicast communication in wireless networks?
87. What is the impact of using different interframe spacing in IEEE 802.11?
88. Why does a node in 802.11 suspend is collision counter when the medium is busy?
89. What is the impact of a contention window in 802.11 that is too small? Too large?
90. Given an example of when the use of 802.11 can lead to unfairness?
91. In a multihop wireless network, why can't the full link bandwidth be utilized on all links?
92. What is back pressure? Why is it only used in short-RTT situations?
93. Why is buffering needed in switches?

94. When does contention happen in a switch?
95. How does a knockout switch solve the scalability problem?
96. Can distance vector routing detect loops? What about when techniques like split horizon and poisoned reverse are used?
97. What is the effect of setting "infinity" to 16 in distance vector routing?
98. Why doesn't link-state routing scale to large networks?
99. Why are metrics like link utilization and delay difficult to use effectively in routing?
100. What could cause the Internet checksum algorithm to fail to detect an error?