University of Illinois Urbana-Champaign
CS/ECE 438 Final Study Guide (Spring 2026)

The on-campus final examination will be held on Wednesday, May 13, 2026, from 7:00pm–10:00pm. The exam will begin and end promptly. Please arrive early 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 **May 1, 2026**.

- You may not consult any materials during the exam: no textbooks, no crib sheets, no calculator, no phones, no smart watches, etc. **Bring pens and blank sheets of paper only.**
- The final will contain **6 long-answer questions** and **26 short-answer questions**.
  - Each long-answer question will consist of multiple parts, and all long-answer 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 problems and questions that are on this study guide.
- To get full credit for your work, you must show all work and reasoning, writing both work and solution legibly. Please 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.

Long-answer Questions

Question 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 (LANs). 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 payload data. In this problem, we assume that frame headers and trailers are removed/added by the network adapters. 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 architecture. 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 over each network link? Express your answer as the total number of bits per second transmitted, including headers and trailers.

Question 2: 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 $E(x) = x^{11} + \cdots + 1$ that cannot be detected by a CRC check (with CRC-8). Essentially, the burst error starts with and ends with a 1 and you need to determine all the bits in-between.

Question 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 1/256).

(a) Using the 8-bit HDLC bit-stuffing protocol where a 0 is stuffed in after every 5 consecutive 1s, 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.

Question 4: 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 an $n \times n$ 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 a Batcher network with $n = 16$.

Question 5: BGP Policy

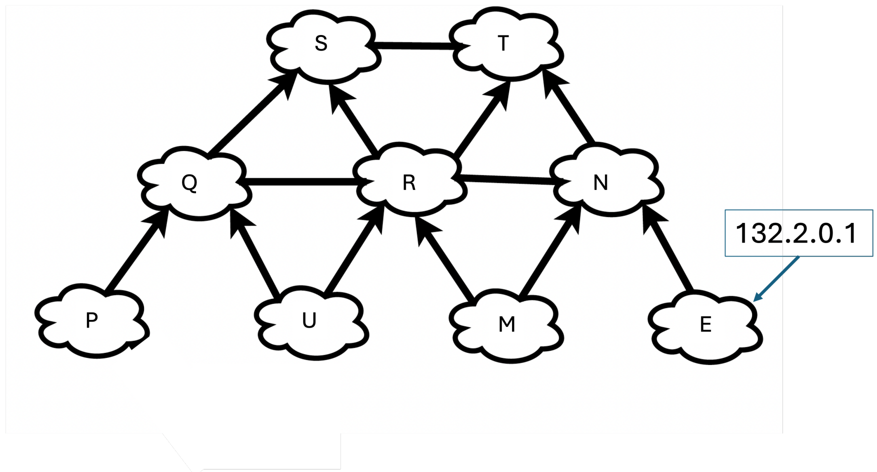

The above figure represents the relationship between ASes. The vertices are individual ASes and edges are links between them and IP is the prefix from AS E. Also suppose that arrows represent customer-provider relationships where the customer points to its provider. An edge without arrows represents a link between peers.

a. Suppose all ASes follow local preference rules that enforce *valley-free* paths: any path must follow a sequence of zero or more provider links, followed by at most one peer link, followed by a sequence of customer links. An AS will route through the *valley-free* path with the least number of hops. List the routes that each AS will follow to reach E in a *valley-free* manner.

AS	Path to E
N	
R	
M	
T	
S	
U	
Q	
P	

b. Suppose AS S does not like AS R. Using only BGP, is it possible for AS S to implement a policy stating that "traffic outbound from my AS should not cross R"? If it is possible, show that S can still reach all ASes using *valley-free* paths that do not cross R. If it is not possible, show that there exists an AS such that any *valley-free* path from S must go through R.

c. Suppose AS R does not like AS P, and therefore decides to not forward any traffic from P. Can AS P deal with this change? If it can, show that P can find *valley-free* paths to all ASes that do not cross R. If it cannot, show that there exists an AS such that any *valley-free* path from P must go through R.

Question 6: SDNs

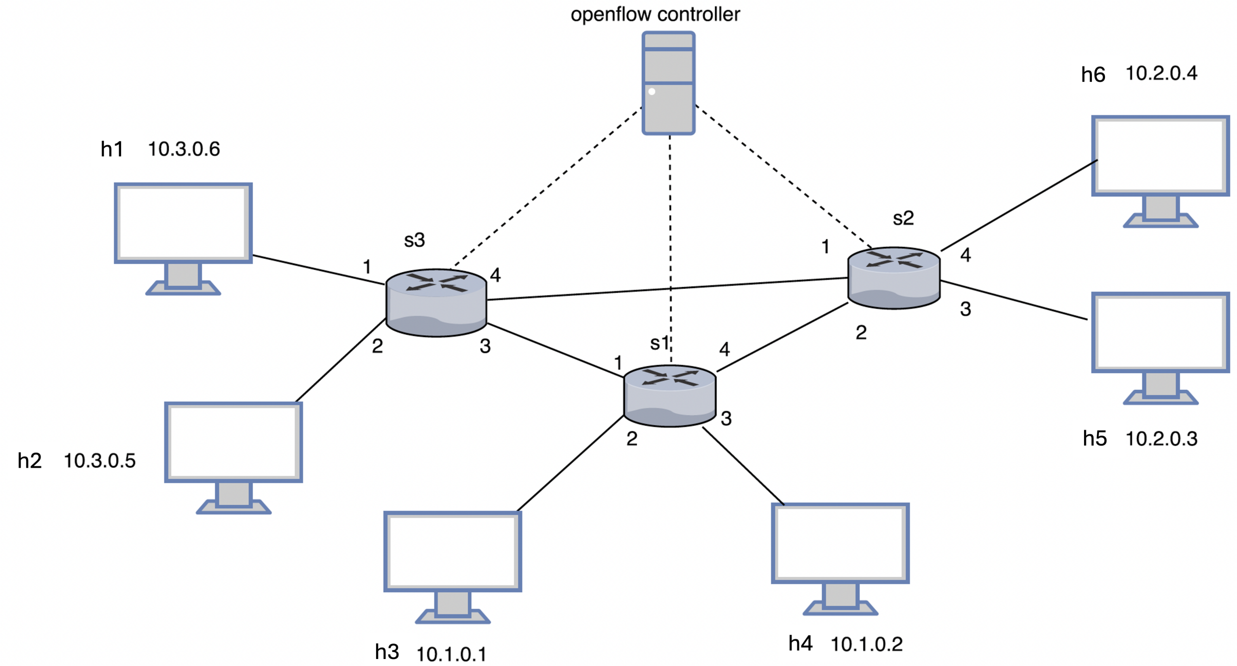

Consider the SDN OpenFlow network shown above. Suppose we want switch s2 to function as a firewall. Specify the flow table in s2 that implements the following firewall behaviors. Specify a different flow table for each of the four firewalling behaviors below. The flow table should only consider delivery of datagrams destined to h5 and h6. You do not need to specify the forwarding behavior in s2 that forwards traffic to other routers. The flow table should show the matching rule and the action taken. Add rows or leave rows empty as needed.

a. Only traffic arriving from hosts h1 and h2 should be delivered to hosts h5 or h6.

Match Rule	**Action**

b. Only TCP traffic is allowed to be delivered to hosts h5 or h6. (i.e., that UDP traffic is blocked.)

Match Rule	**Action**

c. Only traffic destined to h5 is to be delivered (i.e., all traffic to h6 is blocked.)

Match Rule	**Action**

d. Only UDP traffic from h4 and destined to h6 is to be delivered. All other traffic is blocked.

Match Rule	**Action**

Question 7: HTTP

Suppose a webpage has nothing but 30 large images, each 4 MB in size. The size of the webpage itself is negligible. A client wants to access the webpage and load the images in their browser. Between the client and server, the RTT is 60 ms, and the transmission rate is 50 Mbps. How long will it take to load the webpage in each of the following cases?

Be sure to account for the time needed to establish the connection and retrieve the webpage itself. You do not need to account for the delay from closing the connection.

a. Using Non-Persistent HTTP?

b. Using Persistent HTTP?

c. Using Pipelined Persistent HTTP?

Question 8: IP Fragmentation

Consider two hosts, A and B, each on a separate shared Ethernet with MTU = 1500 B. In addition to these LANs, 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 B. 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 B. Also recall that IP fragmentation breaks data along 8 B boundaries.

a. An application on host A passes 2400 B of data to TCP. Following the approach used to draw Figure 73 from Section 3.3.2 of the Peterson & Davie book, 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 B 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 B to an application on host B given $p = 3/4$, 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?

Question 9: Comparison of Multiplexing Strategies

Consider a 15 Mbps 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 bit, 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, counting "time in queue + time in service", but not the propagation delay.

For each of the three scenarios below, find:

i. the mean delay for flow 1,

ii. the mean delay for flow 2, 3, or 4, and

iii. the mean delay, averaged over all packets of all flows.

The last of these can be computed as

$$\frac{1}{2}(\text{flow 1 delay}) + \frac{1}{6}(\text{flow 2 delay}) + \frac{1}{6}(\text{flow 3 delay}) + \frac{1}{6}(\text{flow 4 delay}),$$

or by using Little's Law, as

$$\frac{\text{mean amount of data queued or in transmission}}{\text{sum of arrival rates}}.$$

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.

(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.

(b) Same as (a), but instead of equal allocations, the allocations are proportional to the arrival rates. Thus, flow 1 is allocated 7.5 Mbps, while the other three flows are allocated 2.5 Mbps each.

(c) 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.

Question 10: Token-Bucket Filtering

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

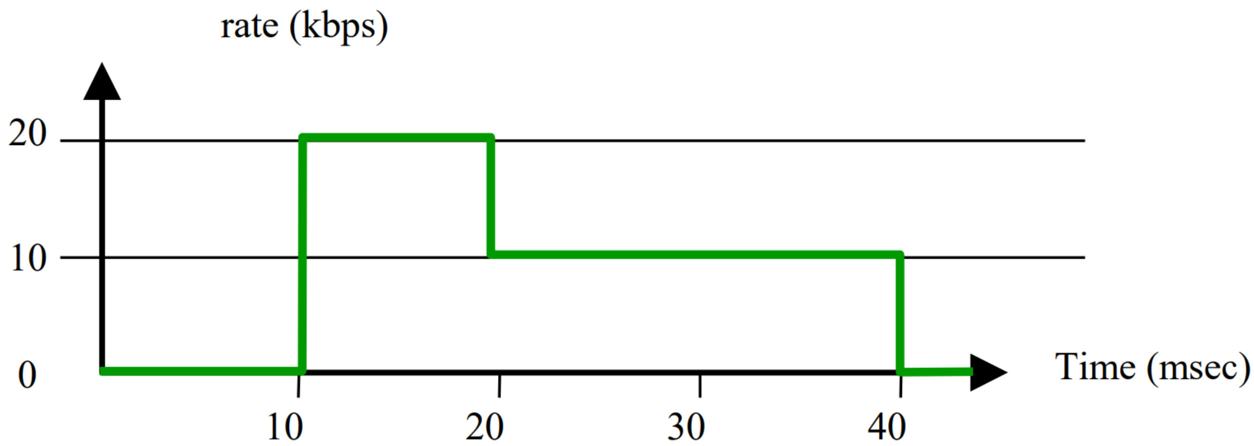

Suppose this stream is passed into a token bucket filter with token generation rate r (where one token is needed per bit) and the token buffer size is B.

(a) Suppose $r = 10$ kbps (i.e. 10 kilotokens/sec). 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 $r = 5$ kbps.

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

Question 11: 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 kB.

Question 12: 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.

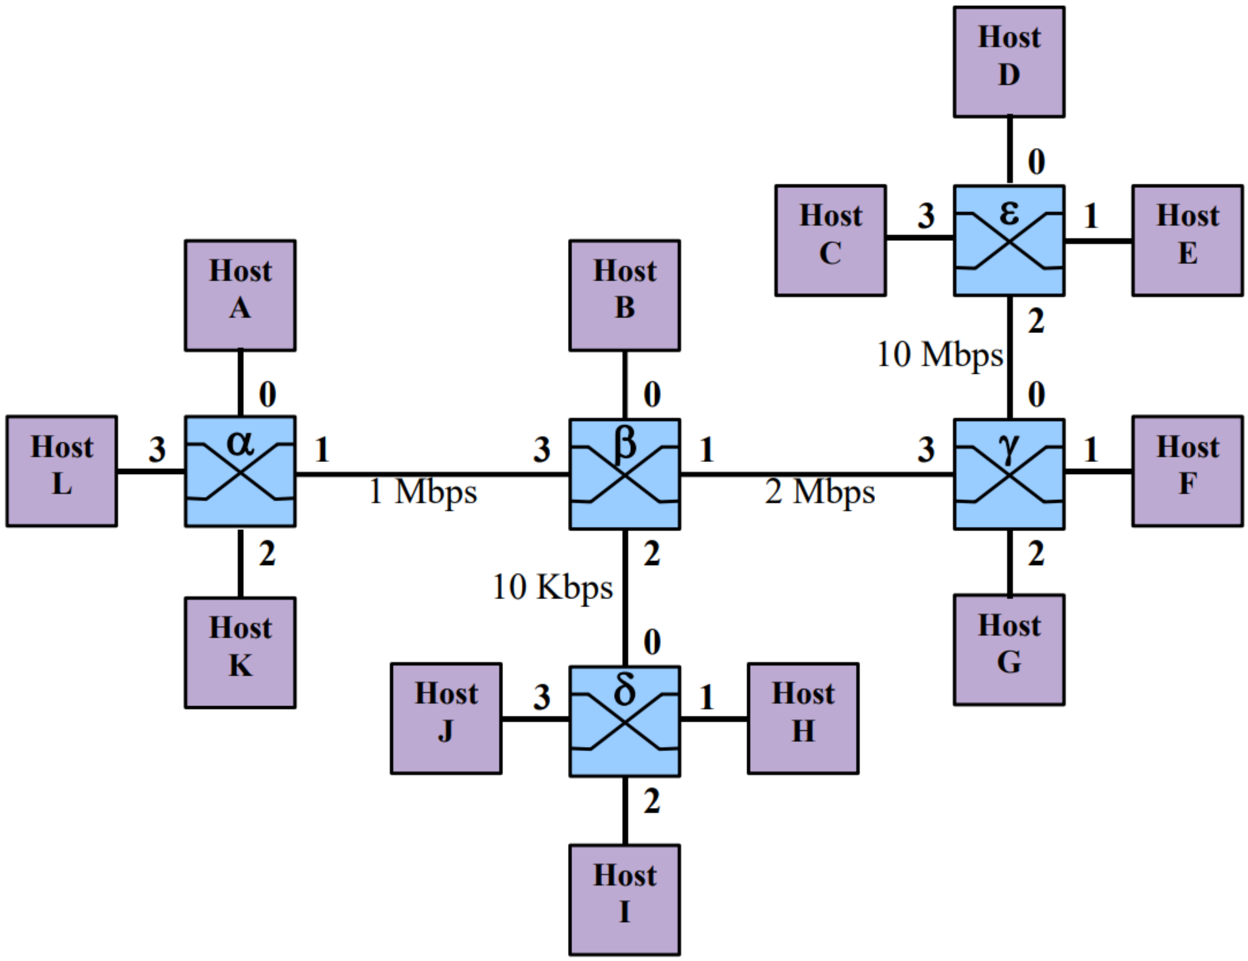

Given the following bandwidth requests for the connections, determine a fair allocation of resources for each connection. For this question, use the definition of max-min fairness discussed in class.

Connection	Requested
$L \to E$	1 Mbps
$A \to I$	5 kbps
$I \to E$	7 kbps
$F \to K$	1 Mbps
$G \to B$	1 Mbps
$C \to F$	10 Mbps

Question 13: 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 $t = 0$.

(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 B_3 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.

Question 14: 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 ms. All packets sent across the link are 1024 B 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 (ACKs) are always received (i.e. 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 bps) achieved for long streams of data.

Question 15: 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 round-trip 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)?

Short-Answer Questions

1. Explain the effect of layering on end-to-end bandwidth.

2. Explain a drawback of forwarding packets with source routing.

3. Explain what frequency-hopped spread spectrum modulation is, and a motivation for using it.

4. Suppose packets on a 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?

5. Name the four components that uniquely specify a TCP connection and state the length of each component in bits.

6. 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?

7. After a collision, how would you determine the probability of a successful transmission in Ethernet?

8. Describe the problem solved by error detection.

9. Explain the benefits gained by framing.

10. Under what circumstances will error detection using CRC fail?

11. Describe the role of the receiver in Ethernet. How is this different from the role of the receiver in IEEE 802.11?

12. What do "learning" bridges actually learn? What do they use this information for?

13. Why is the port number needed for TCP and UDP communication?

14. Why is a second socket needed in traditional client-server communication over TCP?

15. What are the limitations of NRZ and NRZI encoding?

16. Explain the efficiency of Manchester encoding.

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

18. What is the impact of a contention window in 802.11 that is too small? Too large?

19. Give two arguments against IP reassembly in routers.

20. List three components that contribute to end-to-end latency.

21. 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?

22. Explain the circumstances that give rise to the count-to-infinity problem.

23. 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.

24. 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?

25. 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.

26. At what OSI layer do Internet routers typically operate?

27. 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?

28. Why will UDP require a checksum with IPv6?

29. What does TCP use in addition to an estimate of RTT to calculate timeouts for adaptive retransmission?

30. 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?

31. Explain in words (no equations) what the memoryless property of a random, exponentially distributed lifetime is.

32. How does TCP guarantee that new connections do not receive segments from previous incarnations of the connection?

33. Give an argument why the leaky bucket algorithm should allow just one packet per tick, independent of how large the packet is.

34. What is the purpose of the protocol field in the IPv4 header?

35. List five services demanded by many applications but not provided by IP (nor typically provided by user-level code).

36. Explain the fundamental conflict between tolerating burstiness and controlling network congestion.

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

38. 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.

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

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

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

42. 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?

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

44. 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.

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

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

47. 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)?

48. TCP waits until it has received three duplicate ACKs 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?

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

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

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

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

53. 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?

54. 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.

55. 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?

56. 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?

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

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

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

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

61. 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?

62. 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?

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

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

65. 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.

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

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

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

69. 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?

70. 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?

71. 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?

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

73. Explain Karn's algorithm (aka. Karn-Partridge algorithm). Why do we need it?

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

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

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

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

78. Describe the problem solved by sliding window-based ARQ protocols.

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

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

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

82. 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.

83. 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?

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

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

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

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

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

89. Why does BGP use path-vector routing?

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

91. What is the effect of setting "infinity" to 16 in distance vector routing?

92. Explain the difference between flow control and congestion control.

93. Explain the difference between rate-based and window-based flow control.

94. How is self-clocking used in TCP?

95. What protocol was the precursor for both the Ethernet and Wi-Fi MAC protocols?

96. Give one reason why DNS uses UDP by default.

97. During a recursive DNS query, which name servers are contacted by the local DNS server in sequence?

98. How does DNS do a reverse lookup of an IP address 1.2.3.4?

99. Describe the head-of-line (HOL) blocking issue in HTTP/1.1 pipelining that HTTP/2 fixed.

100. Describe the head-of-line (HOL) blocking issue in HTTP/2 that HTTP/3 fixed.

101. Describe how QUIC reduces the connection establishment time.

102. Unlike TCP in kernel space, QUIC over UDP runs in user space. Briefly describe the pros and cons.

103. Explain how HTTP cookies help maintain state between a client and a server.

104. Describe two benefits of web caching on proxy servers.

105. To implement a firewall and a NAT using OpenFlow, what are the example match-and-action rules, respectively?

106. Explain the data plane and control plane in SDN.

107. Explain one advantage and one challenge of a logically centralized SDN controller.

108. Briefly describe one SDN killer app: what is the app, and why it benefits from SDN.

109. Give two reasons why packet loss may not be a good congestion signal.

110. How does Explicit Congestion Notification (ECN) notify a TCP sender of in-network congestion? (Note the relevant network-layer header but the exact header fields are not needed.)

111. Explain the AS-PATH and NEXT-HOP attributes in BGP.

112. How does a BGP router decide path selection and path export, respectively?

113. Name two issues with BGP.

114. Explain how BGP and IGP (e.g., OSPF) work together for internet routing.

115. Describe the pros and cons of NAT.