NET ID ____________

CS519, Prelim (March 17, 2004)

NAME: ____

You have 50 minutes to complete the test.

NET ID ___________

Q1. 2 points

Write your NET ID at the top of every page of this test.

Q2. X points

Name 3 advantages of a circuit network over a packet network:

1 ____

2 ____

3 ____

Q3. X points

Name 3 advantages of a packet network over a circuit network:

1 ____

2 ____

3 ____

Q4. X points

Which of the following best characterizes a “long fat pipe” (i.e. high delay x bandwidth product)? Check one:

_______ OC12 link from between the east and west coasts of the USA

_______ 1Mbps link from New York City to Tokyo Japan

_______ 100 Gbps LAN

NET ID ___________

Q5. X points

Two hosts with private addresses 10.9.8.7 and 10.9.8.8 respectively both have LAN interfaces. Do you know for sure that they are connected to the same LAN? (Assume that they are properly connected and addressed etc.)

________ YES ________ NO

Why or why not? ____

____

Q6. X points

Two hosts with global addresses 20.9.8.7 and 20.9.8.8 respectively both have LAN interfaces. Do you know for sure that they are connected to the same LAN? (Assume that they are properly connected and addressed etc.)

________ YES ________ NO

Why or why not? ____

____

Q7. X points

Why isn't Internet routing run as one huge unified routing domain (rather than the 2-tier intra-domain and inter-domain architecture we have today)? Check each reason that is true for the 2-tier architecture, but would not be true for a unified routing domain architecture.

_______ Routing convergence would be too slow (i.e. after a change)

_______ Sites want to run different routing protocols

_______ External routing errors may effect routes within a site (i.e. between two hosts in the same site)

_______ ISPs do not want to receive routing updates showing the internals of sites (i.e. the subnets in a site)

NET ID ___________

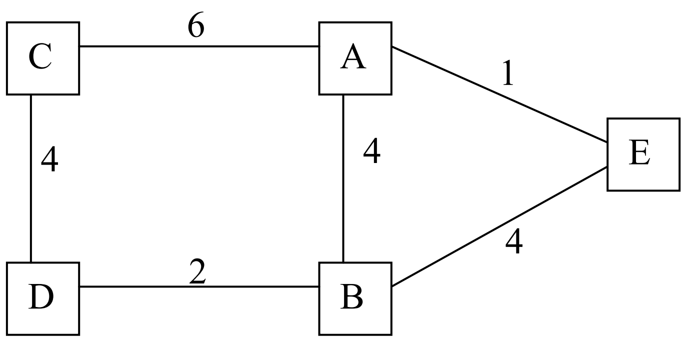

Figure 1: A network with five routers, A - E. Assume that the costs shown are for both directions of the link.

Q8. X points

In the space below, write two RIBs for router E in Figure 1, one of them assuming a Distance Vector (DV) protocol, and the other assuming a Link State (LS) protocol. Clearly label which is which.

NET ID ___________

Q9. X points

Fill in the contents of the confirmed and tentative lists for each step of a shortest path spanning tree algorithm for Router E in Figure 1 using the (Dest, cost, next-hop) notation. Steps 1 and 2 already been filled in. Note there may be fewer than 12 steps.

Step	Confirmed	Tentative	Step	Confirmed	Tentative
1	(E,0,-)		2	(E,0,-)	(A,1,A) (B,4,B)
3			4		
5			6		
7			8		
9			10		
11			12		

NET ID ___________

Q10. X points

Name 3 advantages of TCP over UDP:

1 ____

2 ____

3 ____

Q11. X points

Name 3 advantages of UDP over TCP:

1 ____

2 ____

3 ____

Q12. X points

In a sliding window protocol, normally the receive window is as big as or bigger than the send window. What happens if the send window is bigger than the receive window? Check one:

_______ It can no longer be guaranteed that all bytes will be received.

_______ It can no longer be guaranteed that all bytes will arrive in order.

_______ It can no longer be guaranteed that all bytes will be received or will arrive in order.

_______ None of the above

NET ID ___________

Q13. X points

In class we saw that a small packet (and corresponding small ACK) can be introduced into the TCP "control loop" when the sending application does a send of only a few bytes (and the TCP layer subsequently sends the bytes after a short time waiting for another send). Using only the space provided below, describe another way in which a small packet can be introduced.

____

____

____

____

Q14. X points

What are the three defining characteristics of "quality" or "performance" in packet networks?

____

NET ID ____________

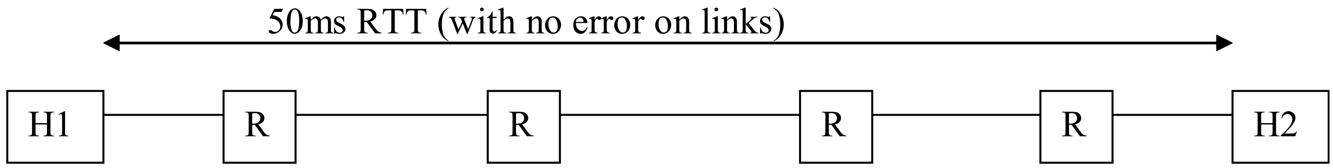

Figure 2:

Q15. X points

In the network of Figure 2, each of the five links exhibits a high error rate (5% packet loss). The RTT between the two hosts is 50ms. The two hosts are running TCP. Imagine that, by introducing a link-level reliable sliding window protocol on each link, you can reduce the error rate to 0% of each link, but in doing so you add 5ms of latency to each link, increasing the RTT to 75ms. Given that the two hosts are running TCP, does the introduction of the link level reliability protocols violate the End-to-end principle?

________ YES ________ NO

Why or why not? ____

____

____

____

____

____

NET ID ___________

Q16. X points

From the point of view of a router in a backbone ISP (i.e. one that does not participate in any site routing), what is the structure of the IP address? Please check one of the following:

_____ | prefix | suffix |

_____ | ISP | Site | suffix |

_____ | ISP | Site | Subnet | Host |

_____ | Class A/B/C Network | Host |

Q17. X points

Image a very large network (thousands of destinations) where only a small fraction of destinations need to be reachable. For instance, this might be the case for a large sensor network where a few locations gather data from the sensors. As a result, the routing algorithm only needs to keep track of a few destinations. Given this, would it make more sense to use a link state routing algorithm or a distance vector routing algorithm?

Q18. X points

What is the primary cause of large routing tables in BGP? Check one:

_______ Large address size

_______ Site multihoming

_______ Increasing number of home networks

_______ Increased NAT usage

NET ID ____________

Q19. X points

When an application does a `send()` on a TCP socket, and the `send()` returns with no errors, what does the application know for sure about the sent data. Check one:

_______ The data has been written into the local TCP's send buffer

_______ The data has been transmitted onto the network

_______ The data has been received by the remote host's TCP layer

_______ The data has been received by the remote host's application

Q20. X points

Of the following TCP socket calls, check those that are used by the client and/or by the server.

Used by client	Used by server	Socket call
		`connect()`
		`listen()`
		`accept()`
		`recv()`
		`send()`
		`close()`

Q21. X points

With that "fast retransmit" method of TCP retransmission (i.e. after three duplicate ACKs), is it possible for a retransmit to happen sooner than one RTT from the time the lost packet was sent?

________ YES ________ NO

Why or why not? ____

____

____

____

NET ID ___________

Q22. X points

With streaming media, the media is downloaded for a while into a local buffer (often for many seconds) before the media starts to play on the local player. This occurs even if the average throughput between the media server and the player is higher than that required for the media player. Why is this necessary? Check one only (the one which most drives the need for buffering).

_______ The buffer eliminates traffic jitter

_______ The buffer allows the player to read data in a uniform way regardless of how the stream was received (i.e. UDP, TCP, or even HTTP)

_______ The buffer hides the fact that throughput may vary considerably over the length of the stream

_______ The buffer allows the player to request retransmissions of lost packets and deliver them in order

Q23. X points

Imagine that an IP layer sends two packets, one fragmented and the other not fragmented. Suppose also that two simultaneous bit errors on the unfragmented packet cause its ID field to match that of the fragmented packet in such a way that the checksum is still correct (admittedly a low probability event, but possible). Finally assume that at the receiver, the unfragmented packet is received after the first fragment of the fragmented packet, but before the last fragment of the fragmented packet (see below). Does the receiver receive the packets correctly or not?

________ Received correctly ________ Not received correctly

Why or why not? ____

____

____

____

NET ID ___________

Q24. X points

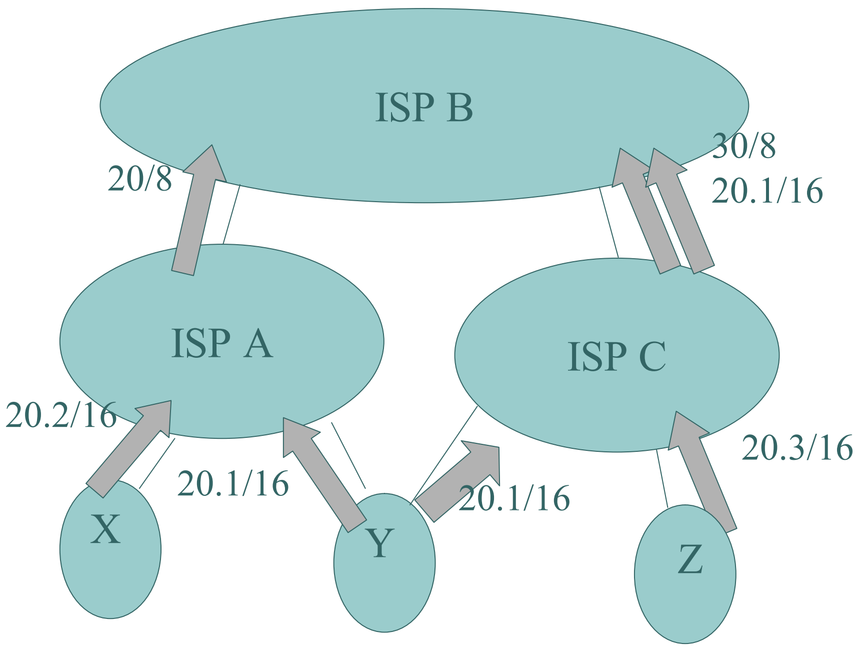

Figure 3

Describe the AS path (assume that all Sites X, Y, and Z, and ISPs are ASs) from each ISP to each Site, using the table below. If there is no path from the ISP to the Site, write "none".

From-to	AS-path	From-to	AS-path
A→X		B→Z	
A→Y		C→X	
A→Z		C→Y	
B→X		C→Z	
B→Y			

NET ID ___________

Q25. X points

Referring to the following figure (which is identical to Figure 1), imagine that there is a lot of traffic from router E to router D, so much so that the E-B and B-D links are saturated with traffic. Modify the link costs so that all traffic except the E to D traffic avoids the two saturated links.

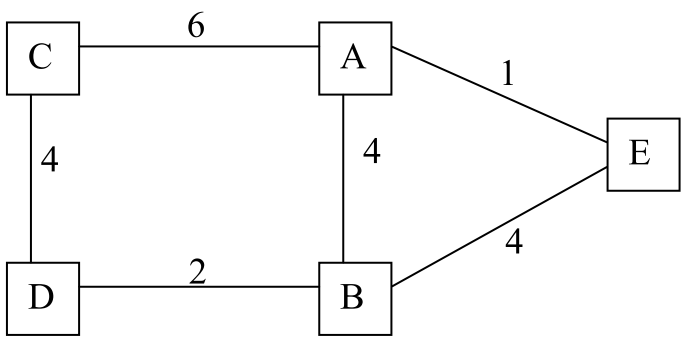

Q26. X points

Given current IP address assignment policies, what is the theoretical smallest routing table size in the backbone ISP routers. Check one:

_______ Number of class A addresses (256)

_______ Number of large sites

_______ Number of backbone ISPs

_______ Number of ISPs of all types

NET ID ___________

Q27. X points

Consider the two traffic sequence diagrams in Figure 4 below. Assume identical time scales, and assume that the same sized data packets are being sent. Assume a send window of 4 packets. Assume that the sender is sending as fast as it can, within the limits of the send window.

Of these two sequences, which has higher bandwidth (select one)?

__________ Left sequence   _________ Right sequence

Of these two sequences, which has longer latency (select one)?

__________ Left sequence   _________ Right sequence

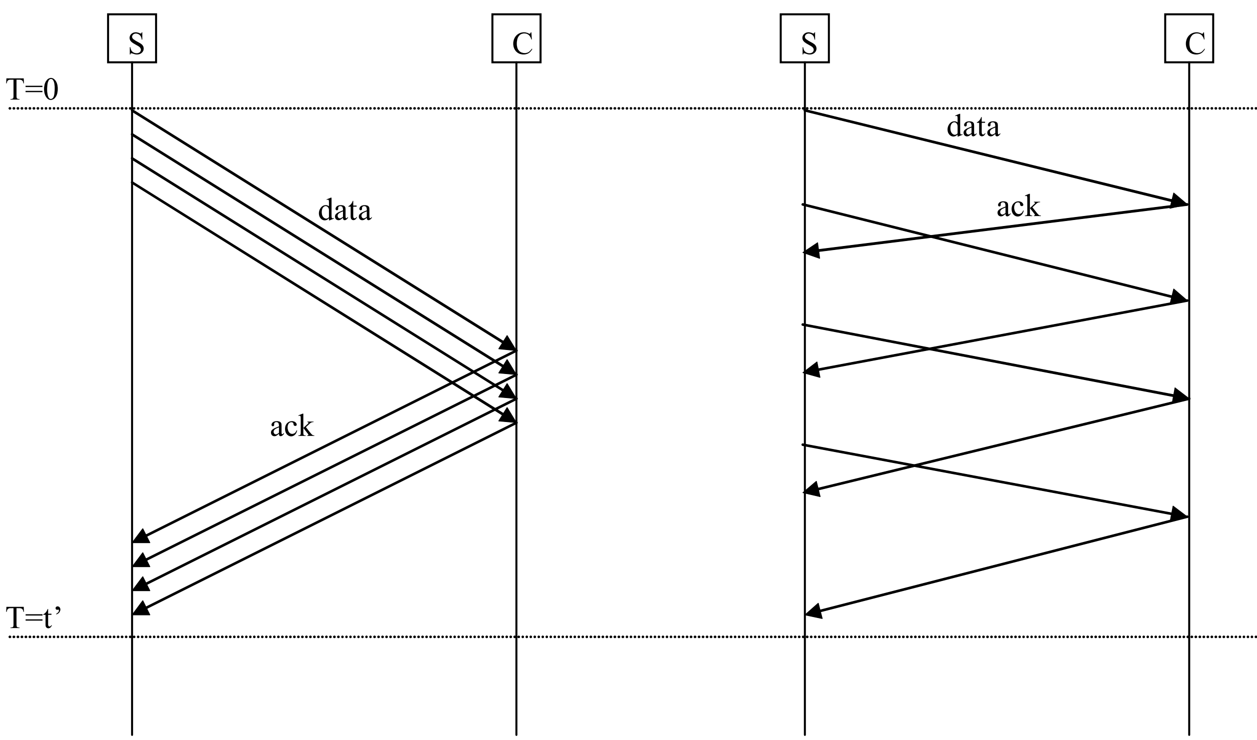

Figure 4

NET ID ___________

Q28. X points

Describe up to five ways by which an application can either close an open TCP socket, or learn that the other end has closed the socket.

1. ____
2. ____
3. ____
4. ____
5. ____

Q29. X points

For each of the following two scenarios, what contributes most to the total delay, queuing, propagation, or transmit?

______________ A 1500 byte packet transmitted over a lightly loaded 10Mbps Ethernet

______________ A 10Gbps transmission from Boston to Los Angeles over a high-speed network

Q30. X points

To which of the following blocks of addresses does the address 64.65.24.7 belong? Check all that apply:

________ 64.65.24.0/28 ________ 64.65.16/21

________ 64.64/12 ________ 1/0

________ 0/0

NET ID ___________

Q31. X points

The following question applies to Figure 1. Assuming a distance vector routing protocol, which node's RIBs *do not* change if link C-D goes down. Which node's FIBs *do not* change?

RIBs that don't change ________________________________

FIBs that don't change ________________________________

Q32. X points

What effect does BGP policies have on paths in the Internet? Check all that apply:

_______ Paths are longer than the shortest possible

_______ Paths have low latency

_______ Paths are asymmetric

_______ Paths cannot always be trace routed.

Q33. X points

Assuming a Go-Back-N acknowledgement scheme (with no selective acknowledgement), if a sender sends two packets (1 and 2), and receives back an ACK for the bytes in packet 2, but not for the bytes in packet 1, what does the sender know for sure? Check one:

_______ The receiver has not received packet 1

_______ Either packet 1 was dropped, or the ACK to packet 1 was dropped

_______ The ACK to packet 1 was dropped

_______ None he above

NET ID ___________

Q34. X points

Assume 1000 byte packets being send over TCP. Assume that the sender has sent 10 packets constituting bytes 10000 through 19999. Assume further that the sender receives a TCP packet with an acknowledgement of 13000, and with a SACK option indicating the ranges 15000-15999, and 18000-19999. Which bytes can the sender remove from its send buffer?

Ans: ____