NAME:

Roll Number:

Computer Science- Network Fundamentals
Mid-Term
March 8, 2025
9.30 AM - 11.30 AM

Question 1: Simple questions – (15 points)

Answer each of the following questions briefly, i.e., in at most a few sentences.

a) We have learned that the Internet provides "best effort service." What is meant by a "best effort" service model? Give an example of another service model that a network might offer.

b) What is the purpose of the HTTP "COOKIE:" field? Are the values in the HTTP message's cookie field stored at the client or server or both? Explain briefly.

c) We saw that TCP and UDP provide two very different service models. Suppose that an application wants all of the functionality provided by UDP but only some of the functionality provided by TCP (e.g., the application wants reliable message transfer and flow control, but not congestion control). How would an application get this different service in today's Internet?

d) Consider a server-side socket that is used to communicate from server to client. In the case of a TCP socket, can data being read from the server-side socket have been sent by more than one client? Explain briefly. In the case of a UDP socket, can data being read from the server-side socket have been sent by more than one client? Explain briefly. UDP, any client can send to the (same) UDP socket on the server.

e) What is the difference between congestion control and flow control

f) What is meant by the term "encapsulation"? wraps it in a new packets with new header

g) In network-assisted congestion control, how is congestion in the network signaled to the sender

Question 2: Simple calculations - (15 points)

1) Suppose there are 5 users whose traffic is being multiplexed over a single link with a capacity of 1 Mbps.

   a) Suppose each user generates 100 kbps when busy, but is only busy (i.e., has data to send) 10% of the time. Would circuit-switching or packet-switching be preferable in this scenario? Why?

   b) Now suppose that the link capacity is still 1 Mbps, but the amount of traffic each user has to send when busy is increased to 1 Mbps, and that each of the 5 users still only has data to send 10% of the time. Would circuit-switching or packet-switching be preferable in this scenario? Why?

2) Suppose a TCP SYN message is sent from a client with IP address 128.119.40.186 and client port number 5345 to a server with IP address 41.123.7.94 and server port number 80 (HTTP)

   a) Once the TCP connection has been established, what will be the client-side IP address, client-side port number, server-side IP address and server-side port number of the TCP segment carrying the HTTP GET message

   b) Will the TCP SYN message and the HTTP GET message be directed to the same socket at the server? Explain in one or two sentences.

3) Consider a sliding window reliable data transfer protocol (e.g., Go-Back-N, or Selective Repeat). Suppose the sender window contains sequence numbers in the range [x,x+N]. Is it possible for the sender to receive an ACK that is for a sequence number that is less than x?

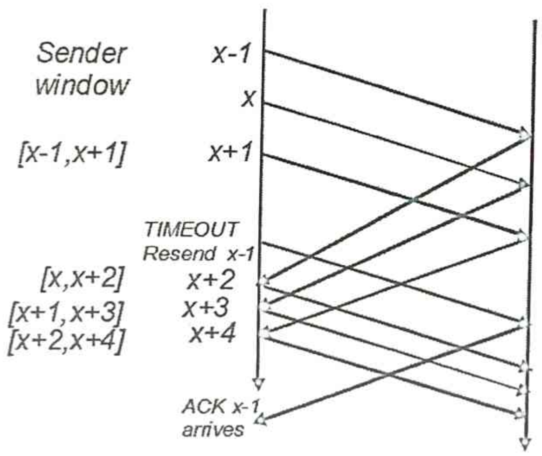

2

Question 2: Delays, Throughput and Caches - (20 points)

Consider the scenario in the figure below in which a server is connected to a router by a 100 Mbps link, with a 100 ms propagation delay. That router in turn is connected to two routers, each over a 25 Mbps link with a 200 ms propagation delay. A Gbps link connects a host and a cache (when present) to each of these routers; this link, being a local area network, has a propagation delay that is essentially zero. All packets in the network are 10,000 bits long.

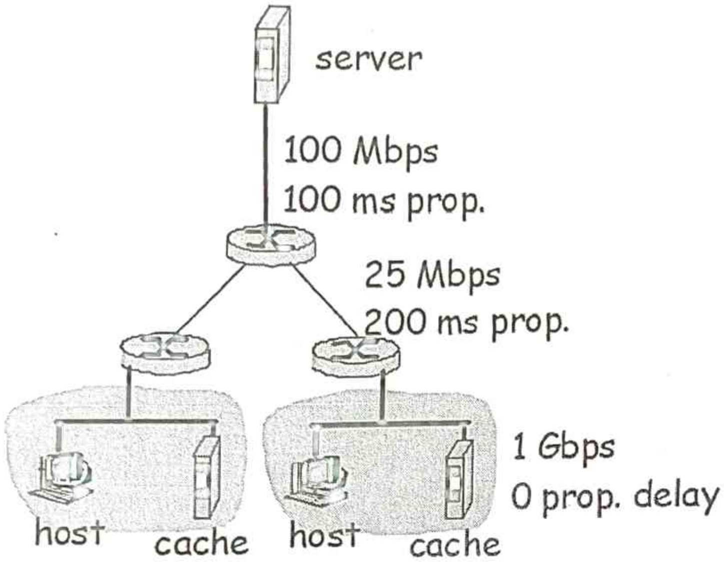

a) What is the end-to-end delay from when a packet is transmitted by the server to when it is received at a host? Assume that there are no caches, that there is no queueing delay at a link, and that the node (router) packet-processing delays are also zero. Answer: If all packets at 10,000 bits long, it takes 100 usec to send the packet over a 100Mbps link, 400 usec to send over a 25 Mbps link, and 10 usec to send over a gigabit link. The sum of the three link transmission times is thus 510 usec. The sum of the propagation delays is 200+100=300 msec.

b) First assume that client hosts send requests for files directly to the server (i.e., the caches are off). What is the maximum rate at which the server can deliver data to a single client, assuming no other clients are making requests.

c) Again assume that only one client is active, but now suppose the caches are HTTP caches and are turned on. A client HTTP GET is always first directed to its local cache. 50% of the requests can be satisfied by the local cache. What is the maximum rate at which this client can receive data in this scenario?

d) Now suppose that the clients in both LANs are active and the HTTP caches are on, as in c) above. 50% of the requests can be satisfied by the local cache. What is the maximum rate at which each client can receive data, in this scenario?

e) Now suppose the 100 Mbps link is replaced by a 25 Mbps link. Repeat question d) above in this new scenario.

1000×0.5 + 0.5×12.5

Question 4: Sliding Window Protocol - (10 points)

a) Consider the sliding window protocol in Figure below Does this figure indicate that Go-Back-N is being used, Selective Repeat is being used, or there is not enough information to tell? Explain your answer briefly.

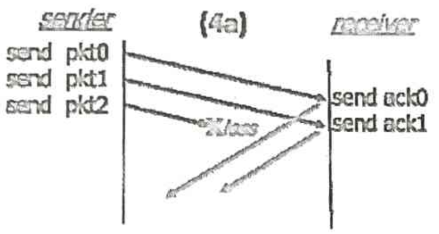

b) Consider the sliding window protocol Figure (4b) below. Does this figure indicate that Go-Back-N is being used, Selective Repeat is being used, or there is not enough information to tell? *Explain your answer briefly.*

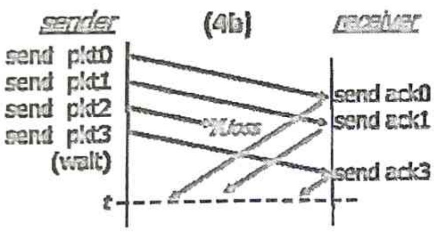

c) Consider Figure (4b) again. Suppose the sender and receiver windows are of size N = 4 and suppose the sequence number space goes from 0 to 15. Show the position of the sender and receiver windows over this sequence number space at time t (the horizontal dashed line).

d) Suppose that it take 1 ms to send a packet, with a 10 ms one-way propagation delay between the sender and receiver. The sliding windows size is again N = 4. What is the channel/link utilization?

4

Question 5: TCP - (20 points)

a) What is the purpose of the receiver-advertised window in TCP ?

b) Identify 4 fields in the TCP header, and give a one-sentence description of each of their uses.

c) Consider two TCP sessions that must share a link's bandwidth. One of the TCP connections has been running for quite some time and has built up a large TCP sending window. The second connection then starts up with an initially small window. Long term, what will be the relative throughput achieved by these two TCP sessions? Explain.

d) Consider a TCP session and a UDP session that must share a link's bandwidth. Of course, both sessions would ideally like to send as fast as they can. Long term, what will be the relative throughput achieved by these two sessions. Explain.

e) Suppose we want to modify TCP so that it counts the number of segments that are lost in transit. What are the difficulties of doing this in TCP?

Question 6: Congestion Control - (20 points)

Consider the following graph of TCP throughput (**NOT DRAWN TO SCALE**), where the y-axis describes the TCP window size of the sender. We will later ask you to describe what happens on the right side of the graph as the sender continues to transmit.

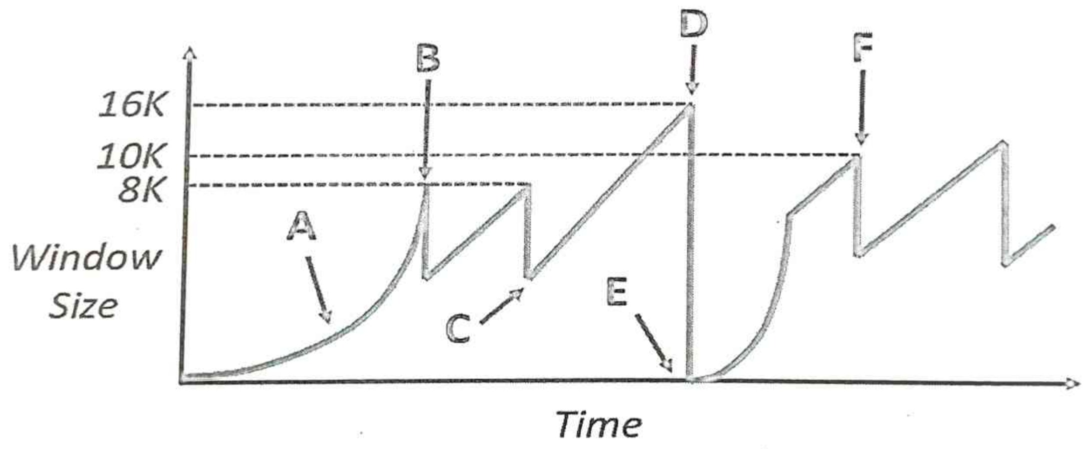

1. The window size of the TCP sender decreases at several points in the graph, including those marked by B and D.

a) Name the event at B that occurs that causes the sender to decrease its window.

5

b) Does the event at B necessitate that the network discarded a packet (Yes or No)? Why or why not?

c) Name the event at D that occurs that causes the sender to decrease its window. Amount

d) Does the event at D necessitate that the network discarded a packet (Yes or No)? Why or why not?

e) For a lightly-loaded network, is the event at D MORE likely or LESS likely to occur when the sender has multiple TCP segments outstanding? (Write "MORE" or "LESS")

2. Consider the curved slope labeled by point A. Why does the TCP window behave in such a manner, rather than have a linear slope? (Put another way, why would it be bad if region A had a linear slope?)

3. Assume that the network has an MSS of 1000 bytes and the round-trip-time between sender and receiver of 100 milliseconds. Assume at time 0 the sender attempts to open the connection. Also assume that the sender can "write" a full window's worth of data instantaneously, so the only latency you need to worry about is the actual propagation delay of the network.

a) How much time has progressed by point B?

b) How much time has progressed between points C and D?

c) How much time has progressed between points E and F?

4. If the sender shares its network with other clients whose traffic traverses the same IP routers, give one explanation for why point D is higher than point B?