QUESTION 2 : TCP and Congestion Control (10 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.

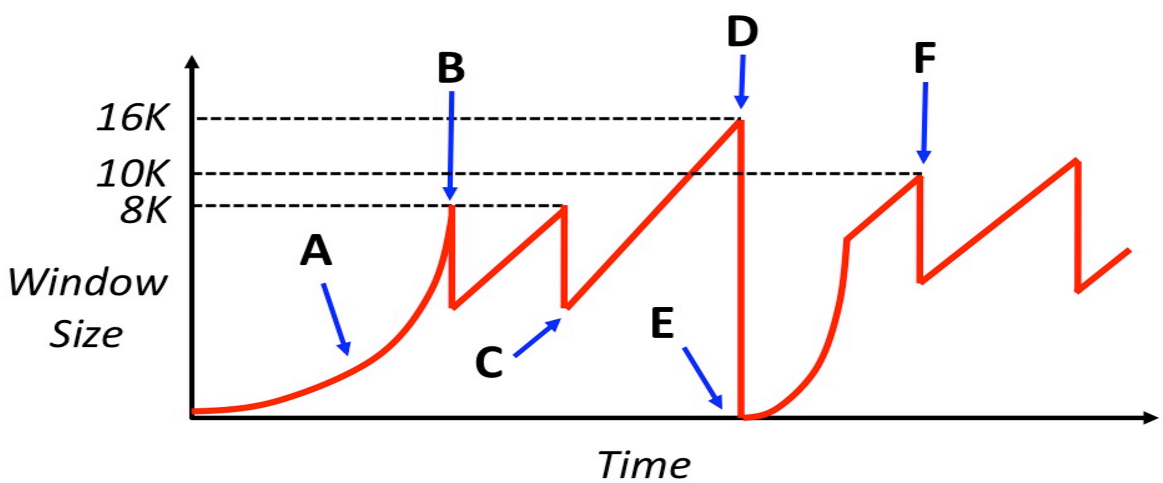

1. [5 points – 1 for each] The window size of the TCP sender decreases at several points in the graph, including those marked by B and D.

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

Triple duplicate ACK

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

No. It could be due to reordering due to queuing or asymmetric paths.

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

Timeout

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

No. Congestion in either direction could cause RTT > RTO (retrans. timeout).

(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”)

LESS

2. [1 points] 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?)

This "slow-start" period quickly discovers the maximum acceptable throughput that the path supports – otherwise, AI (additive increase) could take too long (each a full RTT).

3. [3 points – 1 for each] 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?

It depends if you interpret B as before or after you received the triple duplicate ACK, so we will accept both answers:

1 RTT (TCP handshake) + 3-4 RTT in slow-start (1, 2, 4, (8) MSS)
= 4 or 5 RTT = 400 or 500 ms

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

4 MSS to 16 MSS = 12 periods of RTT = 1.2 s

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

First: slow start to 8K window size (1->2, 2->4, 4->8 MSS), then AI from 8 to 10 MSS window size (8->9->(10) MSS). Again, similar question about "where" F exists as question (a), so we will accept either 5 or 6 RTT = 500 or 600 ms.

4. [1 point] 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?

Changing cross-traffic by other concurrent senders across same routers.

QUESTION 3: Distance-Vector Routing (9 points)

The CS department at Princeton bought new Sun Fire V210 servers. They decided to run a distance-vector protocol for routing between these servers (even though it is a rather small network). They are currently configured as the picture below, with respective edge costs.

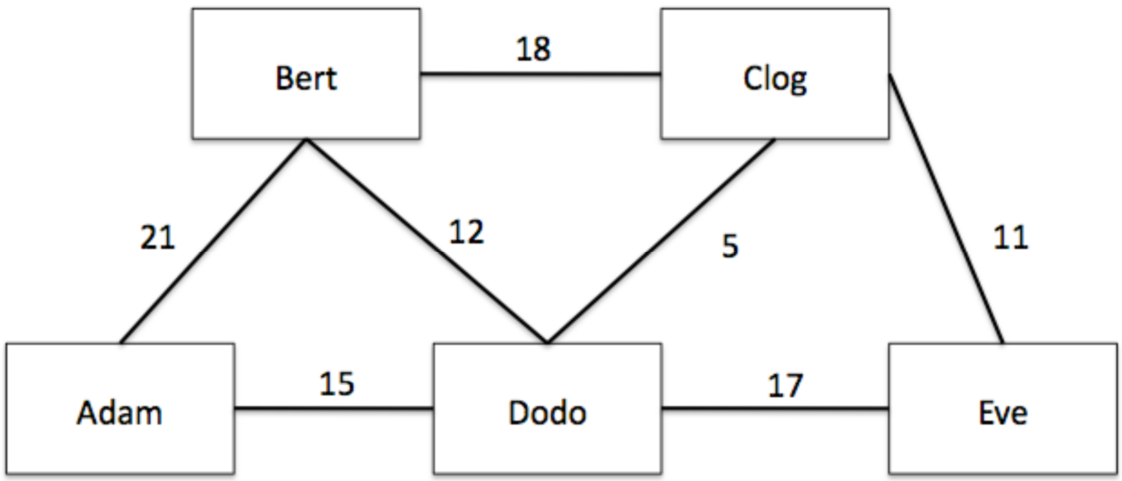

The CS staff asked for your help. Write down each step of building the distance-vector routing table for 'Eve' so they can compare it to the output of their implementation. You can use abbreviations e.g., 'A' for Adam and 'E' for Eve.

The initial routing table at node A is:

Destination	***Cost***	***Next Hop***
B	21	B
C	∞	---
D	15	D
E	∞	---

(a) Show the initial routing table of node E:

Destination	***Cost***	***Next Hop***
A	∞	–
B	∞	–
C	11	C
D	17	D

(b) Show the routing table of node E after one iteration of the algorithm:

Destination	*Cost*	*Next Hop*	
A	17+15 = 32	D	
B	11+18=29 / 17+12=29	C / D	either one
C	11	C	
D	11+5 = 16	C	

(c) Show the routing table of node E after two iterations of the algorithm:

Destination	*Cost*	*Next Hop*	
A	16+15 = 31	C	via D
B	16+12 = 28	C	via D
C	11	C	
D	16	C	

(d) In some failure situations, the administrator notices that it takes an exceptionally long time for the routing protocol to stabilize in this network.

(i) What problem with the distance vector protocol is the cause?

Count - to - infinity problem

(ii) The administrator is told that BGP does not suffer from this problem. What prevents BGP from having this problem?

BGP keeps track of the route, allowing for easy loop detection.