CSE 473 – Introduction to Computer Networks *Jon Turner*

Exam 2 Review Questions

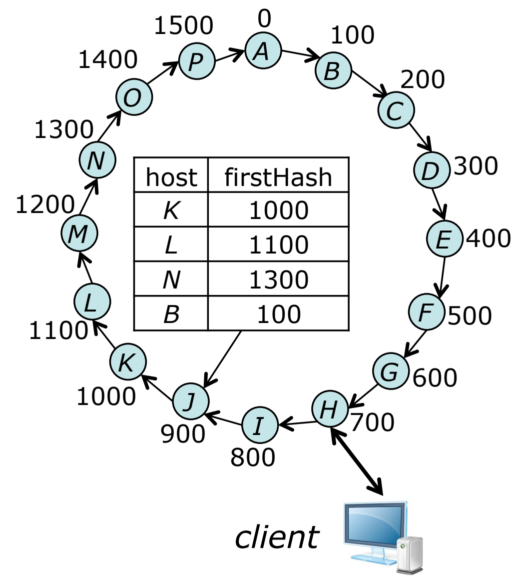

1. The diagram at right shows a DHT with 16 nodes. Each node is labeled with the first value in its range values (so for example, *B* is responsible for hash values 100-199. The routing table for node *J* is shown in the figure. Note that *J* has routes to the node that is 1 hop away, the one that is 2 hops away, the one that is 4 hops away and the one that is 8 hops away. Assume that all nodes have routing tables that are configured similarly.

   Suppose the client shown in the diagram sends a get request to node *H* with a key string of "flapjack", and that *hash*("flapjack")=513. List the servers through which this request would pass, assuming that the key string does not appear in any node's cache.

   What servers would the request pass through if the key string appears in node *M*'s cache?

   What servers would it pass through if the key string appears in node *D's* cache?

   Suppose that "flapjack" is requested frequently. Specifically, each DHT node receives a get request for "flapjack" about once per second. If the system is operated without caches, how many requests per second must the "responsible server" process?

   If caching is enabled, and each cache entry expires 60 seconds after being placed in the cache, approximately how often does the responsible server receive a get request from another server? (Hint: how many other servers send directly to the "responsible server"?)

2. Consider a pipelined, reliable transport protocol that uses go-back-*N* with cumulative acknowledgments. Assume that timeouts trigger retransmissions (duplicate ACKs do not) and that the receiver does not maintain any receive buffer. If the one-way delay between the sender and receiver is 50 ms and every packet is 10,000 bits long, how big must the window be to allow the sender to send at a steady rate of 1 Gb/s under ideal conditions?

   Suppose that approximately one packet in 100,000 is lost. If the sender uses a timeout of 500 ms and a window size of 20,000 packets, how often does sender experience a timeout? How many packets will it retransmit when a time out occurs?

   Assume that after the connection starts at time 0, the 100,000-th packet (call it *p*) is lost. At what time was *p* sent by the sender?

   At what time does the sender re-transmit *p*?

   What happens to the packets sent between the time *p* is sent the first time and the time it is retransmitted?

   Estimate the throughput for this connection, assuming one packet in 100,000 is lost.

3. The diagram at right shows a TCP segment being sent from host *A* to host *B* and an ACK being returned. The numbers on the arrows are the sequence numbers of the data segments and the ACK numbers. Suppose that after receiving the ACK with ack number 20, *A* sends packets with sequence numbers 20, 30, 40, 50, 60, 70, 80, 90 and 100. Some time later, it receives ACKs with sequence numbers 40, 40, 60, 60, 60, 60, 60. (Assume that *A* sends no additional data segments in the meantime.) Complete the diagram in a way that is consistent with the given information and what you know about the way TCP behaves.

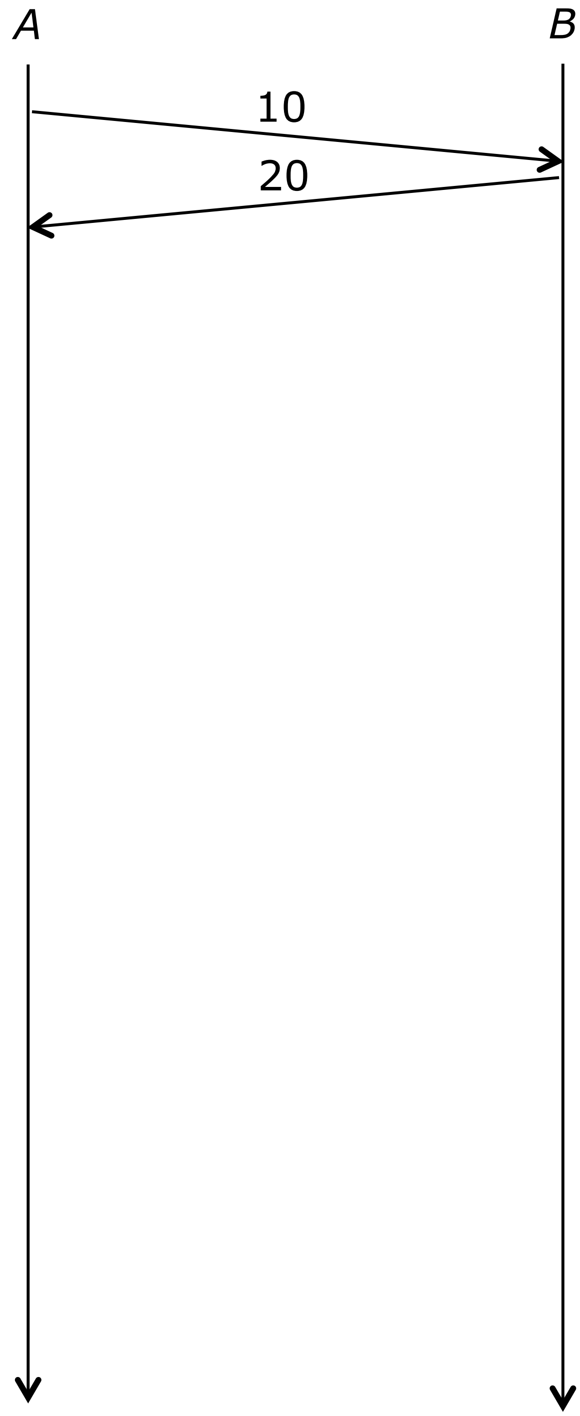

   What sequence number would you expect to see in the next packet sent by *A*?

   What ACK number would you expect in the next ACK? You may assume that all packets sent by *A* carry 10 bytes of data.

4. The diagram at right shows two TCP senders at left and the corresponding receivers at right. Both senders use TCP Reno. Assume that the MSS is 1 KB, that the one-way propagation delay for both connections is 50 ms and that the link joining the two routers has a bandwidth of 2 Mb/s. Let $cwnd_1$ and $cwnd_2$ be the values of the senders' congestion windows and assume that $cwnd_1 = cwnd_2$. What is the smallest value of $cwnd_i$ for which the link joining the two routers stays busy all the time?

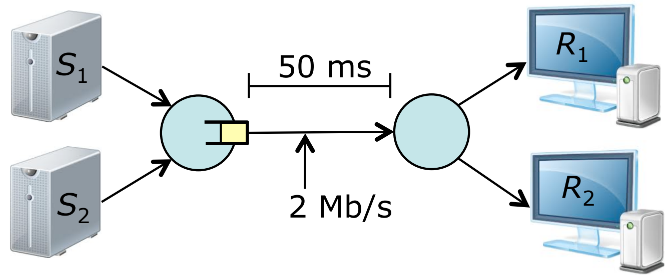

Assume that the link buffer overflows whenever $cwnd_1+cwnd_2 \geq 36$ KB and that at time 0, $cwnd_1=12$ KB and $cwnd_2=24$ KB. Approximately, what are the values of $cwnd_1$ and $cwnd_2$ one RTT later? Assume that all packet losses are detected by a triple duplicate ack.

How many RTTs pass before $cwnd_1+cwnd_2=36$ again? What are the values of $cwnd_1$ and $cwnd_2$ at this point?

Approximately, how many RTTs pass (in total) before $cwnd_2 - cwnd_1 < 2$ KB?

5. The diagram at right shows several autonomous systems and for each one, a prefix that it advertises to other ASs. Assuming that router *C* has IBGP sessions to all four gateway routers in AS1, list all the path vectors that it would receive describing paths to to other ASs. For each path vector identify the gateway router from which *C* receives the path vector, the vector itself, and the associated next-hop address.

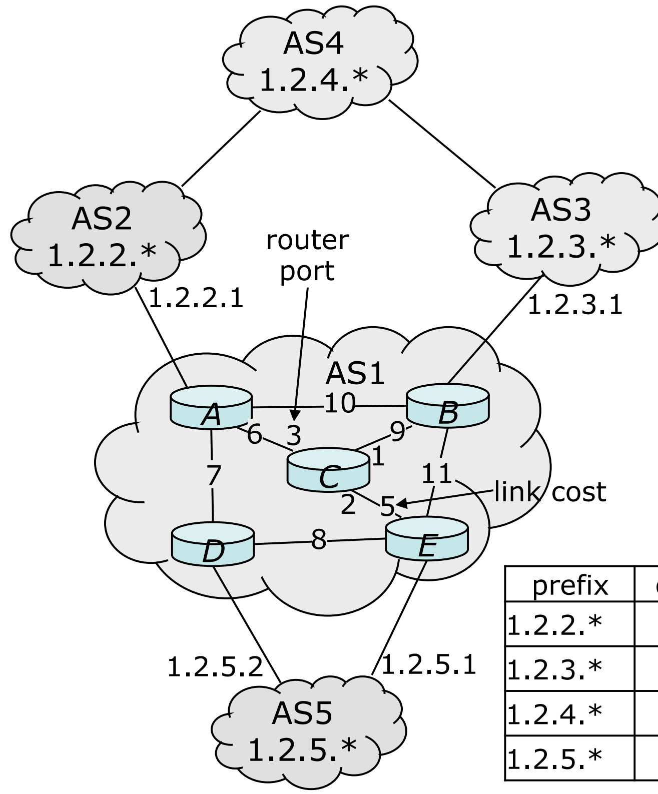

prefix	output
1.2.2.*	
1.2.3.*	
1.2.4.*	
1.2.5.*	

The table at the bottom right of the figure represents a portion of the routing table at *C*. Complete the table, as it would be completed, assuming that AS1 uses OSPF in addition to BGP. For this part assume there are no policy constraints.

Suppose that AS1 is not intended to carry transit traffic. Explain how AS1 can use BGP to enable it to receive traffic from other ASs, while not allowing them to send traffic through it.

address	port

6. The diagram at right shows a switched Ethernet LAN. Assume that the routing tables are empty initially. Show the contents of the routing table at switch v, after host A sends a packet to host B and host B sends a response. You may assume that A and B denote the MAC addresses of the corresponding hosts

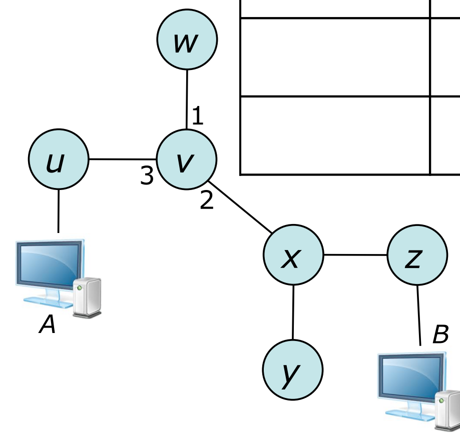

Suppose that at this point, a host at switch w sends a packet to some host other than B at switch z. Which switches in the network receive a copy of this packet?

If a host at switch y sends a packet to host A, which switches receive a copy?

The figure at right represents a switched Ethernet network, where the letters represent the switch identifiers and the numbers on the links represent costs used by the spanning tree algorithm. Assume that the switch identifiers are ordered alphabetically (so x comes before y, for example) and that all switches begin execution of the spanning tree algorithm at the same time and that all of them send an initial message to their neighbors, then receive the messages sent by their neighbors.

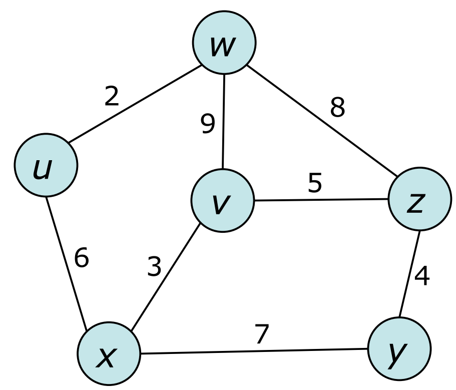

At this point, each switch has a tentative candidate for the spanning tree root. Complete the table below to show the candidate root for each switch.

u	v	w	x	y	z

In the diagram, highlight the edges that are included in the spanning tree when the algorithm completes (make the tree edges "heavier" by drawing a thicker line).

Suppose that all the links in the network operate at 1 Gb/s. Suppose we have a TCP connection between hosts on switches v and w, another between hosts on switches y and z and a third between hosts on switches u and x. If all hosts are trying to send data as fast as they can, what throughput would you expect each to achieve?

What throughput could they achieve if the routing was not limited to the spanning tree?

7. Consider a situation in which 1000 clients are trying to download a 10 MB file from a server. If the server has a 100 Mb/s access link and the clients have access links with a downstream rate of 2 Mb/s, how long does it take to download the file to all clients, under ideal conditions (you may ignore the time to establish a TCP connection to the server).

   Now, consider the peer-to-peer situation, in which there is no server and one peer holds the file to be distributed. Assuming that the upstream rate from each peer is 1 Mb/s and the downstream rate is 2 Mb/s, how long does it take to distribute the file to all peers?

8. (5 points) The diagram at left below shows the state of the sending side of a sliding window protocol with a window size of 4 and the selective repeat feature. The array represents the send buffer and each pair in the buffer represents a packet and its sequence number, together with the time at which it is scheduled to be retransmitted. (so, for example, the pair (*p*3,14) denotes a packet with sequence number 3, which is to be retransmitted at time 14).

Suppose that at time 7, the application passes us a new payload to be sent, at time 8, we receive an ack with sequence number 3, and that at time 9, we receive an ack with sequence number 2. Show the state of the sender at time 10, in the right-hand diagram. Assume that the timeout used for retransmitting packets is 20 time units.

	0			0	
	1			1	
sendBase=2	2	(p2,10)	*sendBase*=_____	2	
	3	(p3,14)		3	
	4	(p4,22)		4	
nextSeqNum=5	5		*nextSeqNum*=_____	5	
	6			6	
time=6	7		*time*=_____	7	

If no additional payload or ack is received before time 25, what is the next thing that should happen and how does it affect the sender's state?

9. Consider two TCP flows that pass through a common link served by a single queue. The first flow uses TCP Tahoe, the second uses TCP Reno. Assume that both flows have the same *RTT.* Suppose that at time t_0, the queue is full, causing both flows to lose a packet. Suppose the senders for both flows detect the lost packet at time *t*. If the first flow has a congestion window size equal to 100**MSS* and the second has a congestion window size of 200**MSS* at time *t*, approximately how many bytes will the first flow send during the next RTT following *t?*

   How many will the second flow send?

   At time *t*+50**RTT*, approximately what will be the size of the congestion window for the first flow?

   What will be the size of the congestion window for the second flow?

10. The diagram below represents a network with the numbers on the edges representing link costs. If the network routes are computed using OSPF, which links would never be used to reach a host Z connected to router *g*? You may simply mark these edges with an X on the diagram. Assume that the topology is stable, that the link costs do not change and that *g* advertises a route to Z.

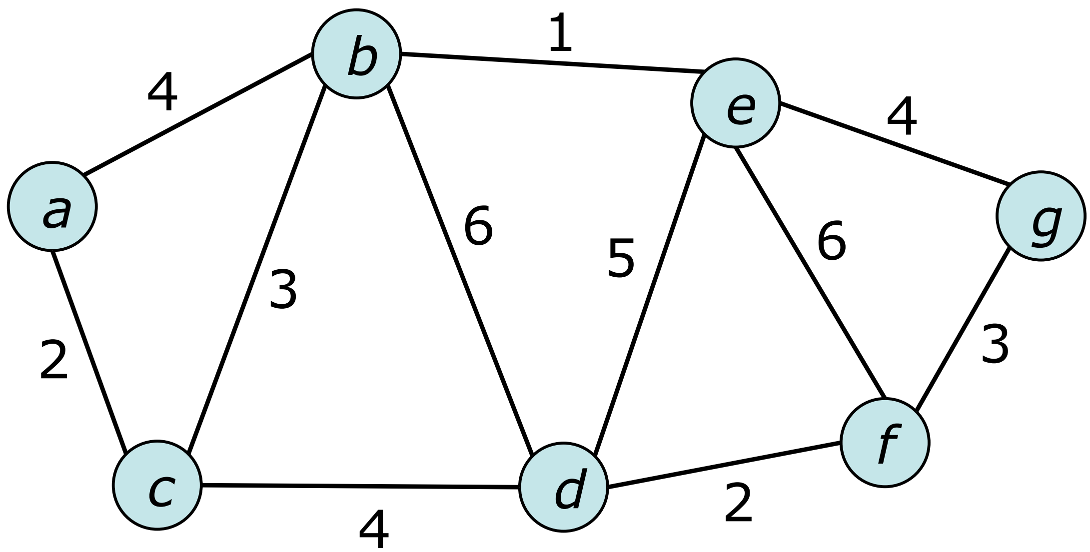

Suppose that router *a* sends an LSA for one of its local subnets. Suppose a copy of this LSA reaches router *e* after passing through *k* other routers. What is the smallest possible value for *k*?

What is the largest possible value for *k*?

Suppose that router *e* receives *n* copies of this LSA, all with the same sequence number. What is the largest possible value for *n*?