Name_______________ Student ID____________________ Department/Year___________

Final Examination

Introduction to Computer Networks
Class#: 901 E31110
Fall 2019

10:20-12:10 Thursday
January 9, 2020

Prohibited

1. You are not allowed to write down the answers using pencils. Use only black- or blue-inked pens.
2. You are not allowed to read books or any references not on the question sheets.
3. You are not allowed to use calculators or electronic devices in any form.
4. You are not allowed to use extra sheets of papers.
5. You are not allowed to have any oral, visual, gesture exchange about the exam questions or answers during the exam.

Cautions

1. Check if you get 20 pages (including this title page), 11 problem sets.
2. Write your **name in Chinese**, student ID, and department/year down on top of the first page.
3. There are in total 100 points to earn. You have 100 minutes to answer the questions. Skim through all questions and start from the questions you are more confident with.
4. Use only English to answer the questions. Misspelling and grammar errors will be tolerated, but you want to make sure with those errors your answers will still make sense.
5. If you have any extra-exam emergency or problem regarding the exam questions, raise your hand quietly. The exam administrator will approach you and deal with the problem.

1. (Transport) Suppose the change of the congestion window size of a TCP connection is as depicted below.

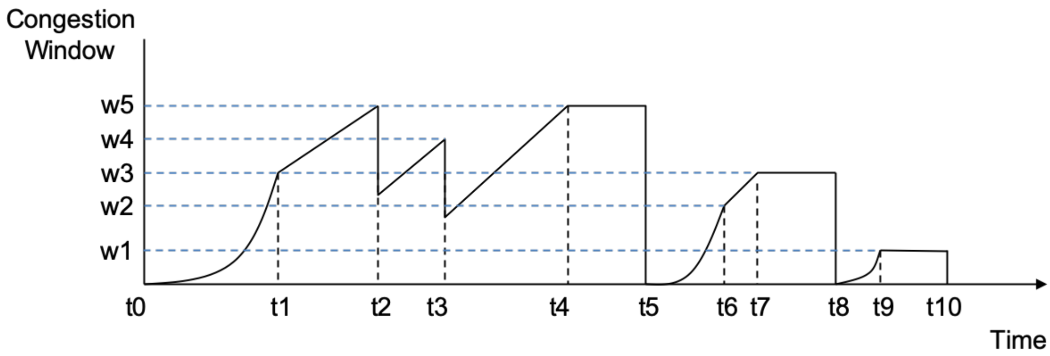

(a) State the time periods that the TCP connection is in the slow start phase. State the time periods in the congestion avoidance phase. (5%)

(b) State the ssthresh during time periods t0-t1, t2-t3, t3-t4, t7-t8, and t8-t9. (5%)

(c) State the TCP version, Tahoe or Reno, this connection is using and how you tell. (3%)

Sample Solution:

(a) Slow start: t0-t1, t5-t6, t8-t10

Congestion avoidance: t1-t5, t6-t8

(b) w3, w5/2, w4/2, w2(or w5/2) , w3/2

(c) TCP Reno. The window size does drop at t2, t3, t5, t8. At t2 and t3 in particular, the window size does not drop all the way down to 1. This indicates TCP Reno behavior, whereas in TCP Tahoe, the window size always drops to 1 and slow-starts again.

2. (Transport) Suppose you work for an Internet radio channel. The company has been transferring audio by UDP until its ISP is concerned about the high-volume audio stream it is sending. The ISP sends the company a notice. If the company would not send the audio stream about the volume everyone else surfing the WWW, the ISP will terminate the contract. As the Internet engineer of the radio channel, you are asked to provide a solution - a transport layer service that does not guarantee reliability but sends traffic at the rate of an average TCP flow.

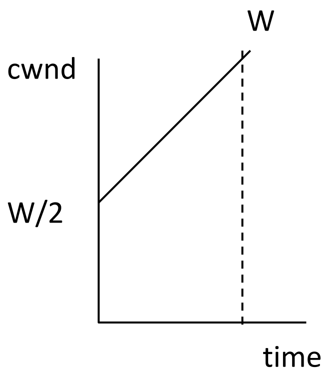

Assume an audio transfer is long and it will be sufficient to transfer at the rate of a long-term TCP connection, which is typically in the congestion avoidance state (the steady state). Assume also (a bit unrealistically) the packet drops are very sporadic and all are discovered by duplicate acknowledgements. The figure above depicts the congestion window size (in number of packets) within an average episode of additive increase in the steady state. Try if you can derive the average TCP throughput: MSS/RTT * $(3/2p)^{1/2}$, where MSS is the packet size, RTT is the round-trip time, and p is the packet drop rate.

(a) Derive TCP throughput in terms of W, MSS, and RTT. (4%)
(b) Identify the relationship between W and p. (4%)
(c) Rewrite TCP throughput in terms of p, MSS, and RTT. (2%)

Sample Solution:

It is assumed that there's only dup-ack drop. The cwnd increases always linearly.
Therefore, the average TCP throughput is half of the min and max of the average episode:
TCP throughput = (W/2*RTT+W/RTT)*MSS/2 = 3/4 W*MSS/RTT ---- (a)

---- (b)

---- (c)

3. (Routing) Based on the algorithm of Distance Vector (DV) routing, fill in the DV tables below. Here we assume the sending, receiving, and computation of DVs are synchronous in an iteration.

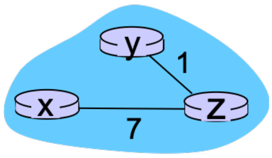

(a) Initialize the DV tables at node x, y and z at t0. (2%)
(b) Develop the DV tables at node x, y, and z at t1. (2%)
(c) Develop the DV tables at node x, y, and z at t2. (2%)

Sample Solution:

node x table

t_0 (from \ cost to)	x	y	z
x	0	∞	7
y	∞	∞	∞
z	∞	∞	∞

t_1 (from \ cost to)	x	y	z
x	0	8	7
y	∞	∞	∞
z	7	1	0

t_2 (from \ cost to)	x	y	z
x	0	8	7
y	∞	∞	∞
z	7	1	0

node y table

t_0 (from \ cost to)	x	y	z
x	∞	∞	∞
y	∞	0	1
z	∞	∞	∞

t_1 (from \ cost to)	x	y	z
x	∞	∞	∞
y	8	0	1
z	7	1	0

t_2 (from \ cost to)	x	y	z
x	∞	∞	∞
y	8	0	1
z	7	1	0

node z table

t_0 (from \ cost to)	x	y	z
x	∞	∞	∞
y	∞	∞	∞
z	7	1	0

t_1 (from \ cost to)	x	y	z
x	0	∞	7
y	∞	0	1
z	7	1	0

t_2 (from \ cost to)	x	y	z
x	0	8	7
y	8	0	1
z	7	1	0

time: t_0 → t_1 → t_2

Note that x's DV goes only to z, y's DV goes only to z, and z's DV goes both to x and y.

4. (Routing) Continue from Problem set #3. Link y-z is suddenly broken. I.e., the link cost c(y, z) is now ∞. Fill in the DV tables below.

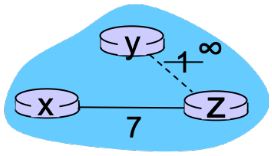

(a) Update the DV tables at node x, y and z at t0, right after the link cost change. (3%)
(b) Develop the DV tables at node x, y, and z at t1. (3%)
(c) Develop the DV tables at node x, y, and z at t2. (3%)
(d) Will the DV table converge? If yes, how many iterations will it take? (4%)

Sample Solution:

(a)-(c)

node x table

t_0 (cost to)

from	x	y	z
x	0	8	7
y	∞	∞	∞
z	7	1	0

t_1 (cost to)

from	x	y	z
x	0	22	7
y	∞	∞	∞
z	7	15	0

t_2 (cost to)

from	x	y	z
x	0	22	7
y	∞	∞	∞
z	7	15	0

node y table

t_0 (cost to)

from	x	y	z
x	∞	∞	∞
y	∞	0	∞
z	7	1	0

t_1 (cost to)

from	x	y	z
x	∞	∞	∞
y	∞	0	∞
z	7	1	0

t_2 (cost to)

from	x	y	z
x	∞	∞	∞
y	∞	0	∞
z	7	1	0

node z table

t_0 (cost to)

from	x	y	z
x	0	8	7
y	8	0	1
z	7	15	0

t_1 (cost to)

from	x	y	z
x	0	8	7
y	8	0	1
z	7	15	0

t_2 (cost to)

from	x	y	z
x	0	22	7
y	8	0	1
z	7	29	0

time

(d) No. $d_x(y)$ and $d_z(y)$ will continue to count (+7 each iteration) to infinity. The phenomenon is known as the "count to infinity problem".

Note that x's DV goes only to z, z's DV goes only to x, and y's DV does not go anywhere.

5. (Routing) The general rule is this – people pay for Internet access. By access, we meant the connectivity to send/receive traffic to/from the rest of the Internet. The customer networks pay their providers for Internet access. The smaller/younger ISPs also pay the larger/older ISPs for Internet access. I.e., the smaller ISP is the customer of the larger ISP.

   The figure below depicts a network of ISPs (provider networks) and their customer networks. Here, A is the largest ISP, B mid-sized, and C the smallest. B pays A for Internet access. C, being the smallest, pays B and A for Internet access. This network of networks is using the traditional BGP routing to discover inter-AS paths.

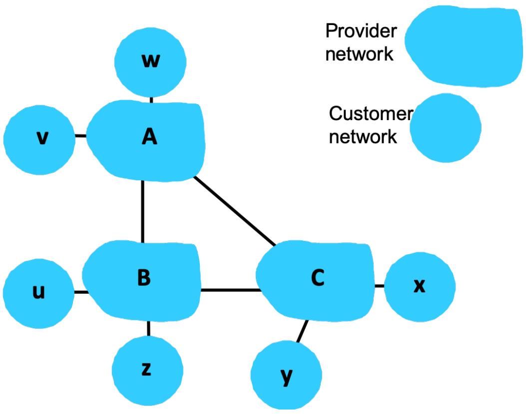

   (a) Suppose A receives AS path 'Cx' from C and 'BCx' from B. Which path you think A would select and advertise to v and w if there is no preference? Why? (3%)
   (b) Suppose A receives AS path 'Cx' from C and 'BCx' from B. Which path you think A would select and advertise to v and w if A charges B more than A charges C for the same amount of traffic? Why? (3%)
   (c) Suppose A receives AS path 'Cx' from C and 'BCx' from B. Do you think A would advertise 'ACx' to B if 'Cx' is selected? Why? (3%)
   (d) Suppose A receives AS path 'Cx' from C and 'BCx' from B. Do you think A would advertise 'ABCx' to C if 'BCx' is selected? Why? (3%)

Sample Solution:

Take your pick and justify. Below are just one possible pick/justification. As the justification is way more important than the answer, if the justification does not make sense, no credits will be rewarded.

(a) 'Cx'. Because it's a shorter AS path
(b) 'BCx'. Because A gets to earn more $$

(c) Yes. Because A as a provider should try to advertise further to allow path discovery
(d) No. Because C is already in the AS path. C wouldn’t select the ‘ABCx’ path anyway.

6. (Forwarding) In traditional BGP routing, an AS typically selects a path for a destination subnet before pass it further. It is not straightforward to set up multiple AS paths to achieve multiple purposes at the same time. SDN being flexible allows finer granularity of path setting. The setting depicted below allows packets from h5 or h6 to travel through ISP B to C before arriving at the destinations.

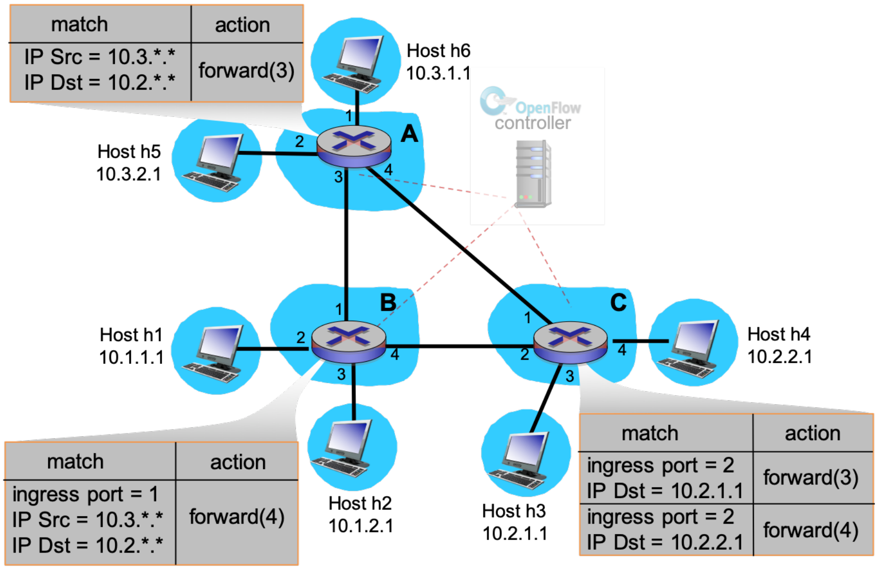

Set the flow table entries in the border gateway router of A, B, and C so that (1) h6 can send video streams to h4 directly through a shorter AS path 'AC', in the meantime (2) h5 can send Web responses to h3 through AS path 'ABC' which might generate a higher revenue. (10%)

Sample Solution:

At A's router:

Match	action
IP Src = 10.3.1.1 1. IP Dst = 10.2.2.1	forward(4)
IP Src = 10.3.2.1 2. IP Dst = 10.2.1.1	forward(3)

At B's router:

Match	action
Ingress = port 1 IP Src = 10.3.2.1 1. IP Dst = 10.2.1.1	forward(4)

At C's router:

Match	action
Ingress = port 1 1. IP Dst = 10.2.2.1	forward(4)
Ingress = port 2 2. IP Dst = 10.2.1.1	forward(3)

7. (Forwarding) Below is a queuing system where incoming packets are classified into two queues, i.e., red and green. Each slot in the queue holds a packet. The number on the slot indicates the packet ID. The link transmits a packet from one of the queues at a time. Based on the principle of priority scheduling, round robin scheduling, and weighted fair queuing, list the packets in the order they will be transmitted onto the outgoing link.

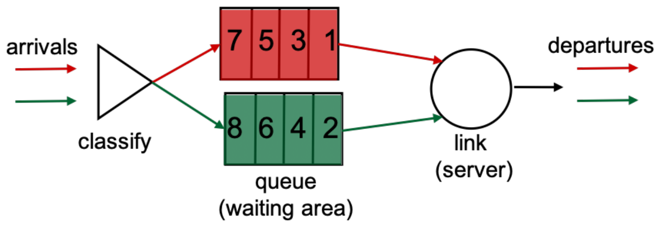

(a) Priority scheduling with the red queue being in high priority (1%)
(b) Round robin scheduling starting from the red queue (1%)
(c) Weighted fair queuing with a weight of 2:1, red queue to green queue and starting from the red queue (1%)

Sample Solution:

(a) 1->3->5->7->2->4->6->8
(b) 1->2->3->4->5->6->7->8
(c) 1->3->2->5->7->4->6->8

8. (Forwarding) Continue from Problem set #7. A traffic intersection is also a queuing system. In the figure below, we see 4 queues in a 4-way interaction. When the north and south (NS) bound lights are green, cars in queue 1 and 2 are served in parallel (i.e., full duplex). During this time, the west and east (WE) bound traffic lights are red and the cars in queue 3 and 4 wait.

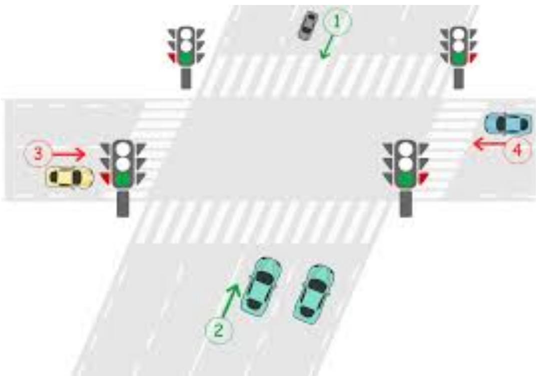

Consider the following traffic light settings, state which of the 3 scheduling mechanisms, priority scheduling, round robin scheduling, and weighted fair queuing, the intersection is implementing and why you come to the finding.

(a) The NS bound lights repeat – on green for 3 mins and red for 3 mins. (2%)
(b) The NS bound lights repeat – on green for 3 mins and red for 1 min. (2%)
(c) The NS bound lights repeat – on green forever. (2%)

Sample Solution:

(a) Round robin or 1:1 weighted fair queuing, where queue 1 and 2 and queue 3 and 4 rotate and equally share the intersection.
(b) 3:1 weighted fair queuing, where queue 1 and 2 gets 3 times more cars crossing the intersection.
(c) Priority scheduling, where the priority is given to queue 1 and 2.

9. (Link) Before being able to send a data frame, a node uses ARP to find the MAC address of the receiving node. Consider the following familiar scenario where the destination is not on the same subnet. Suppose that the ARP tables on all nodes are empty and the data is going from A to B, through router R. Complete the blanks in the ARP and data frames involved in the transmission below. (6%)

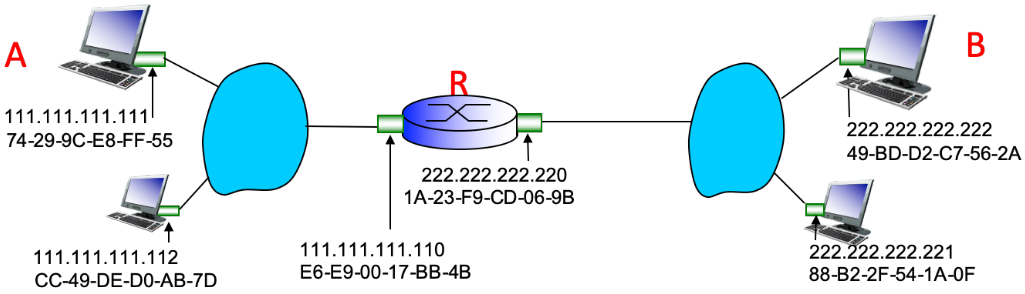

Sample Solution:

A R B

ARP Query →

Source Address	Destination Address	Who has?
74-29-9C-E8-FF-55	FF-FF-FF-FF-FF-FF	111.111.111.110

← ARP Reply

Source Address	Destination Address	Mapping
E6-E9-00-17-BB-4B	74-29-9C-E8-FF-55	111.111.111.110 -> E6-E9-00-17-BB-4B

Data Frame →

Source Address	Destination Address
74-29-9C-E8-FF-55	E6-E9-00-17-BB-4B

ARP Query →

Source Address	Destination Address	Who has?
1A-23-F9-CD-06-9B	FF-FF-FF-FF-FF-FF	222.222.222.222

← ARP Reply

Source Address	Destination Address	Mapping
49-BD-D2-C7-56-2A	1A-23-F9-CD-06-9B	222.222.222.222-> 49-BD-D2-C7-56-2A

Data Frame →

Source Address	Destination Address
1A-23-F9-CD-06-9B	49-BD-D2-C7-56-2A

10. (Link) Take the network from Problem set #9. Replace the router R with a link layer switch S. The data destination B is now on the same subnet. Suppose that the ARP tables on all nodes are empty and the data is going from A to B, through switch S.

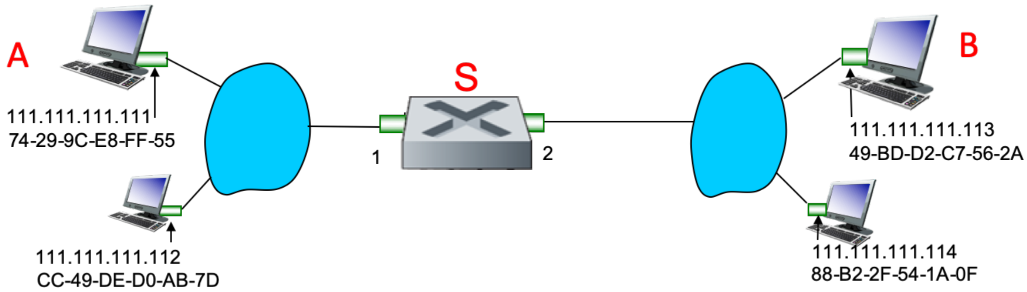

(a) A sends first an ARP query to find out B's MAC address. Is there any new entry being added to S's switch table? If yes, tell the destination address and outgoing interface of the entry. (1%)
(b) Continue from (a). Does S find the ARP query's destination address matching an entry in its switch table? If yes, tell the destination address and outgoing interface of the entry. (1%)
(c) B sends next an ARP reply back to A. Is there any new entry being added to S's switch table? If yes, tell the destination address and outgoing interface of the entry. (1%)
(d) Continue from (c). Does S find the ARP reply's destination address matching an entry in its switch table? If yes, tell the destination address and outgoing interface of the entry. (1%)
(e) A sends finally a data frame to B. Is there any new entry being added to S's switch table? If yes, tell the destination address and outgoing interface of the entry. (1%)
(f) Continue from (e). Does S find the data frame's destination address matching an entry in its switch table? If yes, tell the destination address and outgoing interface of the entry. (1%)

Sample Solution:

(a) Yes. Adding (74-29-9C-E8-FF-55, 1) or (A, 1)
(b) No.
(c) Yes. Adding (49-BD-D2-C7-56-2A, 2) or (B, 2)
(d) Yes. Matching (74-29-9C-E8-FF-55, 1)
(e) No.
(f) Yes. Matching (49-BD-D2-C7-56-2A, 2)

11. (Link) We have classified MAC protocols into 3 categories -- partitioning, random access, and taking turns. Consider the following MAC protocols. Tell which category each of them belongs to and why. (15%)

(a) CDMA
(b) CSMA/CD
(c) Aloha
(d) Token Ring
(e) TDMA

Sample Solution:

(a) Partitioning. A chip sequence is reserved for a device and only this device uses the sequence to encode the data bits.
(b) Random access. A device running CSMA/CD transmits as long as the channel is not busy. There is a chance that two devices on the CSMA/CD link collide.
(c) Random access. An Aloha device will just transmit and see if there's collision.
(d) Taking turns. A device on the token ring transmits whenever it holds the token. Otherwise, it stays quiet. It's not quite partitioning (no idle resource). Not quite random access (no collision)
(e) Partitioning. A time slot is reserved/dedicated to a device in TDMA