Homework 2 assignment for ECE671

Posted: 02/16/21

Due: 02/23/21

Note: *In all written assignments, please show as much of your work as you can. Even if you get a wrong answer, you can get partial credit if you show your work.* If you make a mistake, it will also help the grader show you where you made a mistake.

Problem 1 (5 Points):

Suppose Client A initiates a HTTP session with web server S. At about the same time, Client B also initiates a HTTP session with web server S. Provide possible source and destination port numbers for:

a) The segments sent from A to S.
b) The segments sent from B to S.
c) The segments sent from S to A.
d) The segments sent from S to B.
e) If A and B are different hosts, is it possible that the source port number in the segments from A to S is the same as that from B to S?

Solution:

a) A → S (S)467 (D)80 **b)** B → S (S)513 (D)80 **c)** S → A (S)80 (D)467 **d)** S → B (S)80 (D)513 **e)** yes, because A and B will have different IP addresses and thus the five tuples will be different and allow the server to differentiate the streams.

Problem 2 (25 Points):

For this problem, you should familiarize yourself with Figure 1 first. Now assume that in the network shown in Figure 1 two parallel TCP transmissions are performed. *TCP1* is a transmission between Source A and Sink A that uses *TCP Tahoe*. *TCP2* is a transmission between Source B and Sink B that uses *TCP Reno.* Initial *ssthresh* for both TCP transmissions is set to 32. In this specific scenario, no additional delay through forwarding is introduced. Thus, the RTT is only composed of the sums of the delay indicated on each link, times two.

a. For the *TCP 1* transmission, draw the resulting congestion window, assuming that a packet loss (triple duplicate ACKs) is detected at time t=900ms in Figure 2.
b. For the *TCP 2* transmission, draw the resulting congestion window, assuming that a packet loss (triple duplicate ACKs) is detected at time t=650ms in Figure 2.
c. Describe the benefit of TCP Reno over TCP Tahoe.
d. In general, explain the purpose of the receiver-advertised window in TCP.
e. Assume a TCP sender transmits 5 TCP segments with respective sequence numbers 1200, 2400, 3600, 4800, 6000. The sender receives four acknowledgements with the following sequence numbers, 2400, 2400, 2400,

2400. Complete Figure 3 to show what TCP segments are exchanged between sender and receiver.

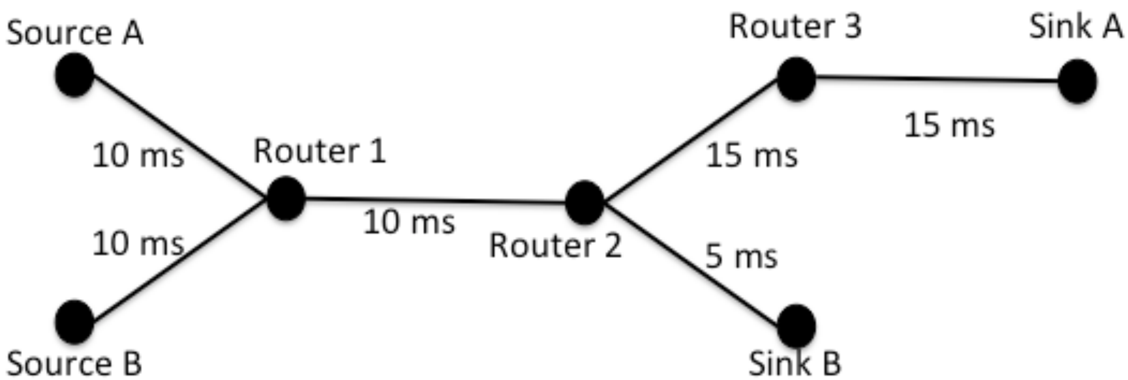

Figure 1. Network Layout

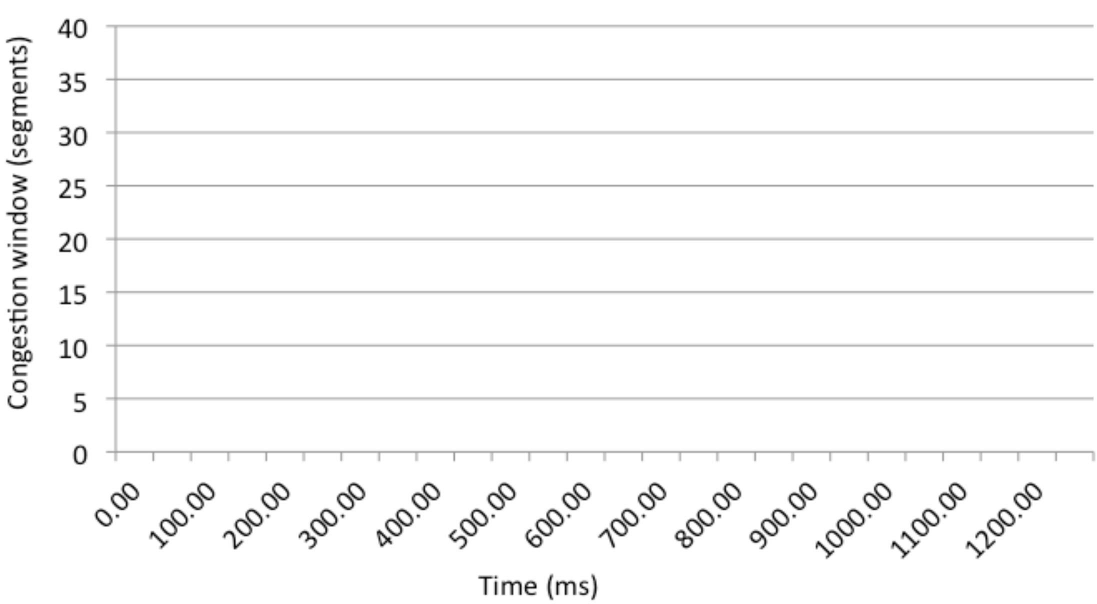

Figure 2. Congestion window of TCP 1 and TCP 2.

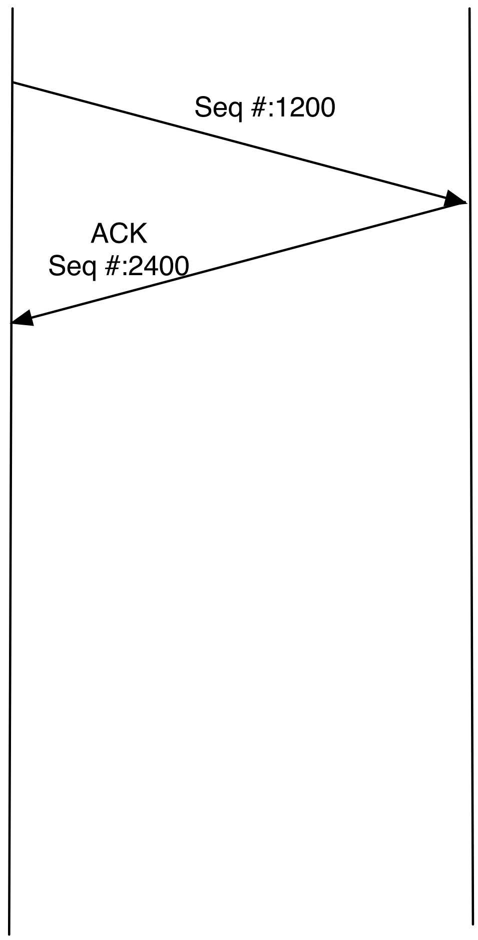

Figure 3. TCP segments exchange

Solution:
c) Reno recovers faster.

d) It allows the receiver to signal the sender how much unacknowledged data can be in flight.

Solution to a) and b)

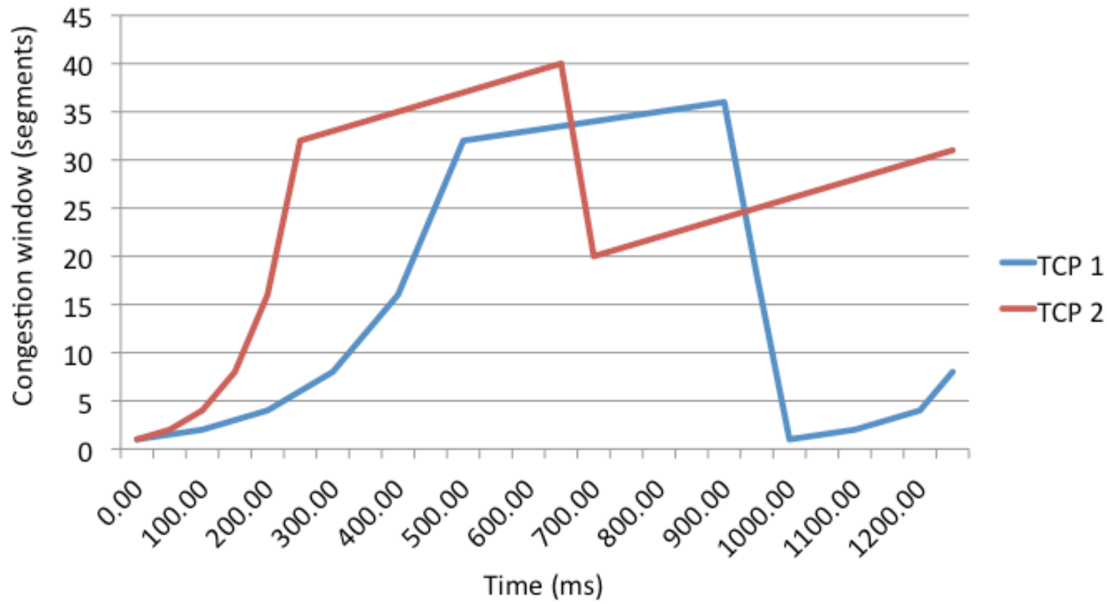

Solution to e)

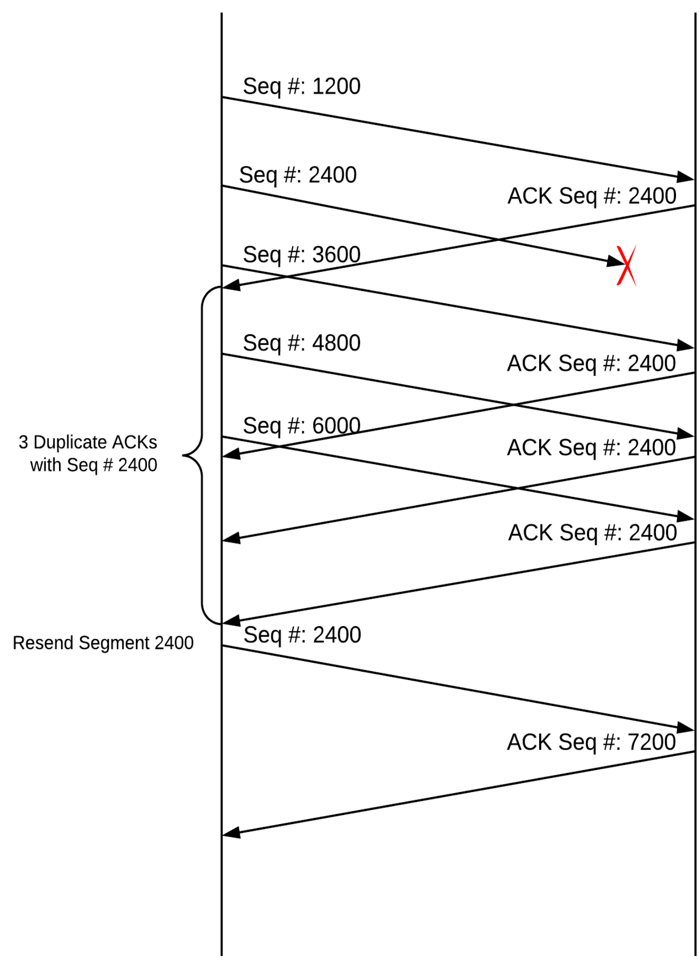

Problem 3 (20 Points):
Suppose two nodes, A and B, are attached to opposite ends of an 1200m cable, and that they each have one frame of 1,500 bits (including all headers and preambles) to send to each other. Both nodes attempt to transmit at time t=0. Suppose there are four repeaters between A and B, each inserting a 40-bit delay. Assume the transmission rate is 100 Mbp, and CSMA/CD with backoff intervals of multiples of 512 bits times is used. After the collision, A draws K=0 and B draws K=1 in the exponential backoff protocol. Ignore the jam signal in this case.

a. What is the one-way propagation delay (including repeater delays) between A and B in seconds? Assume the signal propagation speed is $2*10^8$ m/sec.
b. At what time (in seconds) is A's packet completely delivered at B?
c. Now suppose that only A has a packet to send and that the repeaters are replaced with switches. Suppose that each switch has a 20-bit processing delay in addition to a store-and-forward delay. At what time, in seconds, is A's packet delivered at B?

Solution:
a)
$$\frac{\mathbf{1200}\boldsymbol{m}}{\mathbf{2}*\mathbf{10^8}\boldsymbol{m/sec}}+\mathbf{4}*\frac{\mathbf{40}\boldsymbol{bits}}{\mathbf{100}*\mathbf{10^6}\boldsymbol{bps}}$$

$$=(\mathbf{6}*\mathbf{10^{-6}}+\mathbf{1.6}*\mathbf{10^{-6}})\boldsymbol{sec}$$
$$=\mathbf{7.6}\boldsymbol{\mu sec}$$

b)
First note, the transmission time of a single frame is given by 1500/(100Mbps) =15 microsecond, longer than the propagation delay of a bit.

- At time $t=0$, both A and B transmit.
- At time $t=7.6\mu\sec$, both A and B detect a collision, and then abort.
- At time $t=15.2\mu\sec$ last bit of B's aborted transmission arrives at A.
- The first bit of A's retransmission would arrive at B at time $t=22.8\mu\sec$ but B already starts retransmitting at time *t=20.32us*. To avoid a collision the backoff interval should be at least *760 bit times.*
- Assuming B would not transmit A's packet would be completely delivered time $$t=22.8\mu\sec+\frac{1500bits}{100\times10^6 bps}=37.8\mu\sec.$$

c) The line is divided into 5 segments by the switches, so the propagation delay between switches or between a switch and a host is given by $\frac{\mathbf{1200}\boldsymbol{m}/\mathbf{5}}{\mathbf{2}*\mathbf{10^8}\boldsymbol{m/sec}}=\mathbf{1.2}\boldsymbol{microsec}$.
The delay from Host A to the first switch is given by 15 microseconds (transmission delay), longer than propagation delay. Thus, the first switch will wait 16.4=15+1.2+0.2

(note, 0.2 is processing delay) till it is ready to send the frame to the second switch. Note that the store-and-forward delay at a switch is 15 microseconds. Similarly, each of the other 3 switches will wait for 16.4 microsecond before ready for transmitting the frame. The total delay is:
16.4*4 + 15+1.2 =81.8 microsecond.

Problem 4 (20 Points):
Consider the single switch VLAN in Figure 4, and assume an external router is connected to switch port 1. Assign IP addresses to the EE and CS hosts and router interface. Trace the steps taken at both the network layer and the link layer to transfer and IP datagram from an EE host to a CS host.

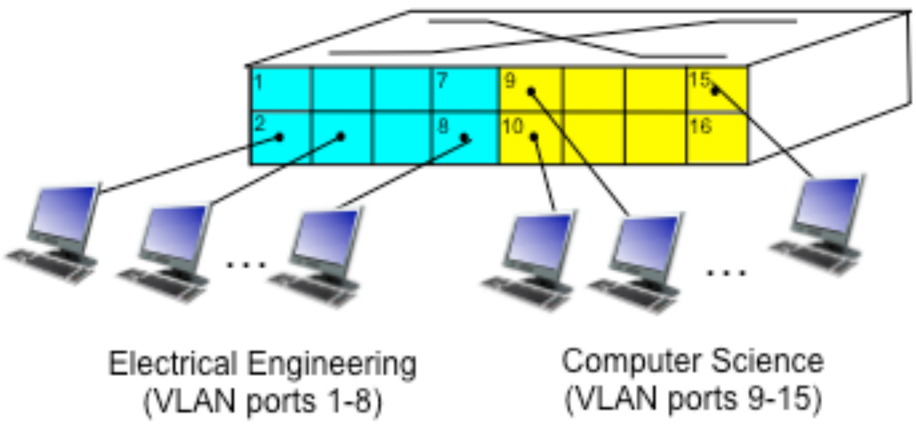

Figure 4

For this problem, keep the following in mind:

- Assume the IP addressing scheme for the EE and CS nodes follows the one indicated in **Error! Reference source not found.**.
- Assume that the EE VLAN has ID 11 and the CS VLAN has ID 12.
- The first figure at the following link gives you an idea of the logical setup for such a scenario: http://gcharriere.com/blog/?p=620

a. Assign IP addresses to the three nodes in the EE VLAN and to the three nodes in the CS VLAN. What're the subnet masks for these two VLAN, if each department should be capable to host 200 hosts?
b. Describe how the router interface has to be set up. What 802.1q VLAN ID will be added to a frame that comes from subnet 111.111.1/24? What 802.1q VLAN ID will be added to a frame that comes from subnet 111.111.2/24?
c. Suppose that host A in the EE department would like to send an IP datagram to host B in CS department. What would be the steps taken at both the network layer and the link layer?

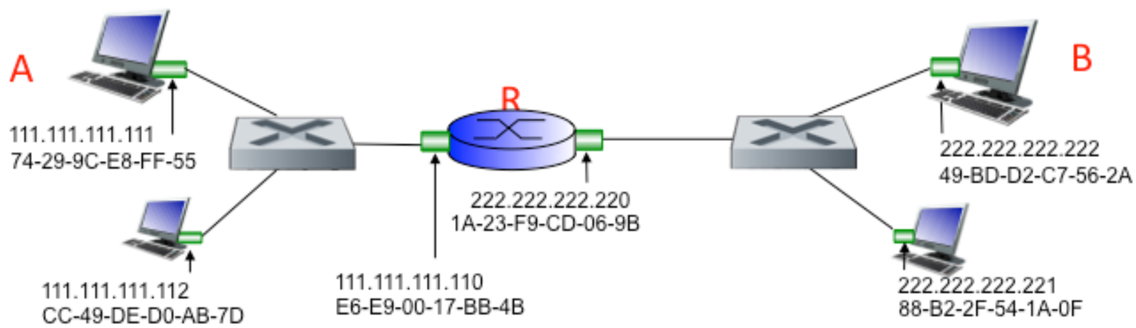

Figure 5

Solution:

a. The IP addresses for those three computers (from left to right) in EE department are: 111.111.1.1, 111.111.1.2, 111.111.1.3. The subnet mask is 111.111.1/24.
   The IP addresses for those three computers (from left to right) in CS department are: 111.111.2.1, 111.111.2.2, 111.111.2.3. The subnet mask is 111.111.2/24.
b. The router's interface card that connects to port 1 can be configured to contain two sub-interface IP addresses: 111.111.1.0 and 111.111.2.0. The first one is for the subnet of EE department, and the second one is for the subnet of CS department. Each IP address is associated with a VLAN ID. Suppose 111.111.1.0 is associated with VLAN 11, and 111.111.2.0 is associated with VLAN 12. This means that each frame that comes from subnet 111.111.1/24 will be added an 802.1q tag with VLAN ID 11, and each frame that comes from 111.111.2/24 will be added an 802.1q tag with VLAN ID 12.
c. Suppose that host A in EE department with IP address 111.111.1.1 would like to send an IP datagram to host B (111.111.2.1) in CS department. Host A first encapsulates the IP datagram (destined to 111.111.2.1) into a frame with a destination MAC address equal to the MAC address of the router's interface card that connects to port 1 of the switch. Once the router receives the frame, then it passes it up to IP layer, which decides that the IP datagram should be forwarded to subnet 111.111.2/24 via sub-interface 111.111.2.0. Then the router encapsulates the IP datagram into a frame and sends it to port 1. Note that this frame has an 802.1q tag VLAN ID 12. Once the switch receives the frame on port 1, it knows that this frame is destined to VLAN with ID 12, so the switch will send the frame to Host B which is in CS department. Once Host B receives this frame, it will remove the 802.1q tag.

Problem 5 (20 Points):

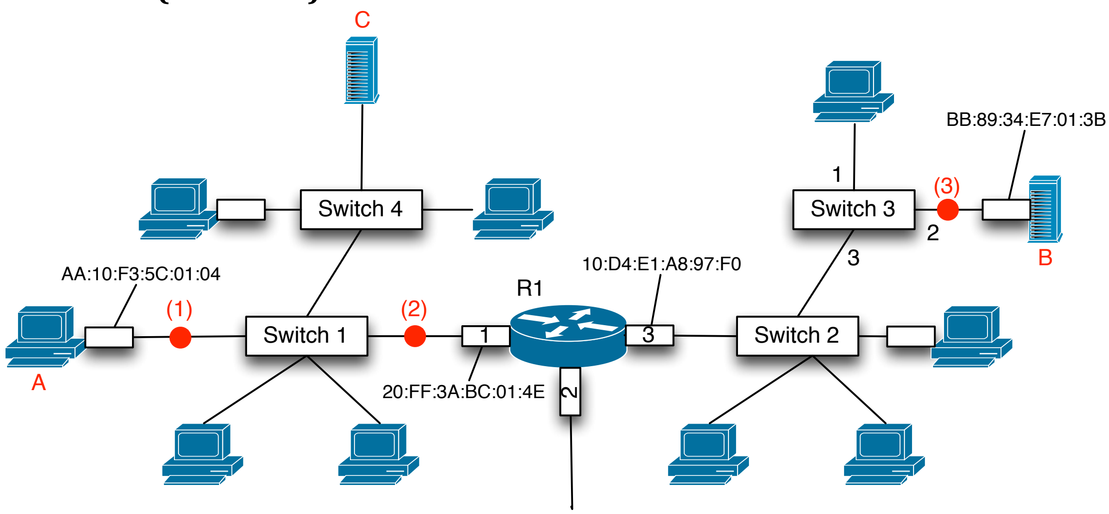

Consider the network shown above.

a. Consider an ARP request send by node B for node A. Whose IP-to MAC address translation is being queried?
b. What is the destination MAC address on the frame containing the ARP request?
c. After B receives the ARP reply, what is contained in switch 3's switching table?
d. Assign IP addresses and subnet masks to hosts A, B, C, and interfaces 1 and 3 of R1. *Note:* Each of the subnets should be able to host a maximum of 17 hosts.
e. Now consider the frame containing the B-to-A IP datagram. What are the MAC source and destination address of this frame and the IP source and destination addresses in the encapsulated IP datagram at points (1), (2), and (3).

Solution:

a. The right router interface "3".
b. The MAC address of the right router interface (10:D4:E1:A8:97:F0).
c. Switch 3 knows that B is reachable via interface 2 (as a result of the ARP request sent by B) and that the router R1 is reachable via interface 3 (as a result of the ARP reply sent by R1)
d. Because each subnet should be able to host a maximum of 17 hosts, five address bits are needed for each subnet. So, let's assign the left subnet XX.YY.ZZ.xx0*****/27, where the XX, YY, and ZZ are 8 bit numbers. Each x is a bit and the five *'s correspond to the five address bits for this network. For the right subnet, well use XX.YY.ZZ.xx1*****/27. Any address in the range XX.YY.ZZ.xx000001 to XX.YY.ZZ.xx011110 can be used for A, C, and interface 1. Any address in the range XX.YY.ZZ.xx100001 to XX.YY.ZZ.xx111110 can be used for B and interface 1.
e. (1) MAC-D: AA:10:F3:5C:01:04, MAC-S: 20:FF:3A:BC:01:4E, IP-D: e.g., XX.YY.ZZ.xx000001 , IP-S: e.g., XX.YY.ZZ.xx100001

(3) MAC-D: 10:D4:E1:A8:97:F0, MAC-S: BB:89:34:E7:01:3B, IP-D: : e.g., XX.YY.ZZ.xx000001, IP-S: e.g., XX.YY.ZZ.xx100001; (2) same as (1).

Problem 6 (10 Points):

This problem focuses on the new approach of Software Defined Networking (SDN).

a. Explain what happens when a packet arrives at the switch and no matching flow table for that packet exists on the switch?
b. What information is contained in a flow table entry?
c. Name the header fields that can be used for matching in the case of OpenFlow?
d. What happens if a flow rule times out?
e. What OpenFlow message is used to add a new flow table in a switch? What message does the switch use to let the controller know that it received a packet for which it does not have a matching rule?

Solution:

a. Packet will be sent from the switch to the controller. Controller checks if a new flow table should be installed in the switch after checking the packet. If so, new flow table entry is sent to the switch and the packet is sent to the switch including information on which port of the switch it should be sent out.

b.

Match	Action	Counter	Priority	Time-out

c.

Switch Port	VLAN ID	VLAN pcp	MAC src	MAC dst	Eth type	IP Src	IP Dst	IP ToS	IP Prot	L4 sport	L4 dport

d. It will be removed from the switch. A packet that would have matched that rule will go through the routine described in a).

e. Flow-mod, Packet-in.