CSC358 Tutorial 10

Julian Sequeira and KyoKeun Park

April 03, 2020

University of Toronto Mississauga

Question 1: Concept Review

(a) How big is the MAC address space? The IPv4 address space? The IPv6 address space?

(b) What's the difference between a hub, a switch, and a router?

(c) What's the difference between CSMA/CD and CSMA/CA?

For the answers, review the lectures, books, go to office hours, and use the discussion board!

Question 2: Money vs Throughput

Consider the institutional network on Page 14 of Week 11 lecture slides. Suppose that all links are 100 Mbps, and any of the 13 hosts (the ones at the bottom) or the 2 servers (mail server and web server) can send to any other host or server.

(a) What is the maximum total aggregate throughput that can be achieved among the 13 hosts and 2 servers in this network?

(b) Now suppose that the institution wants to save money so the 4 switches at the bottom are replaces by hubs (much less expensive). Answer the above question again.

(c) Now suppose that the institution wants to save more money therefore all switches are replaced by hubs. Answer the same question again.

Question 2: Money vs Throughput

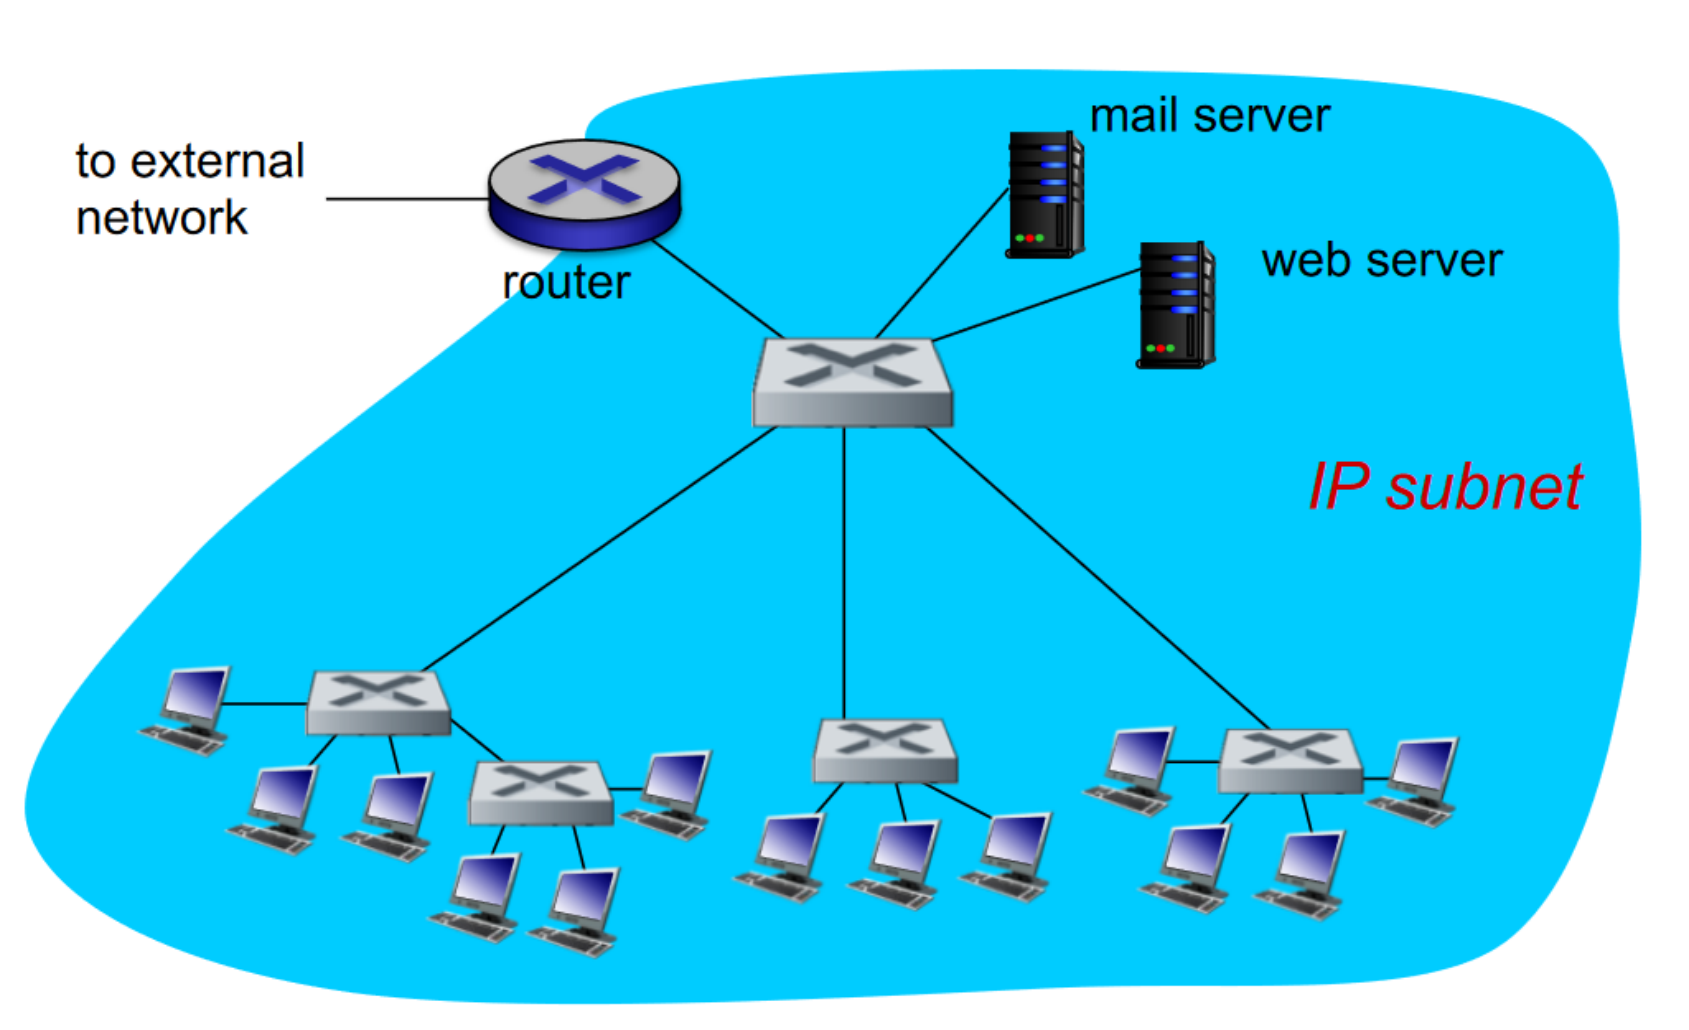

Question 2: Money vs Throughput

(a) If all 15 nodes are using the entirety of their link's bandwidth, a $15x100 = 1500$ Mbps aggregate is possible

Question 2: Money vs Throughput

(a) If all 15 nodes are using the entirety of their link's bandwidth, a 15x100 = 1500 Mbps aggregate is possible

(b) Remember: all links in a hub share the same collision domain. So each hub has a maximum throughput of 100 Mbps. The first two hubs will share a collision domain since they are connected, so account for 100 Mbps for them and 100 Mpbs each for the other two hubs. Then if the web server and mail server send data at their maximum rates of 100 Mbps each, that all adds up to a maximum aggregate throughput of 500 Mbps.

Question 2: Money vs Throughput

(a) If all 15 nodes are using the entirety of their link's bandwidth, a 15x100 = 1500 Mbps aggregate is possible

(b) Remember: all links in a hub share the same collision domain. So each hub has a maximum throughput of 100 Mbps. The first two hubs will share a collision domain since they are connected, so account for 100 Mbps for them and 100 Mpbs each for the other two hubs. Then if the web server and mail server send data at their maximum rates of 100 Mbps each, that all adds up to a maximum aggregate throughput of 500 Mbps.

(c) Now all 15 nodes are in the same collision domain, and altogether can only achieve a maximum total aggregate throughput of 100 Mbps.

Question 3: Learning Switch

Consider the above network. Suppose that the switch tables at all switches are initially empty. Below is the a sequence of events to happen among the switches. Discuss which entries are added to the switch tables after each event.

(a) A sends a frame to G

(b) G replies a frame to A

(c) D sends a frame to F

(d) F replies a frame to D

(e) A send a frame to D

(f) D replies a frame to A

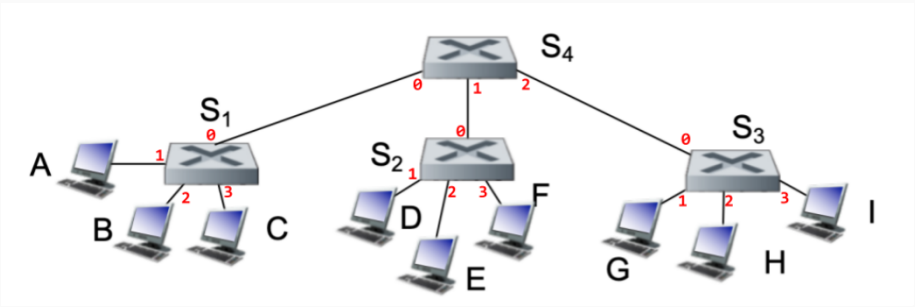

Question 3: Learning Switch

Note that generally, TTL is another variable which is used during learning, but we'll ignore it for this question.

(a) A sends a frame to G

S1: (A, 1)

S2: (A, 0)

S3: (A, 0)

S4: (A, 0)

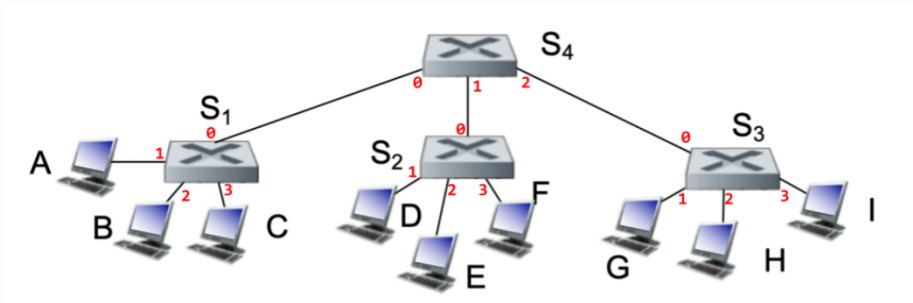

Question 3: Learning Switch

(b) G replies a frame to A

S1: (A, 1), (G, 0)

S2: (A, 0)

S3: (A, 0), (G, 1)

S4: (A, 0), (G, 2)

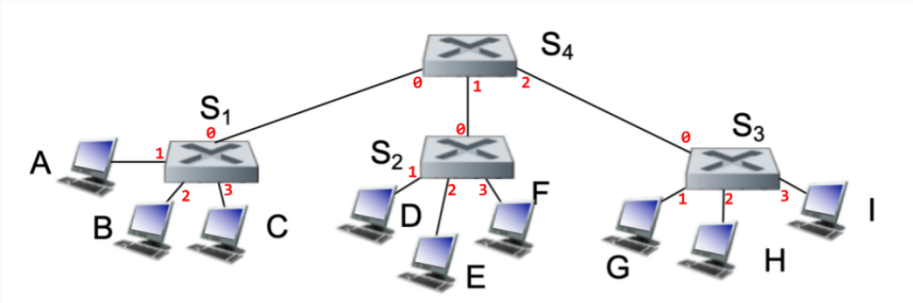

Question 3: Learning Switch

(c) D sends a frame to F

S1: (A, 1), (G, 0), (D, 0)

S2: (A, 0), (D, 1)

S3: (A, 0), (G, 1), (D, 0)

S4: (A, 0), (G, 2), (D, 1)

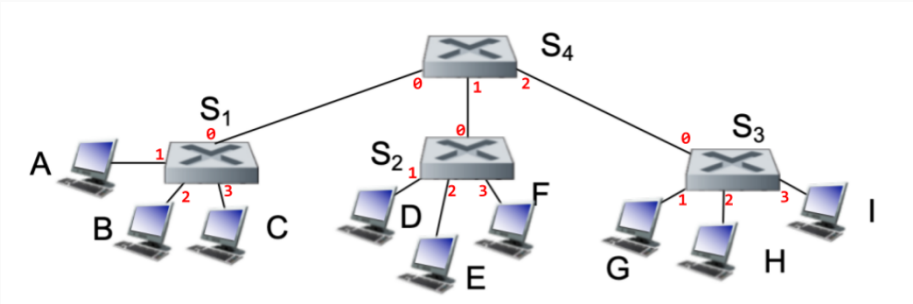

Question 3: Learning Switch

(d) F replies a frame to D

S1: (A, 1), (G, 0), (D, 0)

S2: (A, 0), (D, 1), (F, 3)

S3: (A, 0), (G, 1), (D, 0)

S4: (A, 0), (G, 2), (D, 1)

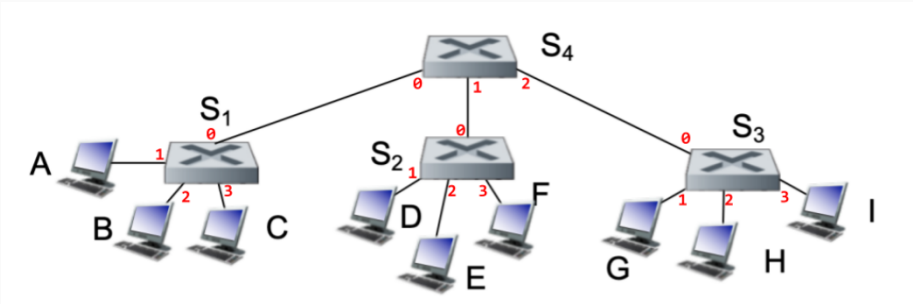

Question 3: Learning Switch

(e) A sends a frame to D

S1: (A, 1), (G, 0), (D, 0)

S2: (A, 0), (D, 1), (F, 3)

S3: (A, 0), (G, 1), (D, 0)

S4: (A, 0), (G, 2), (D, 1)

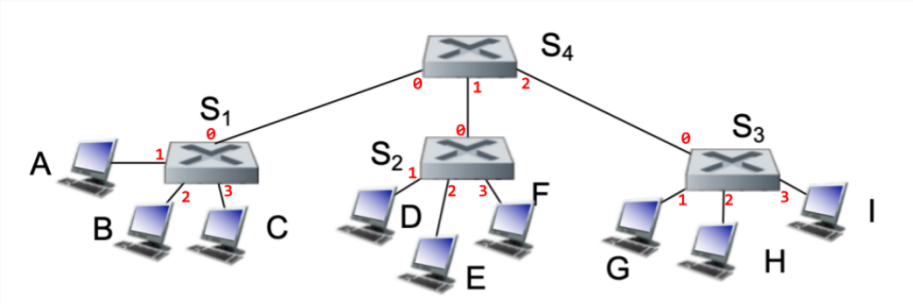

Question 3: Learning Switch

(f) D replies a frame to A

S1: (A, 1), (G, 0), (D, 0)

S2: (A, 0), (D, 1), (F, 3)

S3: (A, 0), (G, 1), (D, 0)

S4: (A, 0), (G, 2), (D, 1)

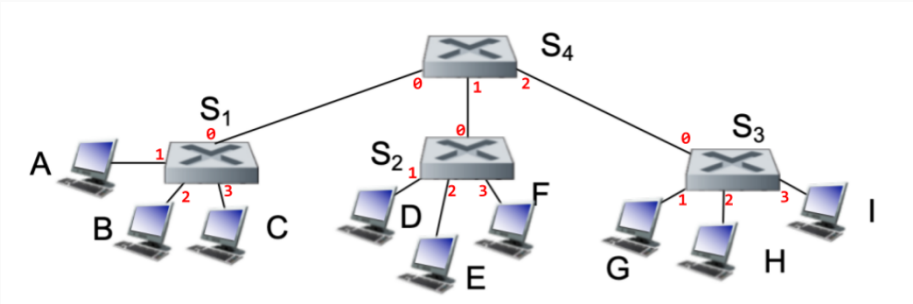

Question 4: Wi-Fi Speed

Suppose your apartment has a Tim Horton's *and* a Starbucks next door, which means that you've got two free Wi-Fi's to choose from (nice neighbourhood!). Suppose that the two coffee shops each operates its own AP and has its own IP address block. Consider the following two scenarios.

(a) Suppose that, by accident, both coffee shops configured their AP to operate over channel 11 (recall that 802.11 divides the 2.4-2.485GHz spectrum into 11 channels). Will the 802.11 protocol completely break down in this situation, i.e., no more free Wi-Fi? Discuss what happens when two hosts (e.g., laptops), each associated with a different AP, attempts to transmit at the same time.

(b) Now suppose that the Tim Horton's AP admin changed its channel from 11 to 1. What difference would you, the freeloader, observe?

A related life pro tip:

https://www.reddit.com/r/LifeProTips/comments/4jcp2o/lpt_download_wifi_analyzer_to_determine_what/

Question 4: Wi-Fi Speed

(a) It will work- the two APs can work in parallel over the same channel. A host will associate with an SSIP (the name of an access point), creating a virtual link between host and AP. Frames will be sent to both APs, only the AP it is addressed to will process them. Label the two APs AP1 and AP2 and suppose the host associates with AP1. When the host sends a frame, it will be addressed to AP1's MAC address and hence be ignored by AP2.

The two APs will be sharing the same wireless bandwidth. So a host H1 transmitting to AP1 and a host H2 transmitting to AP2 at the same time will result in a collision. For 802.11b, the maximum aggregate transmission rate for two APs is 11 Mbps.

Question 4: Wi-Fi Speed

(a) It will work- the two APs can work in parallel over the same channel. A host will associate with an SSIP (the name of an access point), creating a virtual link between host and AP. Frames will be sent to both APs, only the AP it is addressed to will process them. Label the two APs AP1 and AP2 and suppose the host associates with AP1. When the host sends a frame, it will be addressed to AP1's MAC address and hence be ignored by AP2.

The two APs will be sharing the same wireless bandwidth. So a host H1 transmitting to AP1 and a host H2 transmitting to AP2 at the same time will result in a collision. For 802.11b, the maximum aggregate transmission rate for two APs is 11 Mbps.

(b) Now the two APs are not sharing the same wireless bandwidth, so the aggregate maximum transmission rate for the two APs is 22 Mbps (11 Mbps for each AP). There are no more collisions between hosts in different APs transmitting at the same time.