

CS 5480/6480: Computer Networks – Spring 2012
Homework 4 Solutions
Due by 1:25 PM on April 11th 2012

Important: The solutions to the homework problems from the course book have been provided by the authors.

CS5480 total points: 31
CS6480 total points: 40

Question 1 (Slotted Aloha) 5 points (1.25 points for each part): Suppose four active nodes—nodes A, B, C and D—are competing for access to a channel using slotted ALOHA. Assume each node has an infinite number of packets to send. Each node attempts to transmit in each slot with probability p . The first slot is numbered slot 1, the second slot is numbered slot 2, and so on.

- What is the probability that node A succeeds for the first time in slot 5?
- What is the probability that some node (either A, B, C or D) succeeds in slot 4?
- What is the probability that the first success occurs in slot 3?
- What is the efficiency of this four-node system?

a) $(1 - p(A))^4 p(A)$
 where, $p(A)$ = probability that A succeeds in a slot

$$\begin{aligned}
 p(A) &= p(A \text{ transmits and } B \text{ does not and } C \text{ does not and } D \text{ does not}) \\
 &= p(A \text{ transmits}) p(B \text{ does not transmit}) p(C \text{ does not transmit}) p(D \text{ does not transmit}) \\
 &= p(1 - p) (1 - p)(1 - p) = p(1 - p)^3
 \end{aligned}$$

Hence, $p(A \text{ succeeds for first time in slot 5}) = (1 - p(A))^4 p(A) = (1 - p(1 - p)^3)^4 p(1 - p)^3$

$$\begin{aligned}
 \text{b) } p(A \text{ succeeds in slot 4}) &= p(1 - p)^3 \\
 p(B \text{ succeeds in slot 4}) &= p(1 - p)^3 \\
 p(C \text{ succeeds in slot 4}) &= p(1 - p)^3 \\
 p(D \text{ succeeds in slot 4}) &= p(1 - p)^3
 \end{aligned}$$


$$\begin{aligned}
 p(\text{either } A \text{ or } B \text{ or } C \text{ or } D \text{ succeeds in slot 4}) &= 4 p(1 - p)^3 \\
 (\text{because these events are mutually exclusive})
 \end{aligned}$$

$$\begin{aligned}
 \text{c) } p(\text{some node succeeds in a slot}) &= 4 p(1 - p)^3 \\
 p(\text{no node succeeds in a slot}) &= 1 - 4 p(1 - p)^3
 \end{aligned}$$

Hence, $p(\text{first success occurs in slot 3}) = p(\text{no node succeeds in first 2 slots}) p(\text{some node succeeds in 3rd slot}) = (1 - 4 p(1 - p)^3)^2 4 p(1 - p)^3$

d) efficiency = $p(\text{success in a slot}) = 4 p(1 - p)^3$

Question 2 (Ethernet) 6 points (2 points for each part): Consider the figure below. Now we replace the router between subnets 1 and 2 with a switch S1, and label the router between subnets 2 and 3 as R1.



- Consider sending an IP datagram from Host E to Host F. Will Host E ask router R1 to help forward the datagram? Why? In the Ethernet frame containing the IP datagram, what are the source and destination IP and MAC addresses?
- Suppose E would like to send an IP datagram to B, and assume that E's ARP cache does not contain B's MAC address. Will E perform an ARP query to find B's MAC address? Why? In the Ethernet frame (containing the IP datagram destined to B) that is delivered to router R1, what are the source and destination IP and MAC addresses?
- Suppose Host A would like to send an IP datagram to Host B, and neither A's ARP cache contains B's MAC address nor does B's ARP cache contain A's MAC address. Further suppose that the switch S1's forwarding table contains entries for Host B and router R1 only. Thus, A will broadcast an ARP request message. What actions will switch S1 perform once it receives the ARP request message? Will router R1 also receive this ARP request message? If so, will R1 forward the message to Subnet 3? Once Host B receives this ARP request message, it will send back to Host A an ARP response message. But will it send an ARP query message to ask for A's MAC address? Why? What will switch S1 do once it receives an ARP response message from Host B?

a) No. E can check the subnet prefix of Host F's IP address, and then learn that F is on the same LAN. Thus, E will not send the packet to the default router R1.

Ethernet frame from E to F:

Source IP = E's IP address

Destination IP = F's IP address

Source MAC = E's MAC address

Destination MAC = F's MAC address

b) No, because they are not on the same LAN. E can find this out by checking B's IP address.

Ethernet frame from E to R1:

Source IP = E's IP address

Destination IP = B's IP address

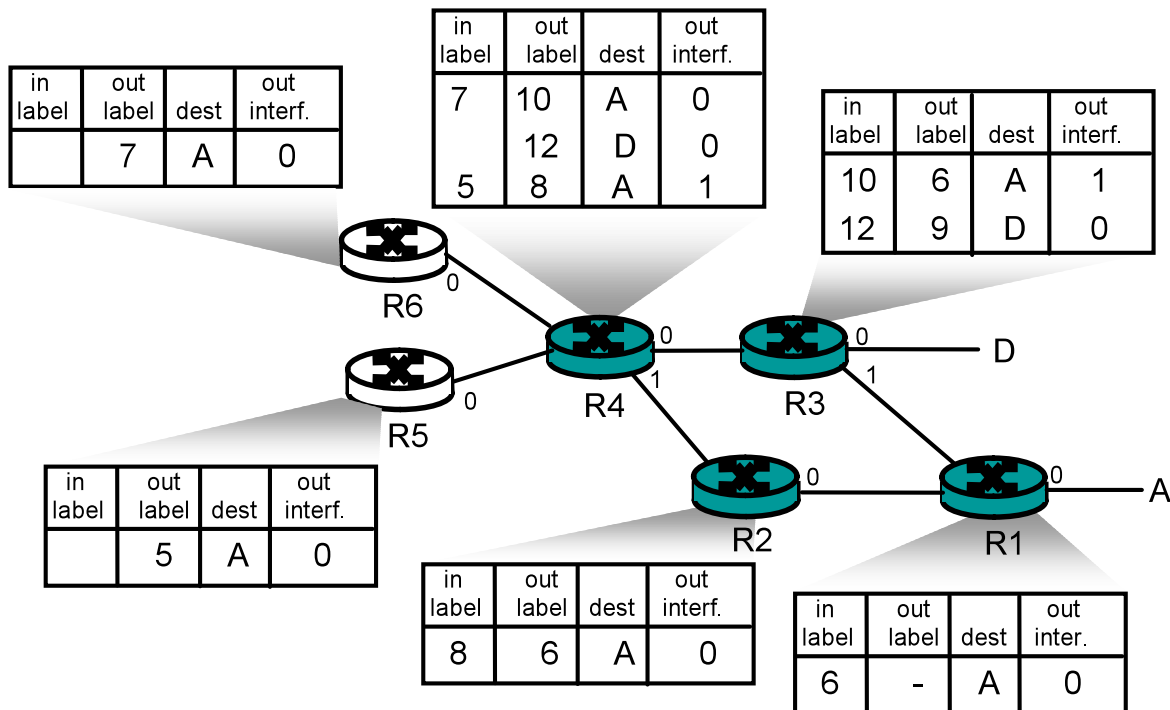
Source MAC = E's MAC address

Destination MAC = The MAC address of R1's interface connecting to Subnet 3.


c) Switch S1 will broadcast the Ethernet frame via both its interfaces as the received ARP frame's destination address is a broadcast address. And it learns that A resides on Subnet 1 which is connected to S1 at the interface connecting to Subnet 1. And, S1 will update its forwarding table to include an entry for Host A.

Yes, router R1 also receives this ARP request message, but R1 won't forward the message to Subnet 3. B won't send ARP query message asking for A's MAC address, as this address can be obtained from A's query message. Once switch S1 receives B's response message, it will add an entry for host B in its forwarding table, and then drop the received frame as destination host A is on the same interface as host B (i.e., A and B are on the same LAN segment).

Question 3 (4 points): Consider the MPLS network shown in slide 5.79, and suppose that routers R5 and R6 are now MPLS enabled. Suppose that we want to perform traffic engineering so that packets from R6 destined for A are switched to A via R6-R4-R3-R1, and packets from R5 destined for A are switched via R5-R4-R2-R1. Show the MPLS tables in R5 and R6, as well as the modified table in R4, that would make this possible.



Question 4 (4 points): Consider the single switch VLAN in the figure below, 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 an IP datagram from an EE host to a CS host (Hint: reread the material on slides 5-45 – 5-50).



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.

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.

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 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.

Question 5 (WiFi) 4 points:

(a) (2 points) Suppose there are two ISPs providing WiFi access in a particular café, with each ISP operating its own AP and having its own IP address block.

a. Further suppose that by accident, each ISP has configured its AP to operate over channel 11. Will the 802.11 protocol completely break down in this situation? Discuss what happens when two stations, each associated with a different ISP, attempt to transmit at the same time.

b. Now suppose that one AP operates over channel 1 and the other over channel 11. How do your answers change?

- a. The two APs will typically have different SSIDs and MAC addresses. A wireless station arriving to the café will associate with one of the SSIDs (that is, one of the APs). After association, there is a virtual link between the new station and the AP. Label the APs AP1 and AP2. Suppose the new station associates with AP1. When the new station sends a frame, it will be addressed to AP1. Although AP2 will also receive the frame, it will not process the frame because the frame is not addressed to it. Thus, the two ISPs can work in parallel over the same channel. However, the two ISPs will be sharing the same wireless bandwidth. If wireless stations in different ISPs transmit at the same time, there will be a collision. For 802.11b, the maximum aggregate transmission rate for the two ISPs is 11 Mbps.
- b. Now if two wireless stations in different ISPs (and hence different channels) transmit at the same time, there will not be a collision. Thus, the maximum aggregate transmission rate for the two ISPs is 22 Mbps for 802.11b.

(b) (2 points) In Step 4 of the CSMA/CA protocol, a station that successfully transmits a frame begins the CSMA/CA protocol for a second frame at step 2, rather than at step 1. What rationale might the designers of CSMA/CA have had in mind by having such a station not transmit the second frame immediately (if the channel is sensed idle)?

Suppose that wireless station H1 has 1000 long frames to transmit. (H1 may be an AP that is forwarding an MP3 to some other wireless station.) Suppose initially H1 is the only station that wants to transmit, but that while half-way through transmitting its first frame, H2 wants to transmit a frame. For simplicity, also suppose every station can hear every other station's signal (that is, no hidden terminals). Before transmitting, H2 will sense that the channel is busy, and therefore choose a random backoff value.

Now suppose that after sending its first frame, H1 returns to step 1; that is, it waits a short period of times (DIFS) and then starts to transmit the second frame. H1's second frame will then be transmitted while H2 is stuck in backoff, waiting for an idle channel. Thus, H1 should get to transmit all of its 1000 frames before H2 has a chance to access the channel. On the other hand, if H1 goes to step 2 after transmitting a frame, then it too chooses a random backoff value, thereby giving a fair chance to H2. Thus, fairness was the rationale behind this design choice.

Question 6 (IP Mobility) 10 points:

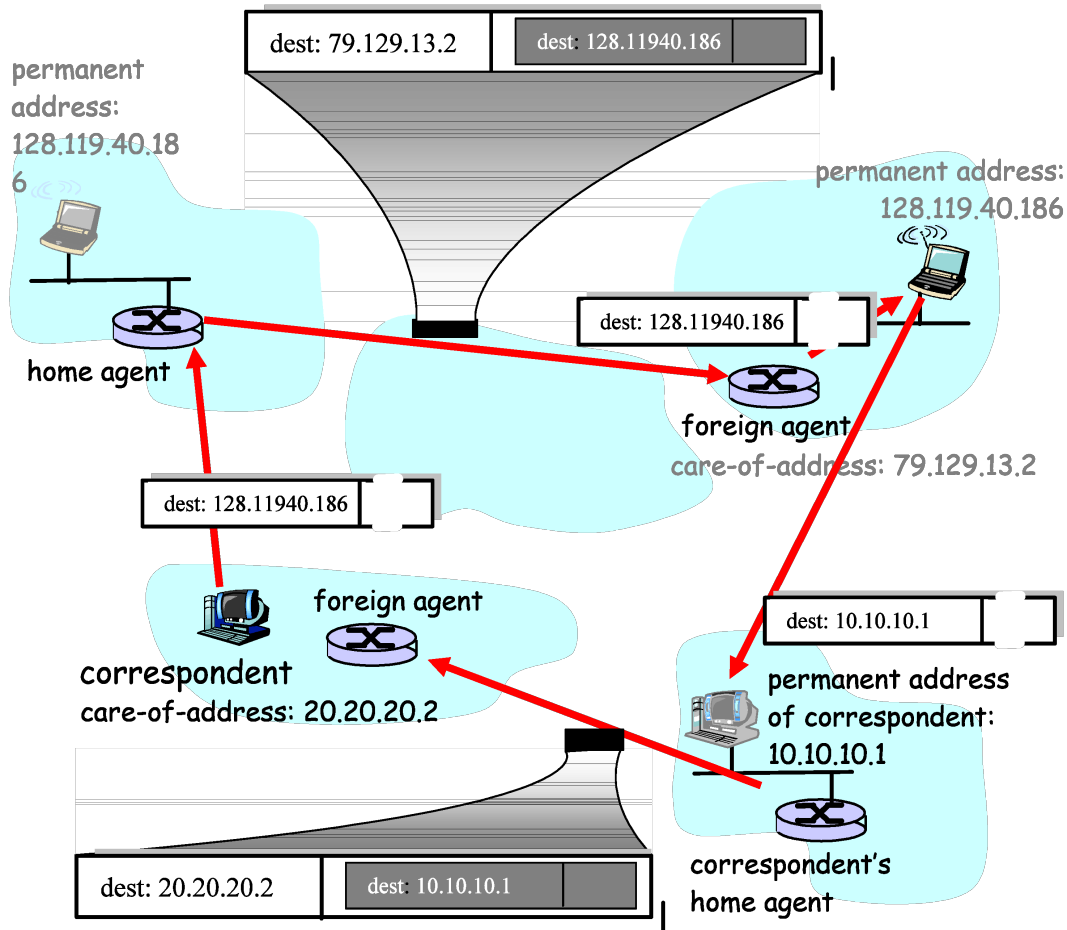
(a) (4 points) One proposed solution that allowed mobile users to maintain their IP addresses as they moved among foreign networks was to have a foreign network

advertise a highly specific route to the mobile user and use the existing routing infrastructure to propagate this information throughout the network. We identified scalability as one concern. Suppose that when a mobile user moves from one network to another, the new foreign network advertises a specific route to the mobile user, and the old foreign network withdraws its route. Consider how routing information propagates in a distance-vector algorithm (particularly for the case of interdomain routing among networks that span the globe).

- a. Will other routers be able to route datagrams immediately to the new foreign network as soon as the foreign network begins advertising its route?*
- b. Is it possible for different routers to believe that different foreign networks contain the mobile user?*
- c. Discuss the timescale over which other routers in the network will eventually learn the path to the mobile users.*

- a. No. All the routers might not be able to route the datagram immediately. This is because the Distance Vector algorithm (as well as the inter-AS routing protocols like BGP) is decentralized and takes some time to terminate. So, during the time when the algorithm is still running as a result of advertisements from the new foreign network, some of the routers may not be able to route datagrams destined to the mobile node.
- b. Yes. This might happen when one of the nodes has just left a foreign network and joined a new foreign network. In this situation, the routing entries from the old foreign network might not have been completely withdrawn when the entries from the new network are being propagated.
- c. The time it takes for a router to learn a path to the mobile node depends on the number of hops between the router and the edge router of the foreign network for the node.

(b) (2 points) Suppose the correspondent in slide 6-43 were mobile. Sketch the additional network-layer infrastructure that would be needed to route the datagram from the original mobile user to the (now mobile) correspondent. Show the structure of the datagram(s) between the original mobile user and the (now mobile) correspondent, as in slide 6-48.



(c) (2 points) Consider two mobile nodes in a foreign network having a foreign agent. Is it possible for the two mobile nodes to use the same care-of address in mobile IP? Explain your answer.

Two mobiles could certainly have the same care-of-address in the same visited network. Indeed, if the care-of-address is the address of the foreign agent, then this address would be the same. Once the foreign agent decapsulates the tunneled datagram and determines the address of the mobile, then separate addresses would need to be used to send the datagrams separately to their different destinations (mobiles) within the visited network.

Question 7 (required for cs6480, extra credit for cs5480) 9 points:

Read the following paper: "RSVP: A New Resource Reservation Protocol," by L. Zhang et al, in the IEEE Communications Magazine, May 2002. This paper is available from http://www.eng.utah.edu/~cs5480/rsvp_network_magazine.pdf. Answer the following questions that are based on this paper.

(a) (2 points) Write two uses of reservation/path refresh messages.

1. Recover from an earlier lost reservation message.
2. Handle receiver/sender that goes away without teardown (caused by the lack of a refresh message).
3. Deal with path changes.

(b) (1 point) What is soft state? Write one advantage of using soft state?

Soft state lives only for a short amount of time unless refreshed. It is robust against equipment failure. For example, if resources are reserved and due to equipment failure they are not explicitly released, these would be wasted. In the case of soft state based resource reservation, these resources will be released unless the node that reserves resources receives a refresh.

(c) (1 point) Does RSVP determine what routes packet will take? Explain in one sentence.

No. That is the job of the routing protocols.

(d) (1 point) What is the difference between the two filter types – no filter and fixed filter?

(e) (3 points) Consider Figure 3 and the no filter example in the paper. Show the tables at S1, S2, and S3 after the path messages from all H1 – H5 hosts, and RESV messages from H1 and H2 have all been seen by these routers. Note that the table at S1 is given in the paper.

S1:

in	L1, L2, L6
out	L1(b), L2(b), L6

S2:

in	L5, L6, L7
out	L5, L6(b), L7

S3:

in	L3, L4, L7
out	L3, L4, L7(b)