CS 5480/6480: Computer Networks – Spring 2012
Homework 3
Due by 1:25 PM MT, Monday March 5th 2012

Important:

- **No cheating will be tolerated.**
- **No Late Submission will be allowed.**
- **Total points for 5480 = 37**
- **Total points for 6480 = 44**

Question 1 (IP/TCP Headers) ***3 points*****:** Look at the 40byte dump of an IP packet containing a TCP segment below.
45 20 03 c5 78 06 00 00 34 06 ca 1f d1 55 ad 71 c0 a8 01 7e
00 50 9a 03 3e 64 e5 58 df d0 08 b3 80 18 00 de 00 02 00 00
Identify all the fields of the IP and TCP header.

Question 2 (IP Fragmentation) ***2 points*****:** Consider sending a 2400-byte datagram into a link that has an MTU of 700 bytes. Suppose the original datagram is stamped with the identification number 422. How many fragments are generated? What are the values in the various fields in the IP datagram(s) generated related to fragmentation?

Question 3 (IP addressing) ***4 points*****:**
(a) (*2 points*) Consider a router that interconnects three subnets: Subnet 1, Subnet 2, and Subnet 3. Suppose all of the interfaces in each of these three subnets are required to have the prefix 223.1.17/24. Also suppose that Subnet 1 is required to support up to 63 interfaces, Subnet 2 is to support up to 95 interfaces, and Subnet 3 is to support up to 16 interfaces. Provide three network addresses (of the form a.b.c.d/x) that satisfy these constraints.

(b) (*2 points*) Consider a subnet with prefix 128.119.40.128/26. Give an example of one IP address (of form xxx.xxx.xxx.xxx) that can be assigned to this network. Suppose an ISP owns the block of addresses of the form 128.119.40.64/26. Suppose it wants to create four subnets from this block, with each block having the same number of IP addresses. What are the prefixes (of form a.b.c.d/x) for the four subnets?

Question 4 (NAT) ***3 points*****:** Suppose you are interested in detecting the number of hosts behind a NAT. You observe that the IP layer stamps an identification number sequentially on each IP packet. The identification number of the first IP packet generated by a host is a random number, and the identification numbers of the subsequent IP packets are sequentially assigned. Assume all IP packets generated by hosts behind the NAT are sent to the outside world.
a. Based on this observation, and assuming you can sniff all packets sent by the NAT to the outside, can you outline a simple technique that detects the number of unique hosts behind a NAT? Justify your answer.

b. If the identification numbers are not sequentially assigned but randomly assigned, would your technique work? Justify your answer.

Question 5 (Routing Protocols) ***18 points*****:**

(a) (*3 points*) Consider the following network. With the indicated link costs, use Dijkstra's shortest-path algorithm to compute the shortest path from x to all network nodes. Show how the algorithm works by computing a table similar to to the table on slide 4-80.

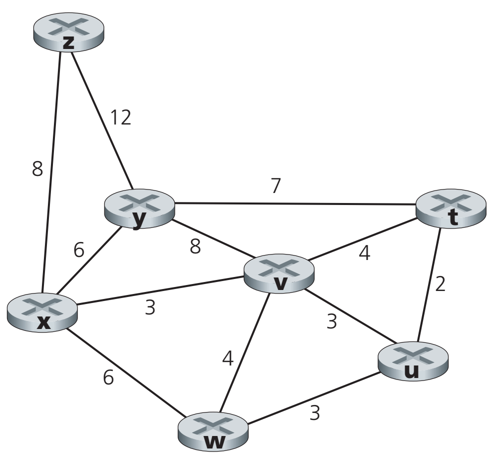

(b) (*3 points*) Consider the three-node topology shown on slide 4-89. Let the link costs be $c(x,y) = 3$, $c(y,z) = 6$, $c(z,x) = 4$. Compute the distance tables after the initialization step and after each iteration of the distance-vector algorithm as done for the example on slide 4-89.

(c) (*3 points*) Consider the count-to-infinity problem in the distance vector routing. Will the count-to-infinity problem occur if we decrease the cost of a link? Why? How about if we connect two nodes which do not have a link?

(d) (*4 points*) Consider the network shown below. Suppose AS3 and AS2 are running OSPF for their intra-AS routing protocol. Suppose AS1 and AS4 are running RIP for their intra-AS routing protocol. Suppose eBGP and iBGP are used for the inter-AS routing protocol. Initially suppose there is *no* physical link between AS2 and AS4.

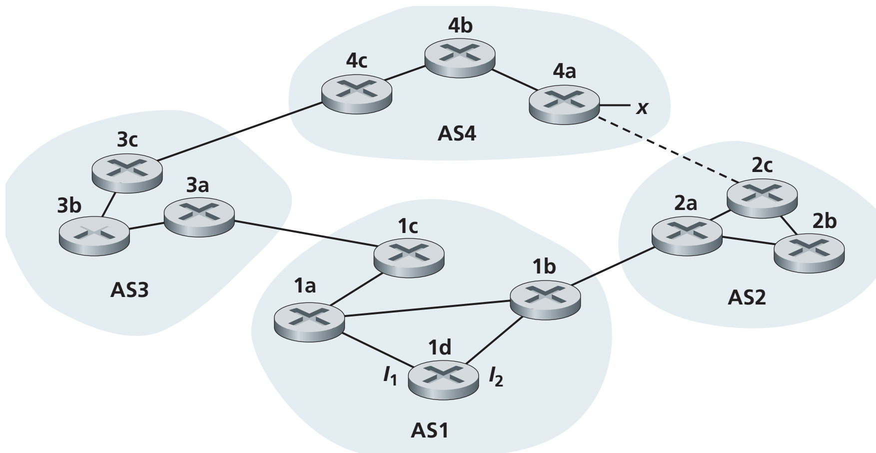

a. Router 3c learns about prefix *x* from which routing protocol: OSPF, RIP, eBGP, or iBGP?
b. Router 3a learns about *x* from which routing protocol?
c. Router 1c learns about *x* from which routing protocol?
d. Router 1d learns about *x* from which routing protocol?

(e) (*3 points*) Referring to the above problem, once router 1d learns about *x* it will put an entry (*x*, *I*) in its forwarding table.
a. Will *I* be equal to *I*1 or *I*2 for this entry? Explain why in one sentence.
b. Now suppose that there is a physical link between AS2 and AS4, shown by the dotted line. Suppose router 1d learns that *x* is accessible via AS2 as well as via AS3. Will *I* be set to *I*1 or *I*2? Explain why in one sentence.
c. Now suppose there is another AS, called AS5, which lies on the path between AS2 and AS4 (not shown in diagram). Suppose router 1d learns that *x* is accessible via AS2 AS5 AS4 as well as via AS3 AS4. Will *I* be set to *I*1 or *I*2? Explain why in one sentence.

(f) (*2 points*) Consider the figure on slide 4-119. B would never forward traffic destined to Y via X based on BGP routing. But there are some very popular applications for which data packets go to X first and then flow to Y. Identify one such application, and describe how data packets follow a path not given by BGP routing.

Question 6 (Broadcast and Multicast Routing) ***7 points*****:**
(a) (*2 points*) Consider the topology shown in Figure 4.44. Suppose that all links have unit cost and that node *E* is the broadcast source. Using arrows like those shown in Figure 4.44 indicate links over which packets will be forwarded using RPF, and links over which packets will not be forwarded, given that node *E* is the source.

(b) (*2 points*) Describe the IP multicast service model.

(c) (*3 points*) Describe the limitations of multicast routing based on flooding and pruning. How can Core-based-tree (CBT) multicast routing help overcome some of these limitations? What are two limitations of CBT multicast routing?

Question 7 (required for CS6480, extra credit for CS5480) ***7 points***:

Read the following paper: "*On Fast and Accurate Detection of Unauthorized Access Points Using Clock Skews*" by Suman Jana and Sneha Kumar Kasera, in the IEEE Transactions on Mobile Computing, March 2010. This paper is available from http://eng.utah.edu/~cs5480/readings/jana2010.pdf.

Answer the following questions that are based on this paper.

(a) (*1 point*) Define clock skew.

(b) (*2 points*) Let us plot the observed clock offset, in microseconds, on the y-axis and the time since the start of the fingerprinting measurements, in seconds, at the fingerprinter, on the x-axis. Let (5, 60) and (7.5, 75) be two points at times 5s and 7.5s, where the clock offset is observed by the fingerprinter to be 60 and 75, respectively. Estimate the clock skew of the fingerprintee from these two points. You can assume that the network delays are negligible.

(c) (*2 points*) What clock skew behavior do the authors observe in the case of virtual access points?

(d) (*2 points*) Why is it not easy to fabricate clock skews of access points?