Homework 2

Due: 11 March 2011, 4pm

Note: no late days discounted until **March 14th, 2011 at 4pm**. *However, no points after that, as we will release the solutions by then. If you hand in by the due date, March 11th, 2011 at 4pm, we will grade your homework before the midterm.*

Problem 1 - Addressing, Aggregation, Forwarding

1. You are the CTO of a growing startup and have to get IP addresses to connect 560 computers to the Internet. You can get IP addresses from two providers, IPMart and EastSideIP. IPMart sells classA, class B and class C blocks, while EastSideIP sells CIDR blocks. As the IPv4 address space is scarce, you want to save money and get the smallest number of addresses possible.

a. If you get one block from IPMart, which class do you have to get? What is the problem with that?

b. If you get one block from EastSideIP, how many bits are there in the mask (e.g., is it a /8, /22)? How many addresses are wasted?

c. Suppose you can get two blocks from EastSideIP, and they can be of different sizes. How many bits are there in the masks for each of the blocks? How many addresses are wasted now?

2. Another customer of EastSideIP needs to get 8000 IP addresses. You work for EastSide, and see that you have the address range from 128.140.80.0 to 128.140.112.255 available.

a. What is the best CIDR block from this range you can allocate the customer?

b. Why is it best to minimize the number of CIDR blocks you allocate?

c. Why is it best to also minimize the size of the address blocks?

3. Suppose you have the following routing table in your router

Destination	Netmask	NextHop	Interface
0.0.0.0	0.0.0.0	100.10.1.1	eth0
128.148.0.0	255.255.0.0	128.148.0.1	eth1
128.148.32.0	255.255.240.0	128.148.32.1	eth2
128.148.34.128	255.255.255.128	128.148.34.129	eth3

What is the next hop for each of these addresses, given that you use longest-prefix matching?

a. 128.148.34.143

b. 128.148.34.12

c. 128.148.38.1

d. 200.192.120.12

e. 128.148.12.2

f. 128.140.0.1

Problem 2 - IP Fragmentation

1. You need to send a packet of 4800 bytes over a network path with 3 links A, B, and C, with MTUs of 1500, 576, and 1500 bytes.

a. If the origin doesn't know anything about the path other than the MTU of the first link, how many packets (and their sizes) that flow through each link? (Ignore the size of the link layer headers)

b. What fields are different in the IP header for the first fragment that goes on link A and on link B?

c. If link B has 0.1% chance of dropping a packet (assume the other links don't drop packets), what is the probability that the original packet will be correctly reconstructed in the destination? *(Hint: it might be easier to think that each fragment has a 99.9% chance of making it through B.)*

d. What does this probability change to if the MTU in link B is increased to 1500 bytes?

Problem 3 - Distance Vector Routing

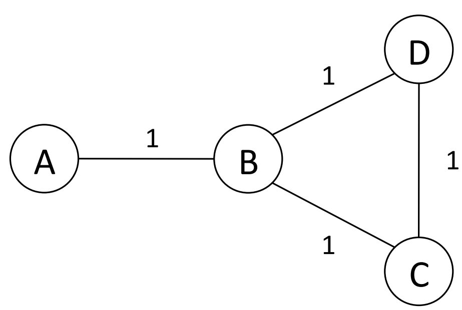

Consider network in the graph above, and only consider the routes with destination A. The network is running a distance vector protocol, and considers infinity to be 16. Use the following notation for a routing table entry: B(A,C,1) means that node B has an entry saying it can reach node A through next hop C, with cost 1. Likewise, a route update from B will say B(A,1). Routing table entries time out and are removed if a node does not get an update from the parent after 5 minutes.

a. If link A-B fails, B immediately sets its routing table entry for A to B(A,-,∞). Give a sequence of events in which a loop and count-to-infinity will occur. Suppose nodes are not using any loop prevention technique.

b. The designers of the protocol decide to add split-horizon to prevent loops. Spoil their party and show a sequence of events that can still cause a loop to form after link A-B goes down, even with split horizon.

c. In a lot of situations the addition of poisoned reverse to split-horizon does not make any difference, other than increasing the size of the routing announcements. Let's look at a scenario in which there *is* a difference. Suppose the following happens after link A-B fails: 1. B sends a route update B:(A,∞), and at the same time C and D send route updates C:(A,2), D:(A,2). 2. C receives the updates from B and D almost simultaneously, and installs C:(A:D,3) in its routing table. Likewise, D installs D:(A,C,3) in its routing table. Explain how this will evolve if the network is using split horizon, and split horizon with poisoned reverse.

d. Explain how a path vector protocol would prevent a loop from forming in this network.