Homework 2

Due: 12 March 2012, 4pm

Note: no late days on this homework, as we will release the solutions promptly at 4pm, so you can read them before the midterm.

Problem 1 - Addressing, Aggregation, Forwarding [11 pts]

1. You are the CTO of a fast growing startup and have to get IP addresses to connect 1200 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? [1 pt]

> You have to get a class B block from IPMart, because a class C only has addresses for 254 computers. (The first address is the network number, and the last one is the broadcast address.) The problem with getting a class B block is that it allocates a space of 2^{16} addresses, and you only need 1200.

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? (Assume that the first and last addresses of the block are not counted as available or used, i.e., pretend, for the calculations, that they don't exist. The first address of a block is used as the network number, and the last one as the broadcast address for the subnet.)[1 pt]

> A CIDR block has to have an integral power of two addresses. The smallest power of two larger than 1200 is 2^{11}, which means that the network mask has $32-11 = 21$ bits. The number of wasted addresses is $2048-2-1200 = 846$. (We also accepted $2048 - 1200 = 848$, if you didn't take into account the first and the last addresses of the block.)

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? [1 pt]

> If we can get two blocks, we should get a /22 block, which will be good for 1024-2 = 1022 computers. For the remaining 178 computers we need to get a /24 block, good for $2^8 - 2 = 254$ addresses. The number of wasted addresses is 254 - 178 = 76. (We also accepted 80 as the number of wasted addresses, if you didn't take into account the first and last addresses of the block for each of the blocks (1024 + 256) - 1200 = 80).

2. You work for EastSideIP and have a large number of consecutive addressed available starting at 198.17.0.0. You receive 2 requests for addresses, 4000 and 2000 and in that order.

a. For each, give the first and last IP address assigned as well as the subnet mask. [1pt]

- 4000: 198.17.0.0 through 198.17.15.255 (a 4096 address block)
- 2000: 198.17.16.0 through 198.17.23.255 (a 2048 address block)

b. Why is it best to minimize the number of CIDR blocks you allocate? [1 pt]

Because the larger the blocks allocated, the smaller the routing tables in upstream routers become.

c. Why is it best to also minimize the size of the address blocks? [1 pt]

Because we improve the utilization of the address space, by reducing wasted allocations.

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

address/mask	next hop
135.46.56.0/22	Interface 0
135.46.60.0/22	Interface 1
192.53.40.0/23	Router 1
default	Router 2

What is the next hop for each of these addresses, given that you use longest-prefix matching? [0.5 pts each]

a. 135.46.63.10

Interface 1

b. 135.46.57.14

Interface 0

c. 192.53.40.7

Router 1

d. 192.53.56.7

Router 2

4. YouTube has prefix 208.65.152.0/22. On 2008/02/24, Pakistan started to advertise routes to prefix 208.65.153.0/24, which caused a lot of users to hit Pakistani servers instead of YouTube servers.

a. Why did routers prefer the Pakistan announcement to the YouTube one? [1 pt]

b. What fraction of the addresses owned by YouTube were hijacked by Pakistan's announcement? [1 pt]

c. If YouTube wanted to stop the attack quickly, which /25 prefixes would it have to announce? [1 pt]

Problem 2 - Distance Vector Routing [7 pts]

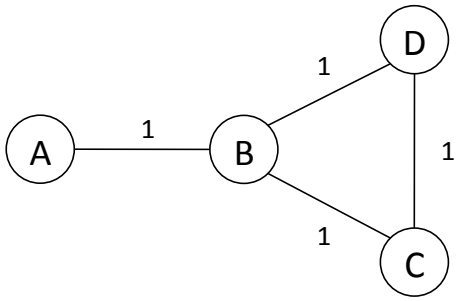

Consider network in the graph above, and only consider the routes with destination A. The network is running a distance vector protocol, and considers infinity to be 16. Use the following notation for a routing table entry: B(A,C,1) means that node B has an entry saying it can reach node A through next 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 3 node loop and count-to-infinity will occur. Suppose nodes are not using any loop prevention technique. [2 pts]

b. The designers of the protocol decide to add split-horizon to prevent loops. Show that this does nothing to prevent the loop created in (a) after link A-B goes down. Why does split-horizon not prevent this loop? What type of loops does this prevent against? Give an example. [3pts]

c. Assume the network uses a path vector protocol. Would this change any of the loops above? If not, what situations is it useful? If so, what are the downsides to this approach[2 pts]

> A permanent loop would not form. In a path vector protocol the entire path to the destination is kept and advertised in a route. If a router sees its address as part of a path that it receives from another one, it knows that choosing that route would create a loop.