

# 15-441:Networking

## Homework 2

Homework 2 is due Friday March 3, 2006, 5pm. Solutions must be submitted electronically, similar to homework 1.

**Solutions for homework 2 will be posted that same evening so you can prepare for the midterm. Thus there will be no late days, i.e. late handins will not be accepted!!!**

### Problem 1: CIDR

Here is short forwarding table:

| Mask        | Network Number | Outgoing Interface |
|-------------|----------------|--------------------|
| 255.0.0.0   | 128.0.0.0      | 1                  |
| 255.240.0.0 | 128.208.0.0    | 3                  |
| default     | -              | 15                 |

Or, equivalently:

| Number bits | Subnet Number | Outgoing Interface |
|-------------|---------------|--------------------|
| 8           | 10000000      | 1                  |
| 12          | 100000001101  | 3                  |
| default     | -             | 15                 |

What interface will the packets with the following destination address be forwarded to:

(a) 128.3.208.23

(b) 123.34.206.9

(c) 128.209.98.9

(d) 128.240.98.126

## Problem 2: Link state routing

Figure 1 shows a network consisting of routers and hosts. Distances between routers are labeled. Interfaces on the routers are also labeled in smaller font. Assume that all these routers are running the Dijkstra algorithm. *You will only show what is going on at router R1.*

For router R1, after *each* iteration of the link-state algorithm show the following information:

1. The tentative set S.
2. The estimated distances D. Specifically:  $D(R1)$ ,  $D(R2)$ ,  $D(R3)$ ,  $D(R4)$ ,  $D(R5)$ .
3. The confirmed set P: Specifically:  $P(R1)$ ,  $P(R2)$ ,  $P(R3)$ ,  $P(R4)$ ,  $P(R5)$ .




Figure 1: *Network of routers and hosts.*

For your answer, please use the format used in class. For each iteration, list the nodes in the “Done” set with their cost and next hop information, list the nodes in the “Horizon” set with their cost and next hop information, and identify the node that will be moved from the Horizon set to the Done set.

## Problem 3: Why traceroute?

The program *traceroute* allows you to find out the path (i.e. sequence of routers) that a packet will follow to a specific destination. The routers along the path are often identified by name, which means that you can learn the identity of the various ISPs your packets travel through.

This problem is very open-ended. Do the best you can. Use the Web. The purpose of this problem is to explore and learn.

For this problem, use *traceroute* to record the path taken by packets from a CMU host to

- [www.cs.stanford.edu](http://www.cs.stanford.edu)
- [www.cisco.com](http://www.cisco.com)
- [www.rug.ac.be](http://www.rug.ac.be)


a) Can you identify the ISPs used in the above queries. (Hint: the router names either contain the ISP's name, or you can find a web site with a similar name, e.g. www.xyz.net, or you can use a search engine).

**Please include the traceroute output in your solution and mark what ISPs use what routers.**

- b) Try as best you can to classify the above ISPs as local, regional, or backbone.
- c) What is the name of the local ISP for CMU?

d) Observe that CMU uses different backbone ISPs depending on the destination. Experiment with different destinations (academic, commercial, national, international, etc.) and see if you can come up with any observations about which backbone CMU routes through for different types of destinations.

### Problem 4: Distance-Vector



We are using a distance vector routing algorithm for the network shown above. Assume that all routers boot at the same time, at time unit 0. Assume that steps are synchronous. Step 1 consists of discovering your distance to all your neighbors. Each subsequent step consists of each router sending an update to all its neighbors and receiving a routing update from all its neighbors.

a) Show the contents of each router's routing table at the end of step 2 by filling in the tables below. The routing table for R1 has been filled out as an example. Each routing table entry consists of a destination address, next hop router and estimated cost for the path to that router.

| R1    |          |      |
|-------|----------|------|
| Dest. | Next Hop | Cost |
| R2    | R2       | 1    |
| R6    | R2       | 2    |
| R3    | R2       | 2    |

| R2    |          |      |
|-------|----------|------|
| Dest. | Next Hop | Cost |
|       |          |      |
|       |          |      |
|       |          |      |

| R3    |          |      |
|-------|----------|------|
| Dest. | Next Hop | Cost |
|       |          |      |
|       |          |      |
|       |          |      |

b) How many steps does it take for the algorithm to converge? Please explain.

c) If router R3 crashes at the end of time unit 17 and is rebooted at the end of time unit 18, at the end of what time unit, will it have a complete routing table, i.e. one that contains routing information on how to reach all other routers? Please explain.