CSEE 4119: Computer Networks — Due on Monday, April 20th, 2015,

#4 – Network Layer

Pooja Shah, Olalekan Afuye, Yelin Hong, Jen-Chieh Huang, Nivvedan Senthamil A. Chaintreau

How to read this assignment **:** Exercise levels are indicated as follows

($\rightarrowtail$) "elementary": the answer is not strictly speaking obvious, but it fits in a single sentence, and it is an immediate application of results covered in the lectures.
Use them as a checkpoint: it is strongly advised to go back to your notes if the answer to one of these questions does not come to you in a few minutes.

($\curvearrowright$) "intermediary": The answer to this question is not an immediate translation of results covered in class, it can be deduced from them with a reasonable effort.
Use them as a practice: how far are you from the answer? Do you still feel uncomfortable with some of the notions? which part could you complete quickly?

($\looparrowright$) "tortuous": this question either requires an advanced notion, a proof that is long or inventive, or it is still open.
Use them as an inspiration: can you answer any of them? does it bring you to another problem that you can answer or study further? It is recommended to work on this question only AFTER you are done with the rest!

Exercise 1: NAT (15pt, included for practice) and ICMP (0 pt) Complete the Wireshark lab for NAT. We also recommend you complete the lab for ICMP as a practice for the next exercice.

Exercise 2: Getting across a router (4 pt)

Consider the following set of packets that reside on the input ports of a router and need to be transferred to the output ports across a crossbar switch, where P_i:A$\rightarrow$B means that the packet P_i must be transferred from incoming port A to outgoing port B:

$$\mid P_1\text{: } 1\rightarrow 2 \mid P_2\text{: } 2\rightarrow 1 \mid P_3\text{: } 2\rightarrow 1 \mid P_4\text{: } 3\rightarrow 2 \mid P_5\text{: } 3\rightarrow 4 \mid P_6\text{: } 4\rightarrow 1 \mid P_7\text{: } 4\rightarrow 3$$

Assuming transfers of type A$\rightarrow$B and B$\rightarrow$A are permitted simultaneously, and that packet P_i is in front of packet P_j on an input queue whenever the incoming port is the same and $i < j$:

1. ($\rightarrowtail$) Assume packets can be processed in any order (i.e., ones at the front of a queue do not have to be processed first) What is the maximum number of transfers that can occur in the first round? Which set of packets achieve this maximum?

2. ($\rightarrowtail$)What if packets must be processed in the order of arrival. What is the maximum number of transfers that can occur in the first round? Which sets of packets achieve this maximum?

3. ($\curvearrowright$) What is the minimum number of rounds needed to transfer all packets across the crossbar? Explain the result in one sentence (i.e, how it's easy to see that fewer rounds could not be used to transfer all packets, regardless of schedule). Give an example.

4. ($\curvearrowright$) Give an example of a poor scheduling choice that maximizes the number of packets that can be sent in parallel across the crossbar each round, but where head-of-line blocking leads to additional rounds being needed beyond the minimum to forward all packets.

Exercise 3: Prefix (2 pt)

As an ISP, you have been able to purchase 6 prefixes with 24 bits: 79.59.179.0/24, 79.59.180.0/24, 79.59.181.0/24, 79.59.182.0/24, 79.59.183.0/24, and 79.59.184.0/24.

1. (↷) To advertise your addresses more efficiently, you would like to aggregate those prefixes in the smallest number of prefixes. What is the smallest number of prefixes that you can use to do so and what are these prefixes?

Exercise 4: Prefix and forwarding table (3 pt)

You operate a very simple ISP that has two customers and one router with 3 ports: port 1 connects to Columbia, port 2 connects to NYU, and port 3 connects to the rest of the Internet. For historical reasons, NYU is assigned the prefix 79.128.72/25 and Columbia receives all remaining addresses in the prefix 79.128.64/18.

1. (↣)What is the range of binary addresses assigned to Columbia? to NYU?
2. (↣)What prefix(es) does the router advertise to the rest of the Internet for Columbia and NYU to be reachable?
3. (↣)Draw and describe the content of the forwarding table in the router.

Exercise 5: Counting bits to configure a network (6 pt)

You operate the network for a 8-floor building. Each floor operates a local network that is connected to a dedicated router, and wishes to have a separate subnet. You have received from your ISP the portion of addresses 199.129.202.000/24 for the whole building.

1. (↣) Assuming that you assign address ranges of equal size to each floor, in increasing order by floor number, what is the subnet of the 3th floor?
2. (↷) How many additional PCs can this floor accomodate?

Exercise 6: Dijkstra's algorithm (3 pt)

We consider the following network, and, as a reminder, a table showing how to present the steps of Dijkstra's algorithm (on a complete different network):

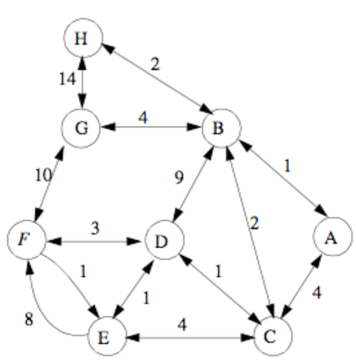

Step	N'	D(v) p(v)	D(w) p(w)	D(x) p(x)	D(y) p(y)	D(z) p(z)
0	u	7,u	3,u	5,u	∞	∞
1	uw	6,w		5,u	11,w	∞
2	uwx	6,w			11,w	14,x
3	uwxv				10,y	14,x
4	uwxvy					12,y
5	uwxvyz					

1. ($\curvearrowright$) Using Dijkstra's algorithm, compute the shortest-path tree from router F to every other router in the network. Present the result in a table where each line denotes a step of the algorithm, and the current estimation for all nodes that remain to be treated by the algorithm, as shown in the example.

Exercise 7: 3 Properties of Bellman Ford Algorithm (6 pt)

In this exercise, we consider a graph G with weighted edges in which shortest paths are computed using the Distance Vector algorithm.

Optimality rule. Shortest paths are composed of shortest paths.

1. ($\curvearrowright$) More precisely, if you consider a shortest path (*i.e.*, a path with minimum cost) $P1$ from A to B. If you assume that P1 goes through an intermediate node C, then the path from A to C used in $P1$ forms a shortest path from A to C. Prove this assertion.

Speed of convergence.

2. ($\curvearrowright$) Assume the algorithm runs in a synchronous round-by-round manner: a node whose shortest path information changes in a round i notifies all of its neighbors of the change in round $i+1$. Consider a node A whose shortest path to a node B goes through k hops, and no path from A to B with fewer than k hops is as short. Prove that if all nodes initially overestimate their shortest path distance to B in the 0_th round, then A learns its shortest path distance to B during the $(k-1)$_th round. (Hint: you may start with small values of k).

Distance vector with multiple metrics. Consider a network where each edge e is annotated with two weights, w_e and v_e, that represents each a possible "cost" metric. We assume a priori no correlation between the value of w_e and the v_e, it is possible that the edge with a large w_e has a small value v_e.

The Distance Vector protocol is implemented to compute routes to each destination, with a small enhancement to account for the two costs metrics. When a node sends to a neighbor its distance vector, the vector contains two distances per destination: the distance of the path using weights w along the edges and the distance of the path using weights v. Upon receiving this vector, the node updates its table accordingly, with two entries for each neighbor/destination, as shown in the example below.

	A	C	D
A	5,2	4,6	7,10
B	4,9	8,3	9,4
C	2,6	5,2	8,8
D	7,10	6,3	5,1

For instance, in this example, the entry (4,9) in column A, row B indicates that according to weight w, the distance to B through A will be 4 and according to weight v, this distance will be 9.

Suppose the node is required to choose a single next hop for each destination. There are many possible approaches for choosing the next hop. Below, three are considered.

RULE A All nodes select the next hop that minimizes the distance along the path using the v weight.

RULE B Some nodes select the next hop that minimizes the distance along the path using the v weight, while others use the w weight.

RULE C All nodes select the next hop that minimizes the distance along the sum of the v weights and the w weights.

For instance, given the above table, if this node N were to choose its next hop on the path to B using weight function w, it would choose node A (with distance 4). Using weight function v, it would choose node C as the next hop, and using v + w, it would also choose C (total weight 11).

3. ($\curvearrowright$) As a general motivation for this problem, can you propose at least four metrics that can be considered in reality to measure the cost of a link in a path?

4. ($\curvearrowright$) For each of the rule A-B-C above, justify either with a counter-example that forwarding loops may occur, or explain using a couple of sentences why all destinations remain reachable?

Exercise 8: Routing is expensive but testing is cheap (1 pt)

Imagine that you have to compute all the shortest paths from a source s in a graph containing N vertices and M edges. Depending on the implementation, doing this computation can take between $O(N^2)$ and $O(M \ln(N))$. You find somewhere on your machine a solution of a previous computation and you are wondering if it is actually the solution that you are looking for.

1. ($\looparrowright$) Prove that it is possible to us only $O(M)$ operations and determine if this solution is exact.

2. ($\curvearrowright$) Imagine the situation occurs regularly. You observe that most of the time the solution you found from previous computation was in fact the correct one, especially for large networks (*i.e.*, you observe that the probability that the solution is wrong is $O(1/N)$). Can you provide an algorithm that provide the solution and use only $O(M)$ operations in expectation.