

Assignment #4 – Network Layer

P. Prakash, W. An, P. Nirantar, R. Yu, X. Zhu (TAs) A. Chaintreau (instructor),

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

(\rightarrow) “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?

(\nrightarrow) “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) and ICMP (0 pt, included for practice) Complete the Wireshark lab for NAT. We also recommend you complete the lab for ICMP as a practice for the next exercise.

Exercise 2: General Multiple Choice Review ($6 \times 1 = 6$ pt)

For each question, indicate all the correct answers (there may be one or several). Correct answers will get full point even without justification. However, if you feel unsure, and would like partial credit, you may justify your choice **on the answer sheet!** (so be brief, no more than a couple of sentences).

1. TCP congestion control divides bandwidth between flows on the Internet. What following aspects makes it easier to state that TCP is ensuring a notion of fairness.
 - (a) TCP uses Additive Increase and Additive Decrease.
 - (b) Inside a period of congestion, flows that have larger windows will be subject to larger absolute window decrease.
 - (c) TCP approximates well max-min fairness, but it cannot be perfect because multiple max-min fairness allocations exist.
 - (d) Because TCP will make sure that UDP flows limit their throughputs.

Explanation: _____

2. In which case Head of line blocking will affect performance the most:

- (a) if a router receives traffic from multiple input interfaces, and forwards all of it to the same output interface.
- (b) if a router receives all traffic from the same input interface, and forwards it to multiple output interfaces.
- (c) if a router receives all traffic from the same input interface, and forwards all of it to the same output interface.
- (d) In all the above cases, Head of line blocking will affect performance in the same way.

Explanation: _____

3. What are the advantages of installing a Network Address Translator inside your home network?

- (a) Whenever a client from your network accesses a web-server s elsewhere, the presence of the NAT is transparent to this client. However, the server s your client contact should be configured in advance to work with this NAT.
- (b) NAT makes it easier for your client to connect to other client outside your network in P2P.
- (c) If your home network has a lot of users, NAT can potentially be useful to protect their privacy.
- (d) It allows you to announce BGP prefix to other autonomous domains, although you do not own the address associated with these prefixes.

Explanation: _____

4. BGP:

- (a) resembles more a Distance Vector protocol than a Link State Protocol.
- (b) always ends up choosing the shortest path to a node
- (c) both of the above
- (d) none of the above

Explanation: _____

5. For BGP:

- (a) Route reflectors are used inside a domain instead of a full mesh, to avoid that peer links are used on paths using other peer links.
- (b) Is using a Path Vector Protocol to avoid that hot potato routing creates assymetry.
- (c) Specify a routing policy that all autonomous domains are required to follow to be interconnected.
- (d) Rules are enforced to guarantee that no two customers of the same provider use a peering links to communicate between each other.

Explanation: _____

6. The following are advantages of Path Vector Protocols:

- (a) As opposed to link state, they do not use a greedy step, hence they are the only ones finding always the shortest path.
- (b) For the same reason, they always avoids loops.
- (c) For the same reason, they are easier to distribute.
- (d) For the same reason, they make it more difficult for an adversary to disrupt routing.

Explanation: _____

Exercise 3: The TCP Throughput formula (4 pt)

Consider the macroscopic description of TCP throughput. Note that, in the period of time in which the connections rate varies from $\frac{W}{2RTT}$ to $\frac{W}{RTT}$ only a single packet is lost (at the very end of the period).

- 1. (\curvearrowright) Show that the loss rate (fraction of packets lost) is equal to $L = \frac{1}{\frac{3}{8}W^2 + \frac{3}{4}W}$.

Hint: Consider the number of packets sent during a period.

- 2. (\curvearrowright) Use the result above to show that if a connection has a small loss rate L , then its average rate is approximately given by $\frac{1.22 \times MSS}{RTT \times \sqrt{L}}$.

Exercise 4: 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 :

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

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. (\rightarrow) 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. (\rightarrow) 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 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. (\rightarrow) 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. (\curvearrowright) How many additional PCs can this floor accomodate?

Exercise 6: Compatibility of NAT and Multi-Homing (4 pt)

1. (\curvearrowright) Many companies have a policy of having two (or more) routers connecting the company to the Internet to provide some redundancy in case one of them goes down. Is this policy still possible with NAT? Explain your answer.

Exercise 7: Result of a distance vector protocol (0 pt, included for practice)

Consider the network shown in the figure. Assume a distance vector routing protocol is used. We assume that the protocol works synchronously in rounds where all messages send a message to their neighbors.



1. (\rightarrow) Assume first that poison reverse is not used. Give the state table for nodes A, B, C and D after the protocol has converged (i.e., containing their current cost and next hop for each destination, along with the vectors they receive from each of their neighbors).
2. (\rightarrow) Same question with poison reverse implemented between the routers?
3. (\curvearrowright) After the protocol has converged the first time, the link C–E breaks. Which values of the tables will converge, with and without poison reverse? What if, instead, the link A–D breaks?

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. (\leftrightarrow) Prove that it is possible to use 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.