

## Communication Networks


Prof. Laurent Vanbever

**Solution:** Exercise 5 – Internet Protocol (IP) forwarding & Convergence

### Internet Protocol (IP) forwarding

#### 5.1 The Art of Defaulting Properly (Exam Style Question)

Consider this simple network configuration between ETH and Swisscom. Assume that ETH owns a large IP prefix 13.1.0.0/17, but only uses 13.1.0.0/24 to address its internal hosts. For simplicity, we assume that ETH and Swisscom operators configure their forwarding table statically and rely on the use of a default route (0.0.0.0/0).



Where are my IP packets going?

a) How many IP addressable addresses does ETH “own” in total?

**Solution:**  $2^{(32-17)} - 2$

b) Give the first and last IP address that ETH can use for addressing a host.

**Solution:** 13.1.0.1 and 13.1.127.254

c) Suppose Swisscom receives a packet for 13.1.0.66 from Deutsche Telekom. What is the path taken by this IP packet?

**Solution:** Swisscom/1 → Swisscom/2 → ETH/0 → ETH/1

d) Suppose Swisscom receives a packet for 13.1.66.1 from Deutsche Telekom. What is the path taken by this IP packet?


**Solution:** Swisscom/1 → Swisscom/2 → ETH/0 → Swisscom/2 → ETH/0 → ...

e) What eventually happens to the packet for 13.1.66.1? As an attacker observing this, could you use this observation to congest the ETH-Swisscom link more easily? Explain why (or why not).

**Solution:** It will eventually be dropped as the TTL reaches 0. Permanent forwarding loops can be used to perform a Denial of Service (DoS) attack with few resources. Here an attacker can simply start sending fake traffic to 13.1.66.1 which will start “piling up” on the Swisscom ↔ ETH link. The actual damages will depend on: *i)* the rate at which the attacker can send; *ii)* the TTL of the packets; as well as *iii)* the actual capacity of the link. Observe that the induced congestion negatively impact *all* traffic, including traffic destined to 13.1.0.0/24.

## 5.2 Changing addresses (Exam Question 2019)

Consider the network depicted in the Figure below which is composed of two hosts along with two routers, one of which acts as Network Address Translator (NAT). Host 1 is located in a private subnet (192.168.1.0/24) and uses 192.168.1.1 as gateway, while host 2 is located in a public subnet (81.0.0.0/24) and uses 81.0.0.1 as gateway. The Figure below also depicts the MAC address of each of the 6 interfaces connected at either end of the three links. The NAT/router performs address translation between the private and the public subnets, translating traffic originating from private IPs to its public one (here, 20.0.0.1), and vice-versa.



A network topology relying on Network Address Translation.

- a) Consider that host 1 tries to open a TCP connection with host 2 on port 80 using 1337 as (random) source port. Write down a possible sequence of packet headers observed at each link for the first two packets (i.e., the SYN sent by host 1, and the SYN/ACK sent by host 2). Fill in the table below to answer. Assume that hosts and routers have the required MAC addresses in their ARP table.

**Solution:**


	src MAC	dst MAC	src IP	dst IP	src TCP port	dst TCP port
link A	00:A1	00:A2	192.168.1.10	81.0.0.2	1337	80
link B	00:B1	00:B2	20.0.0.1	81.0.0.2	rand	80
link C	00:C1	00:C2	20.0.0.1	81.0.0.2	rand	80
link C	00:C2	00:C1	81.0.0.2	20.0.0.1	80	rand
link B	00:B2	00:B1	81.0.0.2	20.0.0.1	80	rand
link A	00:A2	00:A1	81.0.0.2	192.168.1.10	80	1337

- b) Could host 2 initiate a TCP connection to host 1? Briefly explain why/why not.

**Solution:** That is not possible as host 1 is behind a NAT. Packets that reach the public IP of the NAT will not be forwarded to host 1 as there is no corresponding NAT rule available.

# Convergence

## 5.3 Convergence with Poisoned Reverse



Consider the network on the left which uses distance vector routing with poisoned reverse. Each link is associated with a weight that represents the cost of using it to forward packets. Link weights are bidirectional. Assume that the link between X and A fails (as shown in the figure) and use the table below to show the first 8 steps of the convergence process. How many steps does it take until the network has converged to a new forwarding state? Explain your observations.

**Solution:** The network does not converge as the maximum link weight is increased by one in each round ("count to infinity problem"). Poisoned reverse does not solve the problem of counting to infinity if three or more nodes are involved. One possible workaround is to define  $\infty$  as a small value (e.g.  $\infty := 16$ ).

**Solution:**

dst=A	X			Y		Z	
	via A	via Y	via Z	via X	via Z	via X	via Y
$t = 0$ before the failure	1	$\infty$	$\infty$	2	3	2	3
$t = 1$ after X sends its vector	★	$\infty$	$\infty$	$\infty$	3	$\infty$	3
$t = 2$ after Y sends its vector	★	4	$\infty$	$\infty$	3	$\infty$	$\infty$
$t = 3$ after Z sends its vector	★	4	$\infty$	$\infty$	$\infty$	$\infty$	$\infty$
$t = 4$ after X sends its vector	★	4	$\infty$	$\infty$	$\infty$	5	$\infty$
$t = 5$ after Y sends its vector	★	$\infty$	$\infty$	$\infty$	$\infty$	5	$\infty$
$t = 6$ after Z sends its vector	★	$\infty$	$\infty$	$\infty$	6	5	$\infty$
$t = 7$ after X sends its vector	★	$\infty$	$\infty$	$\infty$	6	$\infty$	$\infty$
$t = 8$ after Y sends its vector	★	7	$\infty$	$\infty$	6	$\infty$	$\infty$


Add the distance vectors to this table

## 5.4 Convergence (Exam Style Question)

Consider this simple network running OSPF as link-state routing protocol. Each link is associated with a weight that represents the cost of using it to forward packets. Link weights are bi-directional.

Assume that routers A, B and D transit traffic for an IP destination connected to C and that link  $(B, C)$  fails. Which nodes among A, B and D could potentially see their packets being stuck in a transient forwarding loop? Which ones would not?

**Solution:** Nodes A and B could see their packets stuck in a forwarding loop if B updates its forwarding table before A, which is likely to happen as B would be the first to learn about an adjacent link failure. On the other hand, D would not see any loop as it uses its direct link with C to reach any destination connected beyond it.




Assume now that the network administrator wants to take down the link  $(B, C)$ , *on purpose*, for maintenance reasons. To avoid transient issues, the administrator would like to move away all traffic from the link *before* taking it down and this, without creating any transient loop (if possible). What is the minimum sequence of increased weights setting on link  $(B, C)$  that would ensure that *no packet* destined to C is dropped?

**Solution:** One example of a minimum sequence of weight settings is [1, 3, 5].

*Note:* The problem highlighted above happens because B shifts traffic to A before A shifts traffic to D, hence creating a forwarding loop. By setting the  $(B, C)$  link weight to 3, (only) A shifts from using  $(A, B, C)$  to using  $(A, D, C)$ . Once A has shifted, it is safe to shift B by setting the link weight to 5 (or higher). Once B has shifted as well, the link can be safely torn down.

## 5.5 Convergence of General Topologies



Consider the T-shaped network on the left and a synchronous version of the distance vector algorithm. Suppose that at each iteration, a node exchanges its distance vectors with its neighbors and receives their distance vectors. Assume that the algorithm runs for the first time, and that each node only knows the costs to its immediate neighbors. Derive a formula for the maximum number of iterations required before the distributed algorithm converges.

### Solution:

At each iteration, a node exchanges distance tables with its neighbors. Thus, if you are node A, and your neighbor is B, all of B's neighbors (which will be one or two hops from you) will know the shortest cost path of one or two hops to you after one iteration (i.e. after B tells them its cost to you).

Let  $d$  be the diameter of the network — the length of the longest shortest path without loops between any two nodes in the network. Using the reasoning above, after  $d - 1$  iterations, all nodes will know the shortest path cost of  $d$  or fewer hops to all other nodes.

For our T-shaped network, diameter  $d$  depends on  $n$ , and  $m$ , respectively. For example, one possible diameter passes from the top left node to the bottom right node.

$$d = \max\{n, m, \lceil 0.5n \rceil + m + 1\}$$

Note that this bound only applies when the algorithm runs for the first time. As we have seen in the lecture, if a link changes after the network has already converged, no upper bound exists due to the count-to-infinity problem.