Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Networked Systems
ETH Zürich — seit 2015

Communication Networks

Prof. Laurent Vanbever

Exercise 5 – Forwarding & Routing

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).

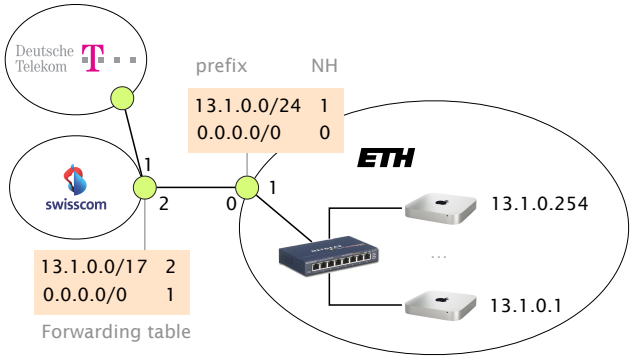

Where are my IP packets going?

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

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

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

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

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).

5.2 Convergence (Exam Style Question)

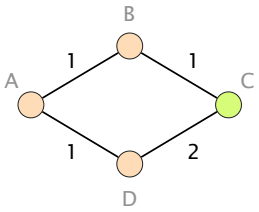

Loopy or not?

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?

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?

5.3 Convergence with Poisoned Reverse

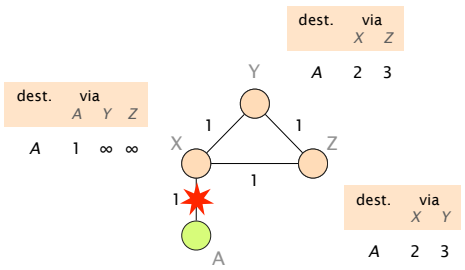

Consider the network on the left which uses distance vector routing with poisoned reverse. Each link is associated with a weight that represents the cost of using it to forward packets. Link weights are bi-directional.

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.

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

Add the distance vectors to this table