


## Communication Networks

Prof. Laurent Vanbever

### Solution: Exercise 6 – Border Gateway Protocol (BGP)

#### 6.1 Traffic Engineering

Assume that ETH has only one prefix: 82.130.64.0/21. As depicted on the left, the ETH network is connected to three providers (Swisscom, Deutsche Telekom and Switch) and the providers are interconnected with each other. The contract with Swisscom is the cheapest one (indicated by the dollar symbols). For this reason, ETH wants to receive all the incoming traffic over the Swisscom link and therefore announces its prefix only to Swisscom.



ETH is connected to three providers with different costs.

- a) Do you think that is a good configuration? What happens if the link between ETH and Swisscom fails?

**Solution:**

Not a good solution. If the link fails, ETH will no longer receive any traffic. ETH is no longer reachable from other networks.

- b) To improve the connectivity in case of a link failure between ETH and Swisscom, ETH wants to optimize its announcements. Write down the prefixes which ETH announces to Swisscom, Deutsche Telekom and Switch. During normal operation (no link failure) ETH should still receive all incoming traffic over the Swisscom link.

**Solution:**

To Swisscom: 82.130.64.0/22 and 82.130.68.0/22 (other splits are also possible)

To Deutsche Telekom: 82.130.64.0/21

To Switch: 82.130.64.0/21

- c) After further investigations, ETH decides that only traffic towards 82.130.68.0/23 has to be received over the Swisscom link. All the other traffic can enter over any of the providers. Which prefixes do you have to announce to achieve this traffic distribution?

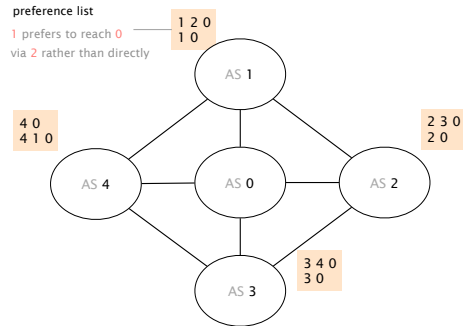
**Solution:**

To Swisscom: 82.130.68.0/23 and 82.130.64.0/21

To Deutsche Telekom: 82.130.64.0/21

To Switch: 82.130.64.0/21

## 6.2 Convergence



Does this network ever converge?

Consider this BGP network composed of 5 ASes. Each AS is assigned with a list of paths which indicates its preferences to reach AS 0.

Considering that only AS 0 originates prefixes, does that BGP network have a unique, stable solution?

- If yes, indicate the path that each AS selects in the stable solution.
- If not, describe an example of oscillation. For instance, by describing a sequence of messages that repeats itself.

**Solution:** This BGP network does have a unique, stable solution in which:

- AS 1 selects [1,2,0] (preferred path);
- AS 2 selects [2, 0];
- AS 3 selects [3, 4, 0] (preferred path);
- AS 4 selects [4, 0] (preferred path).

## 6.3 Not-so-reliable Internet

Consider now the same BGP network composed of 5 ASes but assuming customer-provider and peer-to-peer policies. Providers are connected to their customers with a single-headed arrow pointing to their customers (AS 1 is the provider of AS 4), while peers are connected with double-headed arrows (AS 1 and AS 2 are peers).

Assume that AS 2 is the only one to advertise an IPv4 prefix: 82.130.64.0/21 (to *all* its neighbors) and that the Internet has converged. Which BGP messages are exchanged after the following events happen, one after the other:

- a) the link between AS 0 and AS 2 fails (event 1)

**Solution:**

- (i) AS 0 sends a WITHDRAW for 82.130.64.0/21 to AS 3 (optional);
- (ii) AS 0 sends an UPDATE for 82.130.64.0/21 to AS 3 with AS-PATH [0,1,2].

- b) the link between AS 1 and AS 4 fails (event 2)


**Solution:**

AS 4 sends a WITHDRAW for 82.130.64.0/21 to AS 3.

- c) the link between AS 1 and AS 2 fails (event 3)

**Solution:**

- (i) AS 1 sends a WITHDRAW for 82.130.64.0/21 to AS 0;
- (ii) AS 0 sends a WITHDRAW for 82.130.64.0/21 to AS 3;



Which messages are exchanged?

Is the network still connected at the end? If not, list the ASes that cannot reach the prefix anymore.

**Solution:**

No. The BGP network is not connected anymore. Only AS 3 is able to reach 82.130.64.0/21 via its direct link with AS 2. Observe that the physical graph is still connected yet as BGP policies prevent paths to be used, blackholes appear nonetheless.