

## Communication Networks

Prof. Laurent Vanbever

### Exercise 99 – Additional Practice

## Link State and Distance Vector


### 99.1 Warm-up Questions (Exam Question 2016)

For the following statements, decide if they are *true* or *false*. Motivate your decision.


- a) Consider a positively weighted graph  $G$ . Applying the Bellman-Ford (used by distance-vector protocols) or Dijkstra (used by link-state protocols) algorithm on  $G$  would lead to the same forwarding state.
- b) Link-state protocols (such as OSPF) are guaranteed to compute loop-free forwarding state as long as the link-state databases are consistent on all routers.
- c) Link-state protocols (such as OSPF) require routers to maintain less state than distance-vector protocols (such as RIP).
- d) Poisoned reverse solves the problem of count-to-infinity.
- e) Consider a positively weighted graph  $G$ . Multiplying all link weights by 2 would change the all-pairs shortest paths computed by the Dijkstra algorithm on  $G$ .
- f) Consider a positively weighted graph  $G$ . Adding 1 to all link weights would change the all-pairs shortest paths computed by the Dijkstra algorithm on  $G$ .

## 99.2 A very long downtime (Exam Question 2019)

Consider the network in the Figure below which connects a large cluster of 500k servers (right) with a large pool of clients (left). The network uses OSPF internally. Each of the 500k servers is located in a different IP subnet, while all of the clients are located in the same /8 subnet.



Consider that the link between router A and C fails. As router A is adjacent to the failure, it immediately detects it and starts rerouting the traffic for the 500k prefixes to B. The Figure below reports the evolution of the throughput observed for 3 of the 500k server prefixes: 23.0.0.0/24, 121.0.0.0/24, and 202.0.0.0/24. One can see that the downtime experienced by each prefix is different: 23.0.0.0/24 experiences less downtime than 121.0.0.0/24, which itself experiences less downtime than 202.0.0.0/24.




- Explain why the downtimes experienced by the three prefixes differ.
- Describe two solutions to speed up the convergence for these three prefixes.

## 99.3 Reverse Dijkstra (Exam Question 2024)

You are looking at a drawing of your network but it clearly seems some links are missing. You also have a table which shows the *partial* output of the Dijkstra's algorithm (some cells are empty) as performed in the lecture starting from the node U. For each iteration, the table indicates the weights of the shortest paths found so far. If after one iteration there are multiple nodes with an equally-shortest path, the algorithm continues with the node which comes first in the alphabet. There is at most one link between two nodes and each link has a strictly positive weight  $[1, \infty]$ .

You feel confident that, with the information from the network drawing and the partial output from Dijkstra's algorithm in the table below, you can find some of those missing links and link weights.

- a) First, fill in the highlighted cells in the table below using the information available in the network graph. Write your answers directly in the table.



A network consisting of 10 nodes with missing links and link weights

#	U	A	B	C	D	E	F	G	H	I
1	0	1		$\infty$			$\infty$	$\infty$	$\infty$	$\infty$
2	0	1	5	2	3	3	$\infty$	$\infty$	$\infty$	$\infty$
3	0	1	4	2	3	3	$\infty$	$\infty$	$\infty$	$\infty$
4	0	1	4	2	3	3		$\infty$	8	$\infty$
5	0	1	4	2	3	3	5	$\infty$	8	$\infty$
6	0	1	4	2	3	3	5	10	8	$\infty$
7	0	1	4	2	3	3	5	9	8	$\infty$
8	0	1	4	2	3	3	5	9	8	13
9	0	1	4	2	3	3	5	9	8	12
10	0	1	4	2	3	3	5	9	8	12

For each iteration (1 to 10) the table shows the shortest path found by Dijkstra's algorithm performed on node U towards all other nodes.

**b)** Now that you have completed the table you can use this information to piece together the missing links in the network above. Draw all the links and corresponding weights you can identify directly.

**c)** You now wonder whether you actually found all the links or not. For all the following links, mark if they could exist or not.

Use the following format:

- If the link could exist, indicate the range of weights this link is allowed to have;
- If the link could not exist, write down the iteration in which the link would have showed up.

Link U-C: \_\_\_\_\_

Link B-E: \_\_\_\_\_

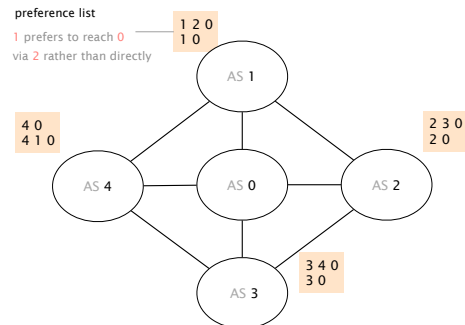
Link U-I: \_\_\_\_\_

Link A-D: \_\_\_\_\_

Link E-G: \_\_\_\_\_

# BGP

## 99.4 Convergence



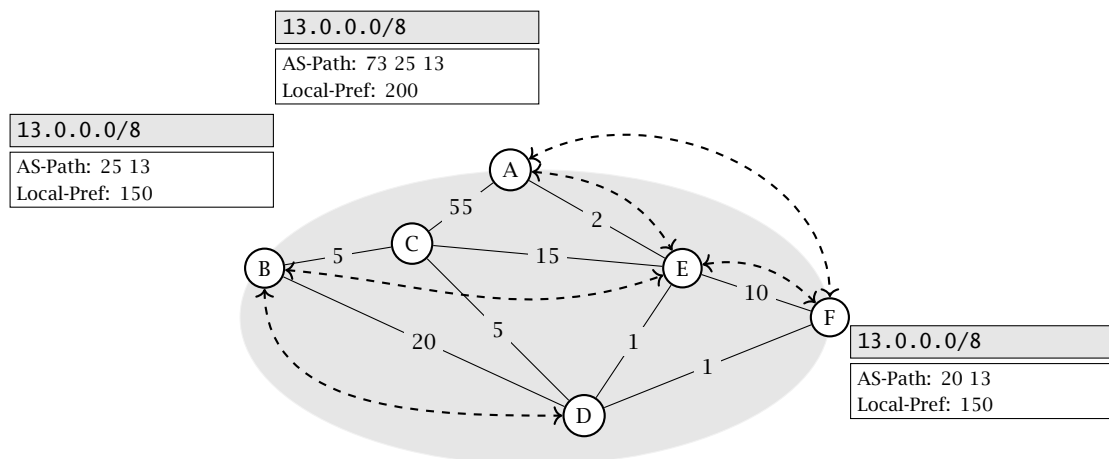
Does this network ever converge?

Consider this BGP network composed of 5 ASes. Assume that each AS has configured its BGP policies in a way that leads to the preference lists shown in the figure. For example, AS 1 is configured to only accept an announcement for AS 0 if it has path [1,2,0] or [1,0]. In addition, AS 1 prefers the path [1,2,0] over the path [1,0].

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.

## 99.5 BGP and IGP: Very creative! (Exam Question 2020)



A simple BGP network **not** forming an iBGP full-mesh.

Consider the AS above with three border routers (A, B, F) and three internal routers (C, D, E). All three border routers receive a route announcement for the prefix 13.0.0.0/8 from their eBGP neighbors (not depicted), which they distribute internally. The iBGP sessions are depicted by double-headed dashed arrows (e.g., router A and F maintain an iBGP session). All routers follow the standard BGP decision process. The three border routers have `next-hop-self` configured on all iBGP sessions.

- a) For every router, list (i) the BGP next-hop, (ii) the path taken by the traffic and (iii) indicate whether the router's traffic can actually reach the destination. If the next-hop is external, put EXT. If there is no next-hop, put NO.

Router	BGP next-hop	Path taken by the traffic	Reachable?
A	EXT	A → EXT	Yes
B			
C	NO	C → ∅	No
D			
E			
F			


- b) Assume the eBGP session of router A fails and consequently, **the external route of A is not available anymore**. List for every router (i) the BGP next-hop, (ii) the path taken by the traffic and (iii) indicate whether the router's traffic can reach the destination. If the next-hop is external, put EXT. If there is no next-hop, put NO.

Router	BGP next-hop	Path taken by the traffic	Reachable?
A			
B			
C			
D			
E			
F			

- c) The network operator reacted and **added a new iBGP session between routers B and C**. The failure still persists, i.e., the external route of A is not available. List for every router (i) the BGP next-hop, (ii) the path taken by the traffic and (iii) indicate whether the router's traffic can reach the destination. If the next-hop is external, put EXT. If there is no next-hop, put NO.

Router	BGP next-hop	Path taken by the traffic	Reachable?
A			
B			
C			
D			
E			
F			

## 99.6 Traffic Engineering



ETH is connected to three providers with different costs.

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.

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


## Reliable Transport

### 99.7 TCP Warm-up (Exam Question 2019)

For the following questions answer either with true or false and give an explanation for your decision. Note that in the actual exam you only had to pick true, false or do not give an answer at all (as wrong answers result in point deductions).

- a) In contrast to the GBN protocol used in the project, TCP's sequence number often increases by more than one between two consecutive data packets.
- b) A client having an ongoing TCP connection to a server (IP 1.2.3.4, port 80) is not able to start a second TCP connection towards 1.2.3.4:80.
- c) A TCP packet with a value of 9 in its header length field (HdrLen, sometimes also called data offset) indicates 16 bytes of TCP options.
- d) Consider an ongoing TCP flow. Whenever the congestion window increases, the sender can transmit additional data segments.

For the following four questions, we consider the congestion window (CWND) evolution observed for a flow  $f$  and depicted in the figure below.




- e) Flow  $f$  was in the slow start phase exactly twice.
- f) Flow  $f$  experiences at least one packet loss between time A and B.
- g) In the future,  $f$  will never be able to experience a higher CWND size than B.
- h) Consider another flow  $f_2$  starting at exactly the same time as  $f$  and traversing the exact same path, then  $f_2$  would experience the exact same CWND evolution.



## 99.8 Congestion Window

Consider the following plot which depicts the evolution of the size of the TCP congestion window of the sender.



What kind of network conditions is this flow seeing?

Describe briefly:


- What happens at point B?
- Does the event happening at point B require the network to discard packets? Why or why not?
- What happens at point E?
- Does the event happening at point E require the network to discard packets? Why or why not?

Consider that the Maximum Segment Size (MSS) of the connection is 1 kB and the Round-Trip Time (RTT) between the two end points is 100 milliseconds. The sender opens the connection at time  $t = 0$ . Transmission delay in this network is negligible, so you should only consider the propagation delay in the following.

- How much time has elapsed at point A?
- How much time has elapsed *between* point C and D?
- How much time has elapsed *between* point F and point G?

Briefly explain how come point D is higher than point B. Would you expect this to happen often?

## 99.9 Fairness



Are you getting a fair share?

Consider the situation in which two hosts, A and B, are concurrently using a 1 Mbps link with a Maximum Segment Size (MSS) of 100 kb.

Assuming that B starts with 500 kbps and A with 200 kbps (see left picture). Describe the evolution of the throughput of the two hosts when:

- A and B rely on Additive Increase Multiply Decrease (AIMD).
- A and B rely on Multiplicative Increase Additive Decrease (MIAD).

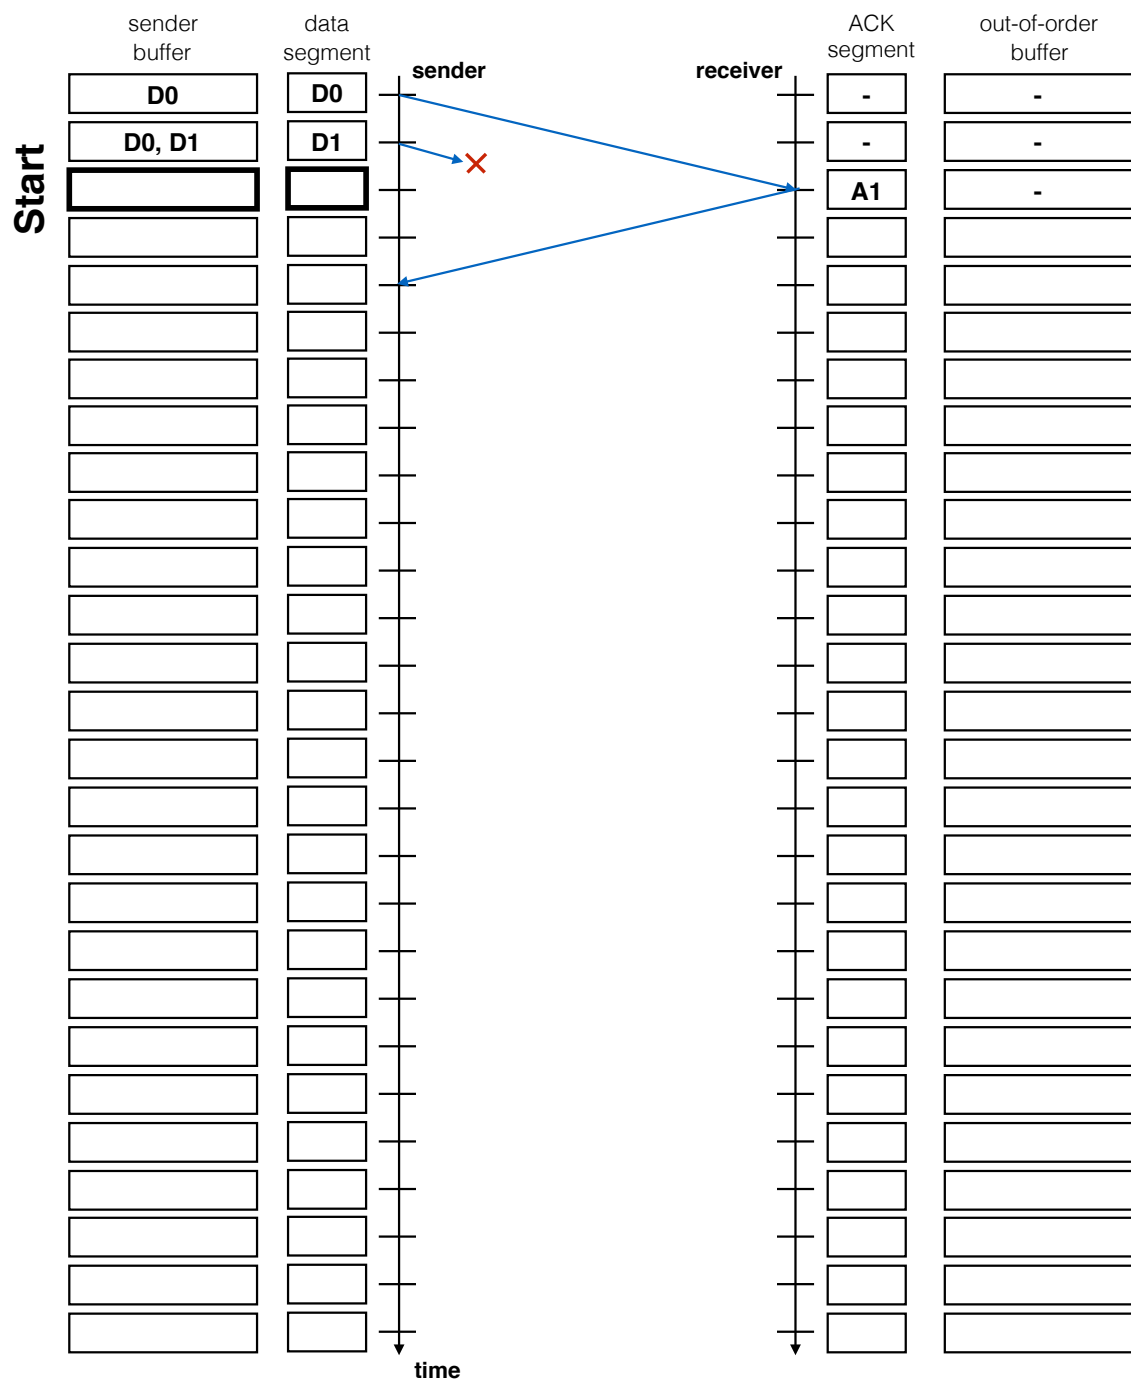
Assume now that only A is malicious, and wants to cheat congestion control to get more throughput. Describe two distinct ways A could do so and what would be the net effect on B's throughput.

### 99.10 Go-Back-N (Exam Question 2017)

Consider a Go-Back-N (GBN) protocol with the following implementation choices for sender and receiver.

- The sender and receiver window have a size of 4;
- The receiver saves out-of-order segments in an (infinite) buffer and removes them as soon as the missing segment(s) arrive;
- The receiver uses cumulative ACKs which acknowledge all previous segments and point to the next expected data segment;
- The sender uses Fast Retransmit. After three duplicate ACKs, the sender immediately retransmits the corresponding data segment. For instance, if the sender gets the following ACKs [A1, A1, A1], it will immediately retransmit the data segment D1;
- For each tick in the diagram below, the sender can send one data segment and the receiver can send one ACK. Sender and receiver will first analyze the incoming packet and then send a data segment/ACK;
- The sender uses a retransmission timer of 5 ticks. Each time it sends a data segment or receives an ACK, the timer is reset. After a timeout, the sender retransmits all current segments in its sender buffer (in order, one segment per tick);
- A data segment or ACK needs two ticks to travel to the other end of the connection. See the given start in the diagram.

**Your task:** Use the diagram on the next page to draw the successful transmission of 6 data segments (D0 to D5) if the **first data segment** (D1, already indicated) **is lost** as well as **ACK A5 is lost the first time it is sent**. For each tick, indicate which data segment or ACK is transmitted (if any) as well as the content of the sender and out-of-order buffer.



## 99.11 Reliable Transport<sup>2</sup> (Exam Style Question)

On the next page you see the beginning of a communication between two end-points using the Go-Back-N protocol with Selective Repeat. Consider that the sender has infinitively many data segments to send and they are immediately available.

We ask you to fill in the missing values in the two tables. Stop if you either reach the bottom of the tables or the sender is no longer able to send new data segments because its buffer is full. Start with the blue row indicated on the left.

**Note: Please read the entire question carefully!**

*Set-up:*

- Every table row corresponds to one time-slot. The sender and receiver can send one data segment respectively ACK segment in every time-slot;
- Consider that the *Sender buffer* contains all the sent but not yet acknowledged segments, while the *Out-of-order buffer* contains all the messages which have been received...out-of-order;
- If the sender receives an ACK in one time-slot, it first processes the ACK (e.g. removes segments from the sender buffer) and then sends the data segment for this time-slot. Similarly, the receiver will first analyse the received data segment and then send a corresponding ACK;
- The link between the sender and receiver is not reliable. The first data segment with a sequence number of 3 and all data segments with a sequence number of 5 are dropped and do not reach the receiver.


*Sender behavior:*

- The sender uses Selective Repeat after receiving 3 duplicate ACKs. That means as soon as the sender receives an ACK with the same sequence number for the third time, it will retransmit the missing segment in the same time-slot (instead of a new data segment);
- The sender can store at most 5 unacknowledged segments in its sender buffer.

*Assumptions:*

- You will not reach the maximal sequence number. No overflow;
- The timeout value is very long and will not occur;
- The receiver out-of-order buffer can store an unlimited number of segments.

### 99.12 Understanding Go-Back-N's Behavior (Exam Style Question)



- a) The sender wants to transmit 10 data segments  $(0, \dots, 9)$  to the receiver. Assume that *exactly* one segment is lost. How many segments has the sender to transmit in the best (resp. worst) case? For each case, indicate which segment was lost.
- b) Once again, the sender wants to transmit 10 data segments  $(0, \dots, 9)$  to the receiver. This time, assume that exactly one ACK is lost. How many segments does the sender have to transmit in the best (resp. worst) case and which ACK was lost?

- c) Assume the sender just transmitted segments 4, 5, 6 and 7 and is now waiting for ACKs from the receiver. It receives three times an ACK with number 4. Therefore, it cannot remove segments from its buffer and eventually the timeout is reached. Following the GBN protocol, the sender will retransmit all four segments.

A friend of yours explains that she improved her GBN algorithm so that, in the case above, the sender would just retransmit data segment 4 (instead of all four segments). She tells you that, quite often, she would then get an ACK with number 8 back (all four packets were successfully transmitted).

Can you explain why your friend believes that only data segment 4 was missing? Under which network conditions would the proposed improvement *not* work (assuming you still get three times an ACK with number 4)?

# IP

## 99.13 IPv6 Calculations

IPv6 addresses have a slightly different notation. Because the addresses are 128 bit long, we switch from a decimal notation to a hexadecimal one. In general, IPv6 addresses are represented by eight colon-separated blocks of up to four hexadecimal digits each. In a block, leading zeros can be omitted. Furthermore, we can use the “::” symbol to compress one or more consecutive zero blocks. However, the “::” symbol can only be used once in a single IPv6 address. As an example, the IPv6 address: 2001:0db8:0000:0000:0000:ff00:0042:8329 can be simplified to 2001:db8::ff00:42:8329.

- a) You are the operator of an enterprise network. Your ISP is giving you a /96 subnet 2001::/96. How many addresses do you have available? Is that a reasonable subnet size compared with the currently available IPv4 addresses?
- b) In your enterprise network, each host machine is identified by a unique ID starting from 1. Now that you have a lot of IPv6 addresses available, you decide to give each host a unique IP address. The host with ID 1 gets the first IPv6 address in your subnet (2001::1), the host with ID 2 the second IP address and so on. Complete the following table:

IPv6 address	host ID
	5
	14
2001::3A5	
	4 294 967 295
2001::3:0	

## 99.14 Detective work

You just started your first job as a network operator of a small network. To get more familiar with the network, you look at a packet trace captured at a switch. The trace contains packets from multiple hosts and one router connected by a (layer 2) switch. The router acts as default gateway, providing access to the Internet and is assigned the first IP address in the subnet. Each row in the following table represents one packet observed at the switch.



SRC MAC Address	DST MAC Address	SRC IP Address	DST IP Address
6a:00:02:49:a1:a0	11:05:ab:59:bb:02	65.222.11.1	65.222.8.2
6a:00:02:49:a1:a0	da:15:00:00:01:11	65.222.11.1	65.222.16.1
da:15:00:00:01:11	11:05:ab:59:bb:02	129.132.103.40	65.222.8.2
11:05:ab:59:bb:02	40:34:00:7a:00:01	65.222.8.2	65.222.15.254
11:05:ab:59:bb:02	ac:00:0a:aa:10:05	65.222.8.2	65.222.9.99
ac:00:0a:aa:10:05	01:05:3c:34:00:02	65.222.9.99	65.222.13.255
6a:00:02:49:a1:a0	da:15:00:00:01:11	65.222.11.1	65.222.8.1

- a) Can you identify all the hosts that are part of the local network?
- b) Can you reconstruct the IP subnet used to address the hosts within that local network?

# DNS

## 99.15 Curious students

Consider that ITET has a local DNS server serving the DNS requests for all students' devices connected in the department. How could you determine if an external website has been visited recently by a fellow colleague of yours? Explain.

## 99.16 Name it or Route it: pick one



Any clear winner?

In the course, we saw two ways to replicate and load-balance content: *(i)* using Anycast routing; or *(ii)* using DNS.

a) List and briefly justify three pros and cons of each.

b) We saw that CDNs often rely on DNS for distributing their load. Could they also use Anycast routing instead? Explain why or why not.

## 99.17 Multiple answers

Whenever a client (e.g., your computer) receives multiple IP addresses as answer to a DNS lookup, it picks the very first one. Only if that one does not work, it tries the next one in order.

When you run `dig yahoo.com`, you receive multiple IP addresses as an answer compared to, for example, `dig google.com`.

Can you think of a reason for providing multiple IP addresses? Run the lookup for `yahoo.com` multiple times.