ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institut für Technische Informatik und Kommunikationsnetze

Department ITET
August 2019

Prof. Dr. Laurent Vanbever
R. Birkner, T. Bühler, A. Dietmüller, T. Holterbach, R. Meier

Exam: Communication Networks

21 August 2019, 14:00–16:30, Room HIL D 15

General remarks:

- Write your **name** and your **ETH student number** below on this front page and **sign it**.
- Put your **legitimation card** on your desk.
- Check if you have received **all task sheets** (Pages **1 - 29**).
- Do **not separate** the **task sheets**.

- Write your answers directly on the task sheets.
- **All answers fit within the allocated space and often in much less.**
- If you need more space, please use your own extra sheets, in which case use a **new sheet of paper** for **each task** and write your name and the exam task number in the **upper right corner**.

- **Read each task completely before you start solving it**.
- **For the best mark, it is not required to score all points.**

- Please answer either in **English or German**.
- **Write clearly** in blue or black ink (not red) using a **pen**, not a pencil.
- **Cancel** invalid parts of your solutions **clearly**.
- At the end of the exam, hand your **solutions in together with all extra sheets**.

Special aids:

- All written materials (vocabulary books, lecture and lab scripts, exercises, etc.) are allowed.
- Using a calculator is allowed, but the use of electronic communication tools (mobile phone, computer, etc.) is strictly forbidden.

Family name: Student legi nr.:

First name: Signature:

Do not write in the table below (used by correctors only):

Task	Points	Sig.
Ethernet & IP	/27	
Intra-domain routing	/26	
Inter-domain routing	/38	
Reliable transport	/27	
Applications	/32	
Total	/150	

leave blank

Task 1: Ethernet & IP **27 Points**

a) Warm-up **(6 Points)**

leave blank

For the following true/false questions, check either *true*, *false* or nothing. For each question answered correctly, one point is added. For each question answered incorrectly, one point is removed. There is always one correct answer. This subtask gives at least 0 points.

true ☐ false ☐ When an Ethernet switch sends an IP packet on a port, it rewrites the destination MAC to the MAC address of the next hop.

true ☐ false ☐ If there are errors in the routing and/or forwarding tables of some routers, then some IP packets might loop forever.

true ☐ false ☐ Assume two layer-2 networks N_1 and N_2 interconnected by an IP router. Hosts in N_1 have an IP address in the private subnet 192.168.0.0/16 while hosts in N_2 have an IP address in the public subnet 12.34.56.0/24. Hosts in N_2 *cannot* communicate with hosts in N_1.

Assume that A and B are two IPv4 hosts connected to the same VLAN. Assume that A's interface is configured with 10.10.1.1/16, while B's interface is configured with 10.10.0.2/24.

true ☐ false ☐ If A sends an IP packet to B, and if the ARP cache at A is empty, then A will first send an ARP packet in order to determine the MAC address of the next-hop router.

true ☐ false ☐ If B sends an IP packet to A, and if the ARP cache at B is empty, then B will first send an ARP packet in order to determine the MAC address of the next-hop router.

true ☐ false ☐ If A communicates with B, the forward and reverse path are identical.

b) Morphing trees **(5 Points)**

leave blank

Consider the switched Ethernet network with 8 switches in Figure 1. Each switch runs the Spanning Tree Protocol (STP), as seen in the lecture. Each link has a unary cost and switch 1 acts as the root switch.

(i) Assume that all the switches boot at exactly the same time, and exchange messages with each other every 3 seconds. Assuming that messages propagate instantaneously, how long does it take for every switch in the topology to learn the shortest path to reach switch 1? Briefly explain your answer. (2 Points)

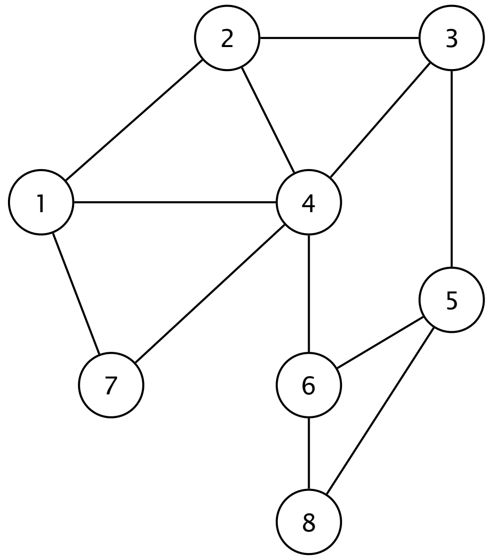

Figure 1: A layer-2 network topology using the Spanning Tree Protocol.

(ii) Draw a topology with the same number of nodes (8) and links (12) that minimizes the time it takes for every switch to learn the shortest path to switch 1. Make sure that your topology is redundant: it should remain connected after any single link failure. Draw your answer directly in Figure 2. (3 Points)

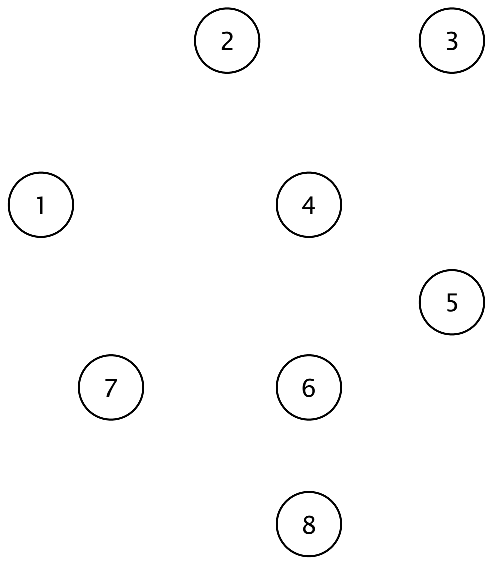

Figure 2: Your redesigned topology which minimizes learning time.

leave blank

c) Changing addresses **(8 Points)**

Consider the network depicted in Figure 3 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. Figure 3 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.

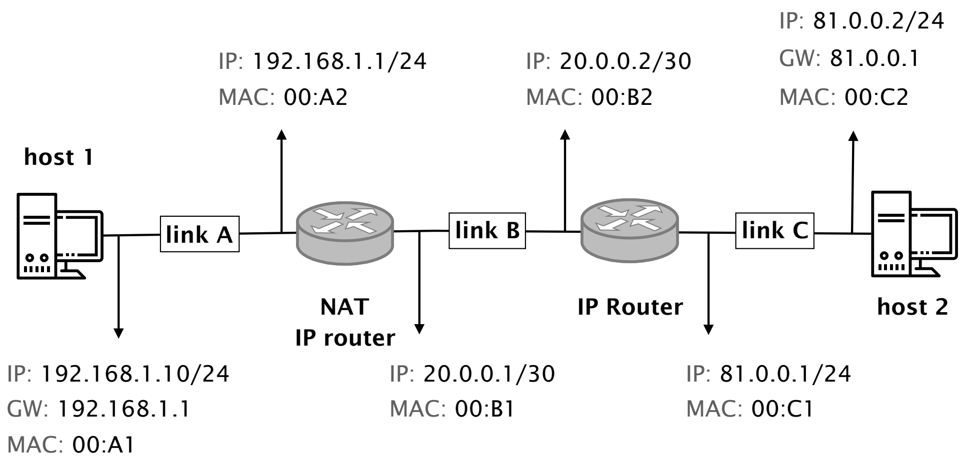

Figure 3: A network topology relying on Network Address Translation.

(i) 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 (rest) of the table below to answer. Assume that hosts and routers have the required MAC addresses in their ARP table. (6 Points)

	src MAC	dst MAC	src IP	dst IP	src TCP port	dst TCP port
link A						
link B						
link C						
link C						
link B						
link A						

(ii) Could host 2 initiate a TCP connection to host 1? Briefly explain why/why not. (2 Points)

d) Less is (often) more **(8 Points)**

leave blank

Consider an IP router with a forwarding table composed of the 16 entries depicted on the left. You can assume that any single IP packet this router will ever receive will be matched by exactly one of these entries. Write down (on the right) an equivalent forwarding table by combining entries together into shorter ones such that the resulting table has the least number of entries possible. Your reduced forwarding table must necessarily contain the default route (0.0.0.0/0) for which you also need to indicate the corresponding next hop.

Hint: The answer requires 8 lines or less.

prefix	next hop
80.160.0.0/16	1
80.161.0.0/16	1
80.162.0.0/16	5
80.163.0.0/16	1
80.164.0.0/16	1
80.165.0.0/16	6
80.166.0.0/16	6
80.167.0.0/16	1
80.168.0.0/16	2
80.169.0.0/16	2
80.170.0.0/16	2
80.171.0.0/16	2
80.172.0.0/16	1
80.173.0.0/16	4
80.174.0.0/16	3
80.175.0.0/16	3

prefix	next hop
0.0.0.0/0	

leave blank

leave blank

Task 2: Intra-domain routing **26 Points**

a) Warm-up **(5 Points)**

For the following true/false questions, check either *true*, *false* or nothing. For each question answered correctly, one point is added. For each question answered incorrectly, one point is removed. There is always one correct answer. This subtask gives at least 0 points.

true ☐ false ☐ Increasing the frequency of OSPF Hello messages helps preventing forwarding loops upon a link failure.

For the following three questions, consider the network in Figure 4 which uses a distance vector protocol without poisoned reverse.

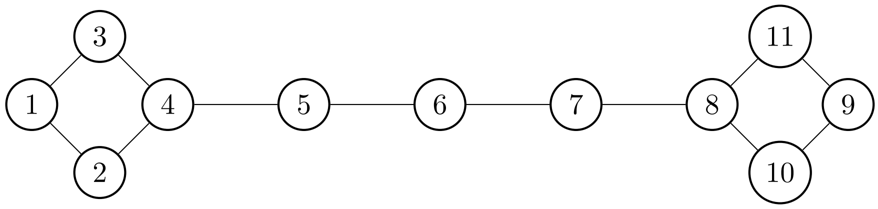

Figure 4

true ☐ false ☐ An infinity value of 4 would result in a better connectivity compared to an infinity value of 64.

true ☐ false ☐ One could enable load balancing to increase the throughput between node 4 and 8.

true ☐ false ☐ It is possible to configure the link weights such that the traffic from node 11 to node 1 does not traverse node 2 whereas the traffic from node 10 to node 1 does not traverse node 3.

For the final question, consider the network in Figure 5 which uses a distance vector protocol with poisoned reverse.

true ☐ false ☐ A failure of the link (1-2) results in a routing loop.

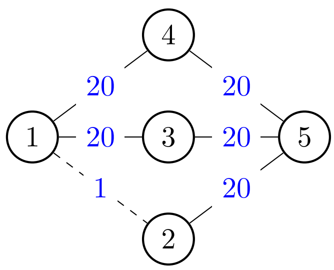

Figure 5

b) Looping packets **(9 Points)**

leave blank

Consider the following topology where OSPF is configured to route packets. In this exercise, we will focus on the traffic destined to the node d.

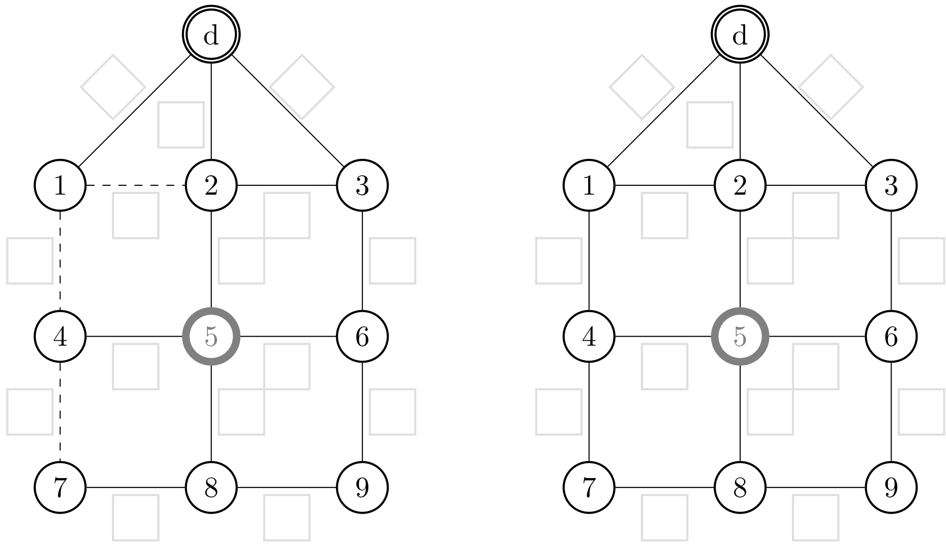

With loops, question (i) Without loops, question (ii)

Figure 6

(i) Find a possible weight for each link such that the traffic to d would be caught in a transient forwarding loop on the links (1-2), (1-4) and (4-7) upon the failure of node 5. Your weights must be such that forwarding loops are only observed on these links. Indicate your weights directly in Figure 6 (left), in the squares next to each link. (7 Points)

(ii) Now, find a possible weight for each link such that the traffic to d would never be caught in a forwarding loop upon the failure of node 5, on any link. Indicate your weights directly in Figure 6 (right). (2 Points)

c) Reverse engineering **(5 Points)**

leave blank

Consider the network topology of Figure 7 in which node A only uses the following forwarding paths:

- To reach B: A–B
- To reach C: A–C
- To reach D: A–C–D *and* A–B–D *and* A–B–C–D

(i) Could a shortest-path-based protocol such as OSPF have computed such paths? Briefly explain your answer. (2 Points)

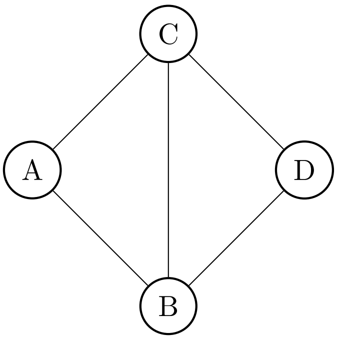

Figure 7

(ii) How would you configure the network such that it computes these paths? Briefly explain. (3 Points)

d) A very long downtime **(7 Points)**

leave blank

Consider the network in Figure 8 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. Figure 9 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.

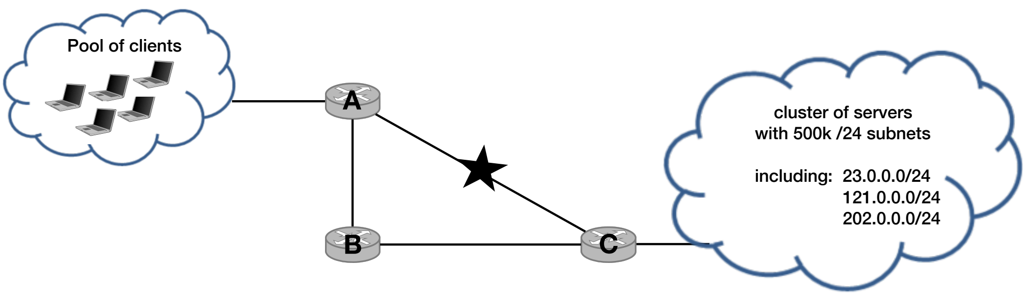

Figure 8

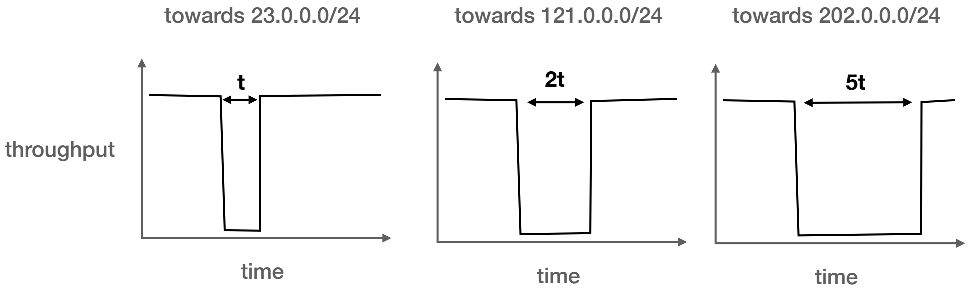

Figure 9

(i) Explain why the downtimes experienced by the three prefixes differ. (3 Points)

(ii) Describe two solutions to speed up the convergence for these three prefixes. (4 Points)

Solution 1:

Solution 2:

leave blank

Task 3: Inter-domain routing **38 Points**

leave blank

a) Warm-up (6 Points)

For the following true/false questions, check either *true*, *false* or nothing. For each question answered correctly, one point is added. For each question answered incorrectly, one point is removed. There is always one correct answer. This subtask gives at least 0 points.

true	false	
☐	☐	Using "AS path prepending", an AS can guarantee that a certain route it propagates will be only used as a last resort.
☐	☐	Configuring a working iBGP full mesh does not require every router to be physically connected to all the others.
☐	☐	If traffic follows the AS path $1 \rightarrow 2 \rightarrow 3$ from AS 1 to AS 3, then the traffic from AS 3 to AS 1 will *always* take the path $3 \rightarrow 2 \rightarrow 1$.
☐	☐	BGP route maps only modify BGP advertisements.
☐	☐	A BGP "hijack" is the attempt of an AS to take over (parts) of the traffic towards a prefix that is not owned by that AS.
☐	☐	A router receiving BGP advertisements verifies their origin and authenticity.

b) BGP policies **(7 Points)**

leave blank

Consider the topology consisting of four Autonomous Systems (ASes) in Figure 10. ASes A and B consist of three (A1, A2, A3) and two (B1, B2) routers, respectively. ASes C and D each consist of a single router. Solid lines represent physical connections, while dashed lines represent the paths for incoming traffic preferred by A. Each AS advertises its own unique prefix. The business relationships of A with B, C, and D are unknown, but the route maps configured by AS A are shown in Table 1. D is a customer of B. All ASes apply standard BGP policies.

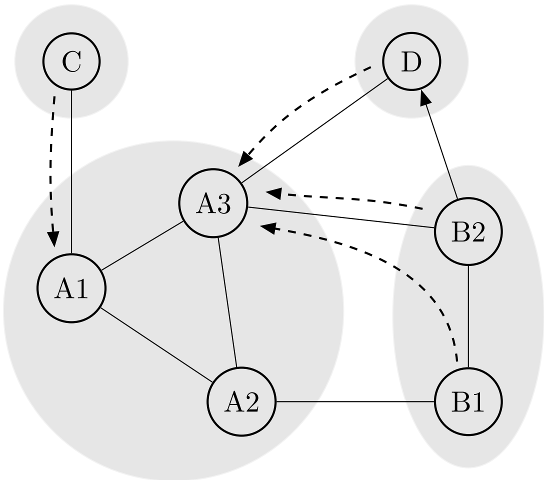

Figure 10: An Internet topology with four ASes A (three routers), B (two routers), C, and D.

Session	Map	Actions
A1–C	IN	set community 3 set local-pref 100
	OUT	match community 3 match community 2 match community 1
A2–B1	IN	set community 1 set local-pref 10
	OUT	match community 3
A3–B2	IN	set community 1 set local-pref 10
	OUT	match community 3
A3–D	IN	set community 2 set local-pref 50
	OUT	match community 3

Table 1: Route maps of AS A.

(i) The route maps are ambiguous, and multiple business relationships are possible. Consider the following three scenarios: Explicitly write if they are possible or not assuming A correctly follows standard BGP policies. Briefly explain your answer. (3 Points)

C is a customer of A: __________

D is a customer of A: __________

Both B and D are peers of A: __________

The router A3 is much more powerful than the router A2. Because of this, the operator of AS A would like to receive all *incoming* traffic from B at A3 (A2 must remain available as a fallback). Traffic from C and D should always use the shortest paths (D → A3, C → A1). The preferred paths are illustrated with dashed lines in Figure 10.

The operator of AS A considers three different ways to achieve this: MED, path prepending, and advertising a more-specific prefix. For the following tasks, assume that A is a customer of both B and D.

(ii) Give an example how the operator of A could use the MED to achieve her goal. Briefly explain what the MED is, and how it can be used. (1 Point)

Explanation: ____________________

____

Example usage: ____________________

____

(iii) As an alternative to the MED, A considers path prepending. Explain why path prepending works for this usecase, and give an example how A could use it. (1 Point)

Path prepending works here, because ____________________

____

Example usage: ____________________

____

(iv) Finally, the operator from AS A considers advertising a more-specific prefix to AS B via A3. Assume A owns the prefix 42.0.0.0/23. From A2 to B1, it advertises 42.0.0.0/23. From A3 to B2, it advertises both 42.0.0.0/24 and 42.0.1.0/24. However, advertising a more-specific prefix has an unintended (and undesired!) side effect in this case. Which one, and why? (2 Points)

Why does this work? ____________________

____

What is the side effect, and why? ____________________

____

____

c) Peering, business, and performance **(11 Points)**

leave blank

Consider the topology with 7 Autonomous Systems (ASes) in Figure 11. After many business deals, complex relationships have emerged: Single-headed plain arrows point from providers to their customers (AS A is a provider of AS C), while double-headed dashed arrows connect peers (AS A and AS B are peers). On each link, packets experience different delays, which are shown as well. Each AS advertises it's own unique prefix, and all ASes apply standard BGP policies according to their business relationships.

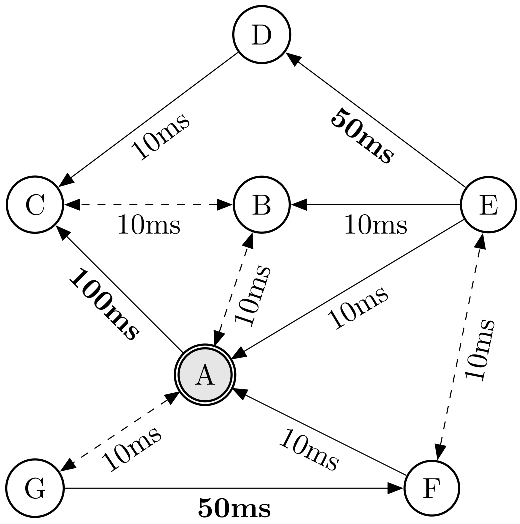

Figure 11: An Internet topology with 7 ASes.

(i) First, analyze all *incoming* advertisements at AS A. For each *other* AS, write down the path *selected* by AS A, compute the cumulative delay along the links and additionally note down whether it is the path with the shortest delay through the network to reach this destination. (4 Points)

Destination	Path	Delay (ms)	Shortest delay?
B	A → B	10	Yes
C	A →		
D	A →		
E	A →		
F	A →		
G	A →		

(ii) Next, analyze all *outgoing* advertisements from AS A. For each *other* AS, write down the path *selected* to reach AS A, compute the cumulative delay along the links and additionally note down whether it is the path with the shortest delay through the network to reach AS A. (4 Points)

Origin	Path	Delay (ms)	Shortest delay?
B	B →		
C	C →		
D	D →		
E	E →		
F	F →		
G	G →		

Consider now that AS A receives a large financial investment and its operators are now willing to minimize delay—for both incoming and outgoing traffic—by re-configuring their local BGP policies, even if that means paying more than before.

(iii) Not all traffic can be redirected over the path with the lowest delay in the network just by re-configuring BGP policies locally. Which paths, and why? (2 Points)

Paths that cannot be improved: ____________________

Why can they not be improved? ____________________

(iv) Aside from changing BGP policies, the operators of AS A also negotiate a peering relationship with AS F. Does this reduce delays while using standard policies? If so, give an example where the delay is reduced. (1 Point)

d) BGP and IGP **(14 Points)**

leave blank

As we have seen in the routing project, advertisements received via iBGP sessions are *not* redistributed via other iBGP sessions. BGP advertisements are indeed only redistributed from eBGP to eBGP, from iBGP to eBGP, and from eBGP to iBGP. Because of this, iBGP sessions are thus usually set up in a full mesh. In this task we investigate why.

Consider the AS with five routers (A to E) shown in Figure 12. The AS uses OSPF as Interior Gateway Protocol (IGP) with link weights between 10 and 100. The AS does *not* have an iBGP full mesh, instead (bidirectional) iBGP sessions are set up only between D and each other router. Finally, the AS is connected to another AS managed by EPFL and an eBGP session is set up with router D. D uses next-hop-self, which means that over the iBGP sessions, D announces itself as BGP next-hop to reach EPFL.

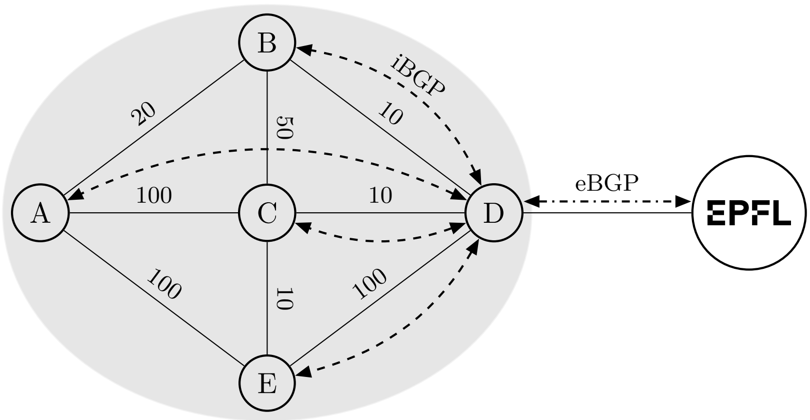

Figure 12: An AS which does *not* run an iBGP full mesh

(i) EPFL advertises its unique prefix to D which distributes it via iBGP. For each router in the AS, indicate below: (i) what is the next hop it learns to reach the EPFL prefix via iBGP; (ii) what is the path taken according to the IGP; and (iii) whether EPFL is reachable or not. If EPFL is not reachable, briefly explain why below the table.

(4 Points)

Router	BGP next-hop	IGP path	EPFL reachable?
A			
B			
C			
D	EPFL	∅	Yes
E			

If EPFL is not reachable from any router(s), explain the reason(s): ____________

(ii) A while later, the AS establishes a new eBGP session with an AS managed by ETH at router A (see Figure 13). ETH advertises its unique prefix to router A, which distributes it via iBGP to router D. ETH's prefix is distinct from EPFL's one. A also uses next-hop-self.

For each router, indicate in the following table: (i) what is the next hop learned for the ETH prefix via iBGP; (ii) what is the path taken according to the IGP; and (iii) whether ETH is reachable or not. If not, briefly explain why below the table. (5 Points)

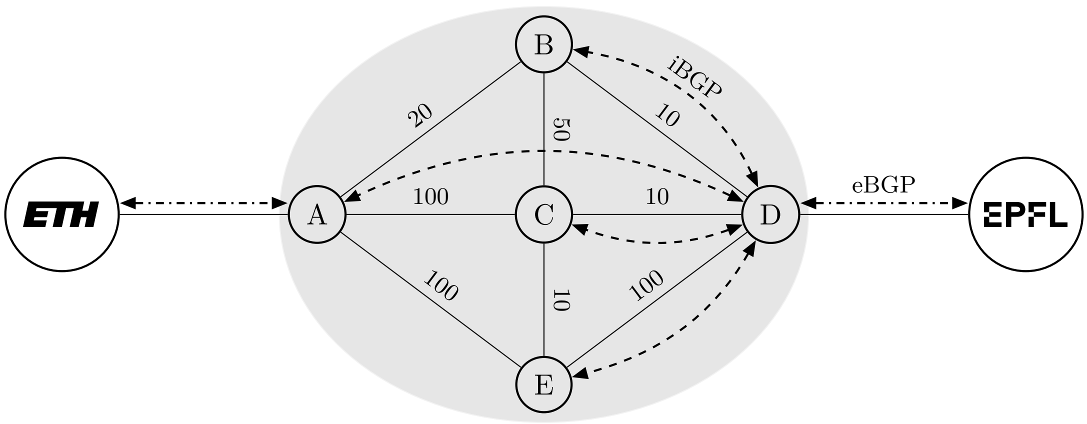

Figure 13: An eBGP session is established with ETH at router A.

Router	BGP next-hop	IGP path	ETH reachable?
A			
B			
C			
D			
E			

If ETH is not reachable from any router(s), explain the reason(s):

(iii) To avoid connectivity problems, the AS operator sets up three additional iBGP sessions: A↔B, A↔C and A↔E (see Figure 14). Additionally, the operator establishes one new eBGP session with ETH via E. Concretely, ETH now advertises its unique prefix to both routers A and E, which distribute it via iBGP. E also uses next-hop-self.

Indicate again for each router in the AS: (i) what is the next hop learned via iBGP for the ETH prefix; (ii) what is the path taken according to the IGP; and (iii) whether the ETH prefix is reachable. If the ETH prefix is not reachable, briefly explain why below the table. (5 Points)

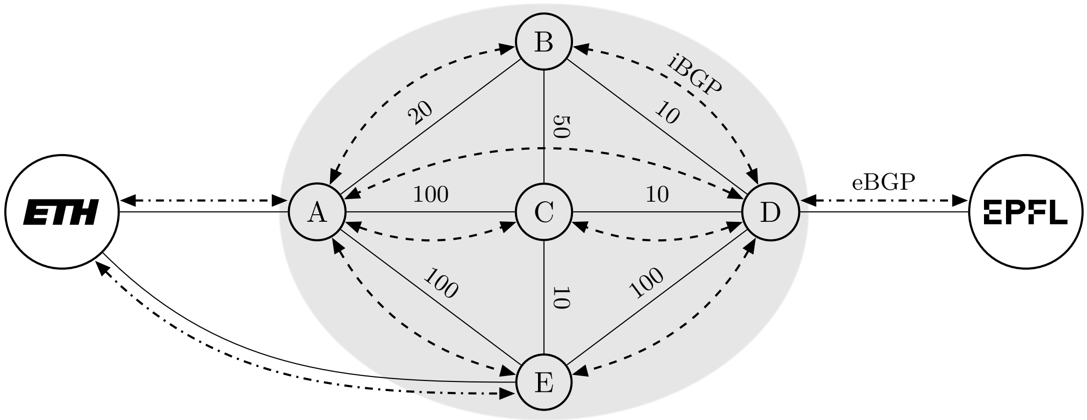

Figure 14: Additional iBGP and eBGP sessions are established.

Router	BGP next-hop	IGP path	ETH reachable?
A			
B			
C			
D			
E			

If ETH is not reachable from any router(s), explain the reason(s): ____________________

____

____

leave blank

leave blank

Task 4: Reliable transport **27 Points**

a) Warm-up **(8 Points)**

For the following true/false questions, check either *true*, *false* or nothing. For each question answered correctly, one point is added. For each question answered incorrectly, one point is removed. There is always one correct answer. This subtask gives at least 0 points.

true ☐ false ☐ 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.

true ☐ false ☐ 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.

true ☐ false ☐ 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.

true ☐ false ☐ 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 Figure 15.

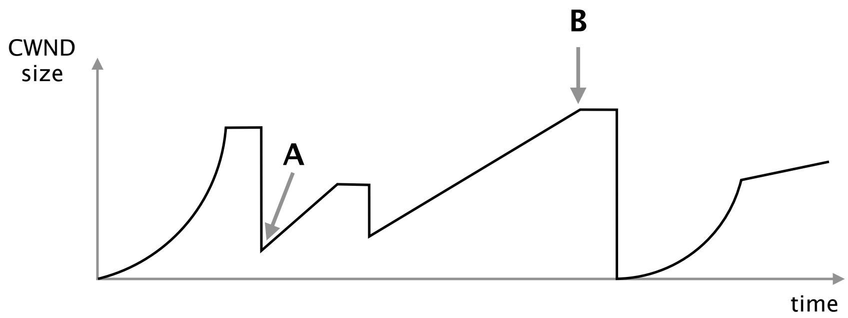

Figure 15

true ☐ false ☐ Flow `f` was in the slow start phase exactly twice.

true ☐ false ☐ Flow `f` experiences at least one packet loss between time A and B.

true ☐ false ☐ In the future, `f` will never be able to experience a higher CWND size than B.

true ☐ false ☐ Consider another flow `f2` starting at exactly the same time as `f` and traversing the exact same path, then `f2` would experience the exact same CWND evolution.

b) GBN protocol (10 Points)

leave blank

You are in charge of implementing a GBN protocol similar to the one we used in the reliable transport project. Each of the following subtasks contains a code snippet (in pseudo-code) that implements a specific functionality of the protocol. Unfortunately, each snippet contains one error.

For each subtask, we ask you to: (i) identify the problem and explain (in words) how it could be fixed; and (ii) describe how it would affect the functionality of the GBN protocol. You can assume that the sequence number can get infinitely large and never overflows.

(i) Algorithm 1 implements a function which retransmits segments upon receiving duplicate acknowledgements. For instance, the segment 3 should be retransmitted upon receiving the sequence [2, 3, 3, 3]. (4 Points)

Algorithm 1 Duplicate detection

```
n_duplicate ← 0
previous ← −1

while receiving ACKs do
    current ← received ACK number

    if previous == current then
        n_duplicate += 1
        if n_duplicate == 2 then
            n_duplicate = 0
            retransmit(current)
    else
        previous = current
```

Explain the error and describe a fix:

Explain the consequence(s) of the error on the GBN behavior:

(ii) Algorithm 2 implements the sender Selective ACKnowledgement (SACK) part where the sender reacts to received SACK blocks. (6 Points)

Algorithm 2 Sender SACK implementation

```
next ← last received ACK number
sender_buffer ← {}            // saves sent but not yet acknowledged segments

while receiving SACK packets do
    ACK ← received ACK number
    for next ≤ i < ACK do
        remove i from sender_buffer
    next ← ACK

    for ACK ≤ i < beginning of first SACK block do
        retransmit(i)
    for all SACK blocks do
        remove block segments from sender_buffer
        retransmit segments between blocks

    while size of sender_buffer <  current window do
        send new data segments
```

Explain the error and describe a fix: ________________________________

Explain the consequence(s) of the error on the GBN behavior: ________________________________

c) Fairness **(9 Points)**

leave blank

Figure 16 shows a network with 6 nodes and unknown link bandwidths. You would like to figure out the bandwidth of as many links as possible. Another operator, who knew all the link bandwidths, used the min-max fair allocation algorithm to compute a fair bandwidth allocation for 9 flows shown in Table 2. None of these 9 flows started *and* ended at the same nodes, but it is possible that two or more flows started (resp. ended) at the same node and ended (resp. started) at different ones.

We ask you to find (possible) link bandwidths and write them down in Figure 16. Should you not have enough information to uniquely identify the bandwidth of a link, indicate the minimal bandwidth required to achieve the min-max fair allocation in Table 2. **In addition**, we ask you to find a corresponding path for each of the 9 flows and report it in the last column of the table.

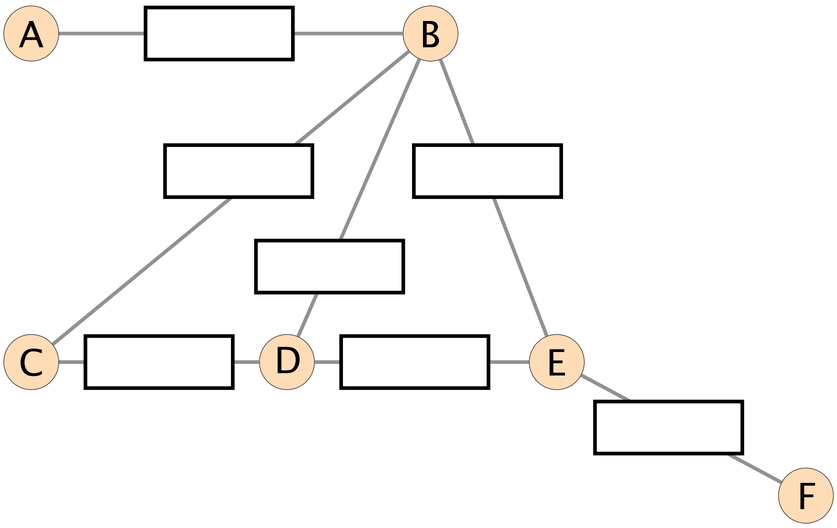

Figure 16

	Step 1	Step 2	Step 3	Step 4	Step 5	**Possible flow path** (e.g., A-B-C)
Flow 1	1	3	7	11	**15**	
Flow 2	**1**					
Flow 3	1	**3**				
Flow 4	1	3	**7**			
Flow 5	1	3	**7**			
Flow 6	1	3	7	**11**		
Flow 7	**1**					
Flow 8	1	3	**7**			
Flow 9	1	3	7	11	**15**	

Table 2: Min-max fair allocation algorithm output. The bold number at the end of each row indicates the bandwidth that eventually was allocated to each flow. **Fill the last column in!**

leave blank

Task 5: Applications **32 Points**

leave blank

a) Warm-up **(7 Points)**

For the following true/false questions, check either *true*, *false* or nothing. For each question answered correctly, one point is added. For each question answered incorrectly, one point is removed. There is always one correct answer. This subtask gives at least 0 points.

true	false	
☐	☐	HTTP is a text-based protocol to transport the text of websites. Images need to be transmitted via other protocols.
☐	☐	Performing a traceroute to one of the DNS root-servers (e.g., `traceroute a.root-servers.net`) would reveal the path to this root-server and allow you to see the locations of local DNS servers.
☐	☐	`k.root-servers.net` points to an IP address which is used simultaneously by multiple servers in different countries.
☐	☐	Content Delivery Networks (CDNs) can deliver contents faster because they are physically closer to the client.
☐	☐	Sending an e-mail from a@ethz.ch to b@epfl.ch does not trigger DNS queries if the sender visited www.epfl.ch before and no caches expired since then.
☐	☐	MIME relies on Base64 to encrypt the content of an email.
☐	☐	An e-mail is bigger (in terms of bits) for the receiver than for the sender.

b) Network debugging **(8 Points)**

leave blank

In this question, you will try to infer parts of the Internet topology using `traceroute`, a network debugging tool. You first run `traceroute` at home and observe the output in Figure 17 (`traceroute` from host A to `x.ch`). You then run `traceroute` from your university and observe the output in Figure 18 (`traceroute` from host B to `x.ch`).

```
user@A:~$ traceroute x.ch
traceroute to x.ch (9.9.9.9), 30 hops max, 60 byte packets
 1  1.1.1.1              1 ms   1 ms  1 ms
 2  7.7.7.7              4 ms   4 ms  4 ms
 3  2.2.2.2             10 ms  10 ms
    5.5.5.5             20 ms
 4  4.4.4.4             16 ms  16 ms 16 ms
 5  9.9.9.9             18 ms  18 ms 18 ms
```

Figure 17: `traceroute` output at host A

(i) Based on the output in Figure 17, how many links does a packet from host A traverse until it reaches `x.ch`? Explain your answer. (1 Point)

```
user@B:~$ traceroute x.ch
traceroute to x.ch (9.9.9.9), 30 hops max, 60 byte packets
1  8.8.8.8              2 ms    2 ms  2 ms
2  3.3.3.3              6 ms    6 ms  6 ms
3  4.4.4.4              8 ms    8 ms  8 ms
4  9.9.9.9              10 ms  10 ms 10 ms
```

Figure 18: `traceroute` output at host B

(ii) Use Figure 19 to draw a topology which is compliant with the combined outputs of Figure 17 *and* Figure 18. Whenever possible, label each node with its IP address and each link with its propagation time.

Note: You can assume the following:

- The network is correctly configured and there is no congestion or link failure.
- The propagation time of packets over links is the *only* cause of network delays.

(6 Points)

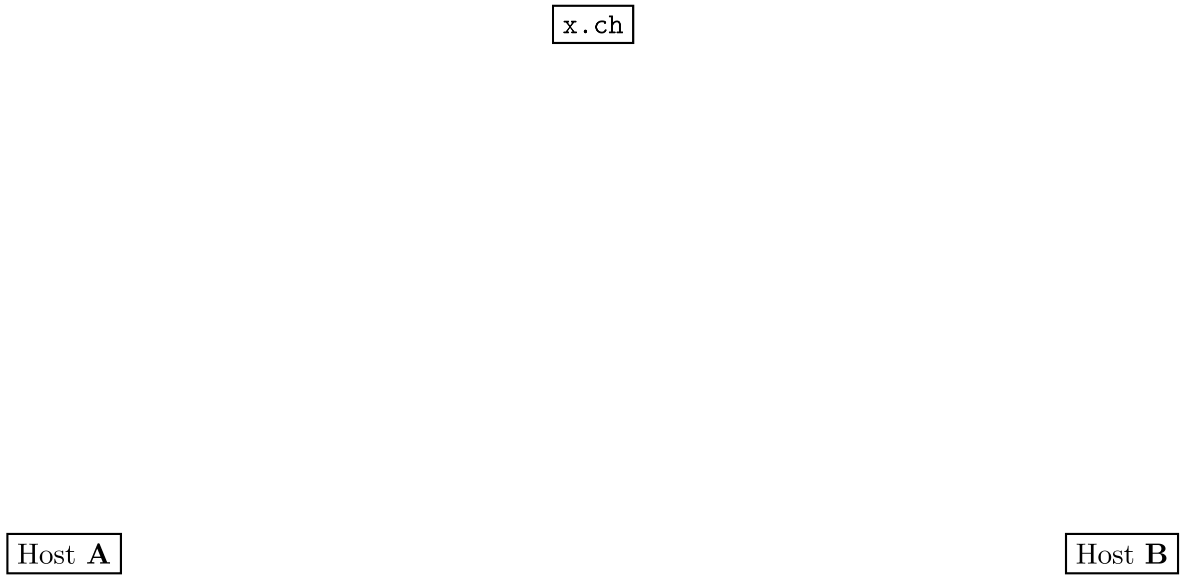

Figure 19: Use this figure to draw the topology

(iii) Does `traceroute` allow you to estimate the number of (layer-2) switches in a network? Explain your answer. (1 Point)

____

____

c) DNS basics **(6 Points)**

In this task, we look at how your web browser figures out how to connect to `http://www.mail.ethz.ch/login`.

leave blank

Note: You can assume the following:

- Each level of the domain hierarchy uses a separate DNS server.
- All caches are initially empty.

(i) List all the DNS servers involved when resolving `http://www.mail.ethz.ch/login`. Sort them according to the order in which they receive queries if you use an *iterative* resolver (start with the one which receives the first query). Write your entries in the form "*DNS server for domain x*". *Hint:* The answer requires 7 lines or less. (2 Points)

A Local DNS resolver installed on the client

B *DNS server for* ______________________

C *DNS server for* ______________________

D *DNS server for* ______________________

E *DNS server for* ______________________

F *DNS server for* ______________________

G *DNS server for* ______________________

(ii) List all DNS queries (including their source and destination) for a *recursive* resolver to determine the IPv4 address hosting `http://www.mail.ethz.ch/login`. You can use the letters A–G from the list above to specify the DNS servers and X for the client. *Hint:* The correct answer requires 10 lines or less. (3 Points)

1. source: X, destination: A, query: ______
2. ______
3. ______
4. ______
5. ______
6. ______
7. ______
8. ______
9. ______
10. ______

(iii) Immediately after you determined the IPv4 address for `http://www.mail.ethz.ch/login`, you want to determine the IPv4 address for `http://www.ethz.ch`. List all DNS queries (including their source and destination) for a *recursive* resolver to determine the IPv4 address hosting `http://www.ethz.ch` *assuming that all the responses from above are still in the caches*. You can use the letters A–G from the list above to specify the DNS servers and X to for the client. *Hint:* The correct answer requires 10 lines or less. (1 Point)

1. ______
2. ______
3. ______
4. ______
5. ______
6. ______
7. ______
8. ______
9. ______
10. ______

d) Setting up a video streaming service **(11 Points)**

leave blank

In this task, you will set up a brand new video streaming service.

(i) You first create a website (`http://www.yourvideo.ch`) where visitors can buy subscriptions. To improve the performance of your website, you want to make it small enough that it fits in one single IPv4 packet. Assuming the maximum size of an IP packet is 1500 bytes, how many bytes are remaining for the application layer at most? Explain your answer. (1 Point)

(ii) Can you use all the available bytes that you computed above for the source code of the website? Explain your answer. (1 Point)

(iii) After a user bought a subscription, you want to show her the list of all available videos. As you know, the HTTP protocol is stateless. Which technique can you use to maintain state about whether a user is logged in or not? Mention the name of the technique, where the state is stored and how your web server knows whether a user is logged in upon receiving a HTTP request from her. (2 Points)

Name of the technique: __________

Where the state is stored: __________

How the server knows the state: __________

(iv) Unfortunately, you now realize that the website and the 10 images that it contains do not fit in a single IP packet and you therefore need to think of other ways to improve the page load speed. You come up with the following solution: Instead of storing everything on the same server, you add 10 additional servers to your data center and store each image on one of them (i.e., image 1 is on 123.0.0.101/image1.jpg, image 2 is on 123.0.0.102/image2.jpg and so on). Each server has a high bandwidth Internet connection.

Compared to the initial solution (everything on the same server), does this new approach have an impact on the number of established connections if the client uses HTTP 1.0 (also referred to as naive HTTP in the lecture) or if the client uses HTTP 1.1? Explain your answers. (2 Points)

Impact on the number of established connections if HTTP 1.0 is used: __________

Impact on the number of established connections if HTTP 1.1 is used: __________

Now that you've finished your website, you can finally start adding videos that your customers can stream.

(v) You heard that your competition provides the videos in multiple resolutions and bitrates. Is this a good idea for your case too, even if you do not use a CDN? Explain your answer. (1 Point)

(vi) Assuming you decided to also provide the videos in different resolutions and bit-rates. Explain how the player knows which versions are available as well as how and when it selects the best version? (2 Points)

How the player knows the available versions: __________

How the player selects a version: __________

When the player selects a version: __________

(vii) One day customers tell you they often observe "rebuffering". Explain what this is and describe one situation in which it can occur. (2 Points)

Explanation: __________

Example of a situation: __________