Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institut für Technische Informatik und Kommunikationsnetze

Department ITET
August 2018

Prof. Dr. Laurent Vanbever
Tobias Bühler, Rüdiger Birkner, Thomas Holterbach, Roland Meier

Exam: Communication Networks

22 August 2018, 14:00–16:30, Room HIL G 61

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 & Switching	/30	
Intra-domain routing	/25	
Inter-domain routing	/41	
Reliable transport	/29	
Security & Applications	/25	
Total	/150	

leave blank

Task 1: Ethernet & IP forwarding **30 Points**

leave blank

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	
☐	☐	In contrast to packet switching, circuit switching does not require switches to know any information about the network topology to function correctly.
☐	☐	In packet-switched networks, IP packets belonging to the same TCP flow will not necessarily be forwarded along the same path.
☐	☐	Two hosts belonging to different VLANs cannot exchange IP traffic.
☐	☐	End-hosts connected to a switch access port can discover the VLAN they belong to by observing the received Ethernet frames.
☐	☐	Adding an extra link to an existing spanning tree (e.g. by activating a blocked port) would necessarily create a cycle/loop.

b) A small detour (6 Points)

leave blank

Consider the layer-2 network composed of 8 switches in Figure 1. The network interconnects two hosts and one router.

Each of the 14 links is identified with a letter (from A to N). The network uses two VLANs: VLAN 10, connecting host 1 and the router, and VLAN 11, connecting host 2 and the router. Switches maintain per-VLAN spanning trees with unary link cost and tie-break based on the smallest switch ID. Switch 1 is configured as root switch for VLAN 10 and switch 6 is configured as root switch for VLAN 11.

Host 1 and host 2 are located in different IP subnets (10/24 and 11/24) and use the router as default gateway.

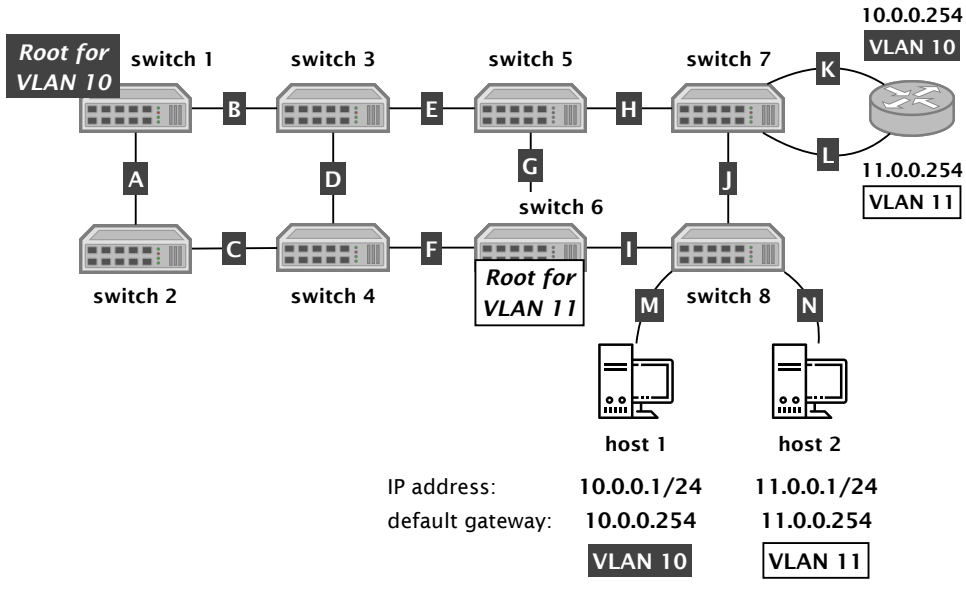

Abbildung 1: A layer-2 network topology relying on VLANs

(i) Indicate the sequence of links followed by IP packets going from host 1 to host 2. Your answer should be the corresponding sequence of letters. There is no need to indicate the path from host 2 to host 1. (4 Points)

(ii) Where would you place the root switches to minimize the number of links being used when host 1 sends traffic to host 2? Indicate the sequence of links followed by IP packets from host 1 to host 2 under your new configuration. (2 Points)

c) When redundancy is *not* beneficial **(9 Points)**

leave blank

It is your first day as a network engineer and you're already called for helping to troubleshoot a problem of duplicate IP addresses. The problem is illustrated in Figure 2 where the administrator of host 2 has statically assigned an IP address (10.0.0.1) which has already been allocated by the DHCP server to host 1. The two hosts and the DHCP server are connected together via a layer-2 network (no VLAN). An IP router, which acts as Internet gateway for the two hosts, is also connected to the network.

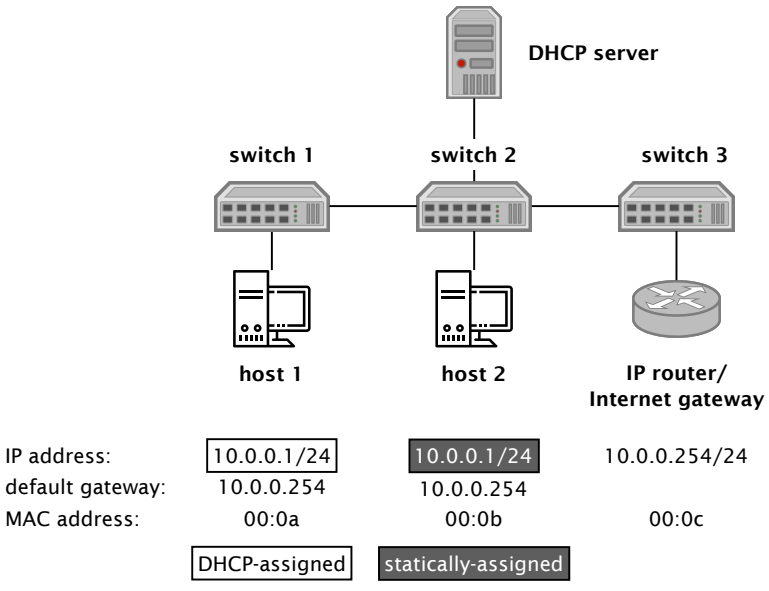

Abbildung 2: A network topology with two hosts sharing the same IP address

(i) Briefly describe the connectivity issues that host 1 and host 2 are experiencing due to the duplicated IP allocation. (2 Points)

(ii) Does one host experience more difficulty to access the Internet than the other? Briefly explain why or why not. (2 Points)

(iii) You are asked to troubleshoot the network. Describe one procedure to obtain the list of hosts that share the same IP address *and* one procedure to locate the physical port on which these hosts are connected. It is safe to assume that MAC addresses are unique.

As a network operator, you have full access to the network devices and can observe the network traffic on any link but you do *not* have access to the end-host. Design your procedures accordingly. (5 Points)

Detection procedure:

Localization procedure:

leave blank

d) Summer pruning **(10 Points)**

(i) Consider an IP router with a forwarding table composed of the 9 entries depicted on the left. 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. Your reduced forwarding table should be such that the forwarding decision made by the router for any IP packet is equivalent to the initial one. (8 Points)

prefix	next-hop
82.130.32.0/20	1
82.130.64.0/20	1
82.130.80.0/20	2
82.130.96.0/20	1
82.130.112.0/21	1
82.130.120.0/21	1
82.130.122.0/24	1
82.130.123.0/24	1
82.130.124.0/24	2

prefix	next-hop

(ii) Answer to the same question as above considering the forwarding table below instead (left) in which an extra default route is defined. (2 Points)

prefix	next-hop
0.0.0.0/0	1
82.130.32.0/20	1
82.130.64.0/20	1
82.130.80.0/20	2
82.130.96.0/20	1
82.130.112.0/21	1
82.130.120.0/21	1
82.130.122.0/24	1
82.130.123.0/24	1
82.130.124.0/24	2

prefix	next-hop

leave blank

Task 2: Intra-domain routing **25 Points**

leave blank

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 ☐ Static routes can be used in addition to a dynamic routing protocol.

true ☐ false ☐ Consider a network where half of the routers use link-state protocols, while the remaining half uses distance-vector. If each link has the same weight in both protocols, the forwarding state obtained by each router will be equivalent to the one obtained if all routers were to use the same protocol.

true ☐ false ☐ When a router uses a distance-vector protocol, it knows the path(s) the traffic uses to reach the destination.

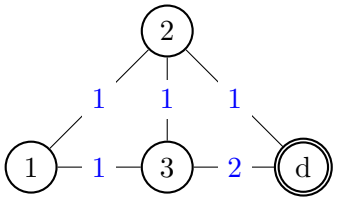

Abbildung 3: Topology with 4 nodes. Link costs are symmetric.

true ☐ false ☐ Consider the scenario in Figure 3, and assume a distance-vector protocol is used with poisoned reverse. To reach d, the node 3 knows three paths: one with next-hop 2 and cost 2, one with next-hop d and cost 2, and one with next-hop 1 and cost ∞.

true ☐ false ☐ Consider the scenario in Figure 3, and assume a distance-vector protocol is used *without* poisoned reverse. Upon failure of the link between node 3 and d, the nodes 1, 2 and 3 will experience the count-to-infinity problem.

leave blank

b) Where does the traffic go? **(8 Points)**

We consider the network in Figure 4 where each router runs OSPF and performs per-packet load balancing. Concretely, whenever a router knows several shortest paths to a destination, it chooses one randomly for each packet going to that destination, irrespectively of the flow the packet belongs to. In this task, we focus on the traffic from s to d.

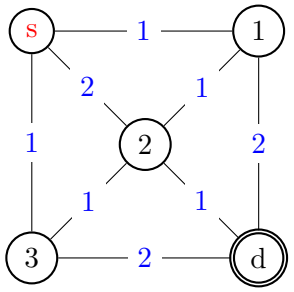

Abbildung 4: Topology with 5 nodes. Link costs are symmetric.

(i) Given that s is sending 12 Mbps to d, what is the expected load (in Mbps) on the link between 2 and d?

(2 Points)

(ii) The link between 2 and d suddenly breaks. During the convergence, can a forwarding loop including router 1 and router 3 be formed? Explain your reasoning. Feel free to draw directly on Figure 4.

(4 Points)

(iii) Instead of a sudden failure, consider that the link between node 2 and d has to be shut down for maintenance. Describe a procedure that you could apply before the maintenance in order to prevent transient loops.

(2 Points)

c) Playing with Dijkstra (12 Points)

leave blank

In the following, we ask you to propose or reason about different modifications made to the Dijkstra algorithm (i.e. the algorithm used by OSPF).

(i) Describe a (simple) modification to the Dijkstra algorithm so as to make it return the path(s) containing the smallest number of links. (1 Point)

(ii) One of your friend proposes you to modify the Dijkstra algorithm to find the longest paths instead of the shortest ones by considering the inverse of each link cost (e.g. $\frac{1}{2}$ instead of 2). Explain with an example why this would not work. (2 Points)

(iii) After realizing his mistake, the same friend now proposes you to take the negative of each link cost (e.g. -2 instead of 2) to find the longest paths instead of the shortest ones. Explain with an example why this would not work. (4 Points)

(iv) Yet another friend of yours modified the Dijkstra algorithm so that every router selects and sends traffic along its second best path (i.e. the path with the second lowest sum of link costs) instead of the shortest path. Explain with an example why this modification can lead to forwarding loops and therefore does not work (again!). (5 Points)

leave blank

Task 3: Inter-domain routing **41 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 ☐ For an iBGP full-mesh to work, every router must have a direct physical connection with every other router.

true ☐ false ☐ Configuring `next-hop-self` on eBGP sessions is required for them to work properly.

true ☐ false ☐ Tier-2s only have Tier-3s as customers.

Consider the simple BGP network in Figure 5. Single-headed plain arrows point from providers to their customers (AS A is the provider of AS D), while double-headed dashed arrows connect peers (AS B and AS C are peers). Each AS in the network originates a unique prefix that it advertises to all its BGP neighbors. Each AS also applies the default selection and exportation BGP policies based on their customers, peers and providers.

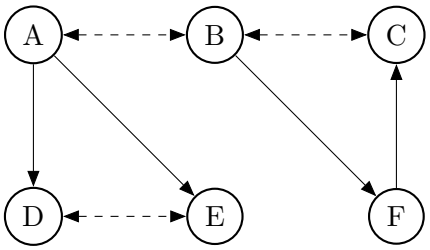

Abbildung 5: A simple BGP network.

true ☐ false ☐ AS D has two routes available to B: $A\ B$ and $E\ A\ B$.

true ☐ false ☐ AS B prefers the route $F\ C$ to AS C over the shorter direct one C.

true ☐ false ☐ Every AS has at least one route to every other AS.

leave blank

b) Path Selection **(6 Points)**

Consider the AS in Figure 6. It has three border routers (A, B and C) and three internal routers (D, E, F). The routers are connected through an iBGP full-mesh. OSPF is used internally with the given link weights. Each border router receives two eBGP advertisements from its eBGP neighbors with the depicted attributes (AS path, MED and local-pref).

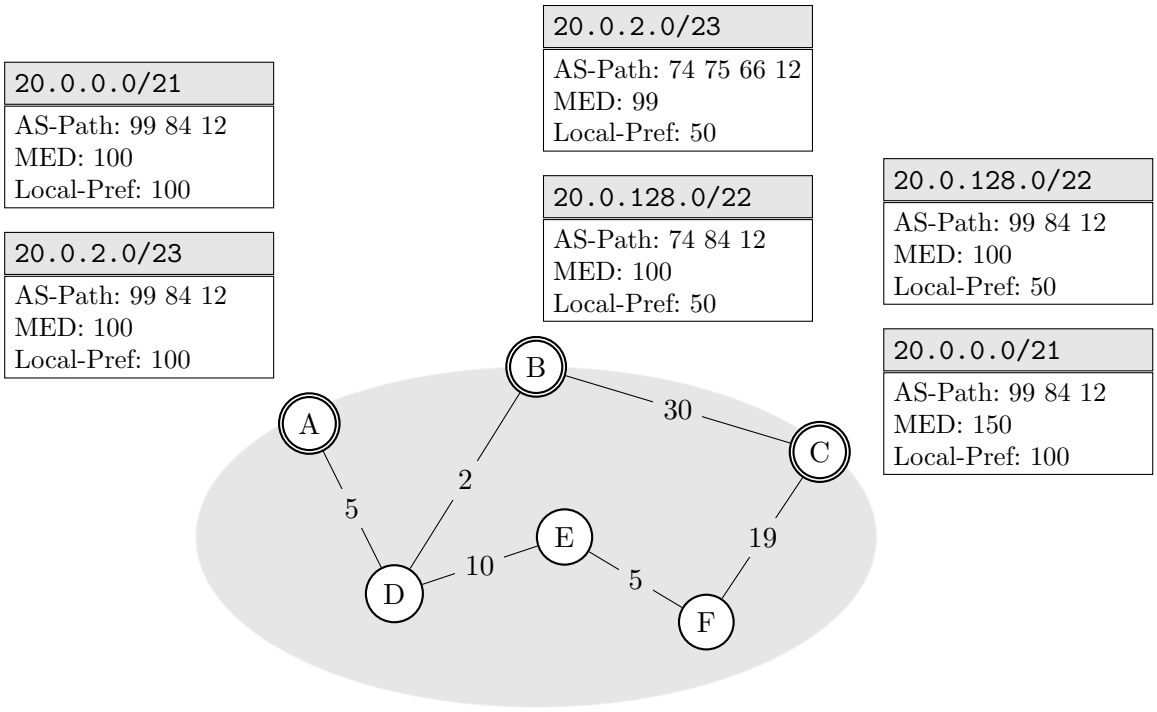

Abbildung 6: A simple BGP network forming an iBGP full-mesh.

Complete the routing table of each router by indicating the router ID of the selected egress (A, B, C) using the provided template:

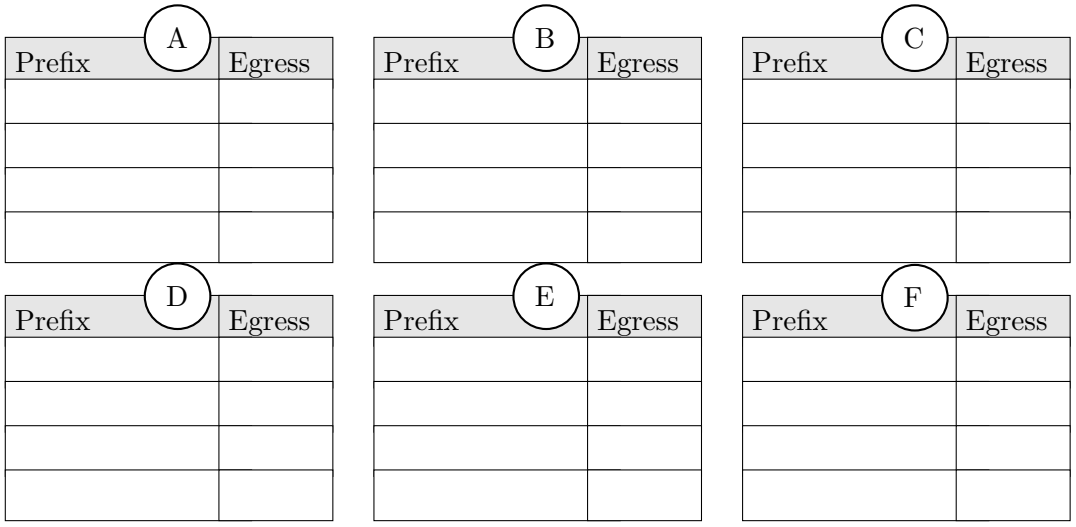

leave blank

c) BGP Hijack **(8 Points)**

Consider the Internet topology consisting of 11 Autonomous Systems (ASes) in Figure 7. Single-headed plain arrows point from providers to their customers (AS A is the provider of AS E), while double-headed dashed arrows connect peers (AS A and AS B are peers). Each AS is made up of a single BGP router and applies the default selection and exportation BGP policies based on their customers, peers and providers.

In this task, the routers break ties using the AS number of the neighbor: in case multiple routes are equally good, the router selects the route of the neighbor with the lowest AS number (alphabetical order).

AS G is the origin of prefix 20.0.0.0/22 and advertises it to its neighbors. Independently of what the external advertisements are, AS G ***always*** prefers its internal route to reach any IP destination in 20.0.0.0/22.

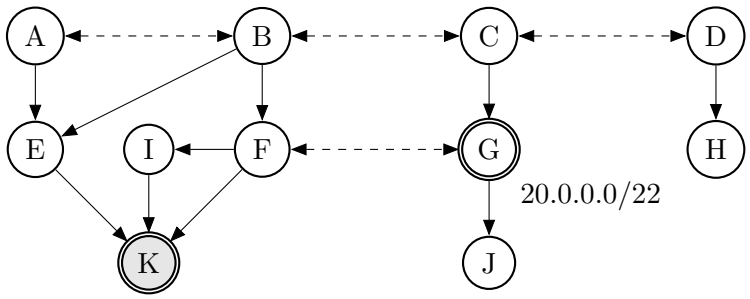

Abbildung 7: An Internet topology with 11 ASes. AS K aims at hijacking traffic destined to AS G.

(i) AS K wants to hijack all the traffic going to AS G for 20.0.0.0/22. It starts advertising the exact same prefix. From which ASes is it able to attract the traffic? (1 Point)

(ii) AS K is not satisfied by the result. What can it do to attract traffic destined to AS G from more of the ASes? List the ASes from which it is able to attract the traffic and explain why this works. (2 Points)

(iii) The ASes from which AS K manages to attract the traffic realize what is happening as all their traffic to 20.0.0.0/22 goes to a dead-end (AS K).

Show how AS K could still deliver the traffic to the real destination (AS G) by poisoning the AS path while attracting as much traffic as possible. In addition, list the ASes from which it can attract the traffic. (2 Points)

(iv) Can you think of a different way for AS K to achieve similar results as in *(iii)* without poisoning the AS path? Explain. (3 Points)

d) Path Hunting/Exploration **(7 Points)**

leave blank

Consider the Internet topology in Figure 8 from the perspective of AS 10. Single-headed plain arrows point from providers to their customers (AS 21 is the provider of AS 60), while double-headed dashed arrows connect peers (AS 10 and AS 32 are peers). Each AS is made up of a single BGP router and applies the default selection and exportation BGP policies based on their customers, peers and providers.

In this task, the propagation of a BGP message (withdrawal) between two directly connected ASes takes $5s$ and a router detects a failure after $5s$. AS 60 announces a route to the prefix 20.0.0.0/24 to its directly connected neighbors that propagate it.

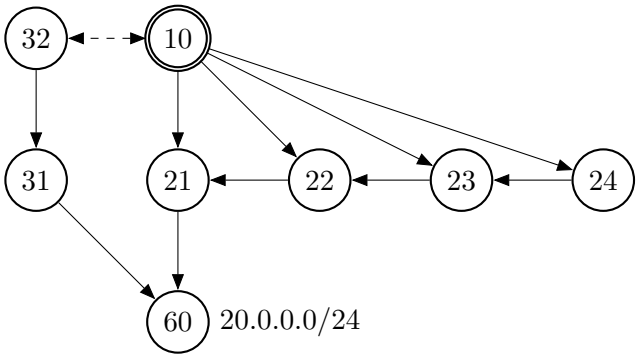

Abbildung 8: AS 60 announcing prefix 20.0.0.0/24 to its neighbors.

(i) The Internet has converged: Each AS knows at least one route to the prefix 20.0.0.0/24 advertised by AS 60. List all of the paths that AS 10 knows to the prefix 20.0.0.0/24. (2 Points)

(ii) All of a sudden the link between AS 60 and AS 21 fails. In the following, indicate each subsequent *withdrawal* that AS 10 observes as a consequence of that failure, the time at which it receives it, and the route that AS 10 is using after processing it. (2 Points)

Time	**Event/Withdrawal**	**AS** 10 **Path**
0	Failure	21 60

(iii) Can you think of a way to prevent the above situation and switch to the working path as soon as possible? (3 Points)

leave blank

e) Small Internet **(14 Points)**

Consider a small Internet topology made up of 8 Autonomous Systems (ASes). Each AS in this Internet consists of a single BGP router and applies the default selection and exportation BGP policies based on their customers, peers and providers.

(i) You are given the full routing tables of the ASes C and F in Figure 9. AS F, for example, knows two paths to B: A B, and D A B. Draw a connection between all ASes that are directly connected using Figure 10 below. (2 Points)

C

Dst	Path
B	B
E	E
G	E G
H	H

F

Dst	Path
A	A; D A
B	A B; D A B
D	D; A D
E	D E
G	G; D E G

Abbildung 9: Routing Tables of ASes C and F.

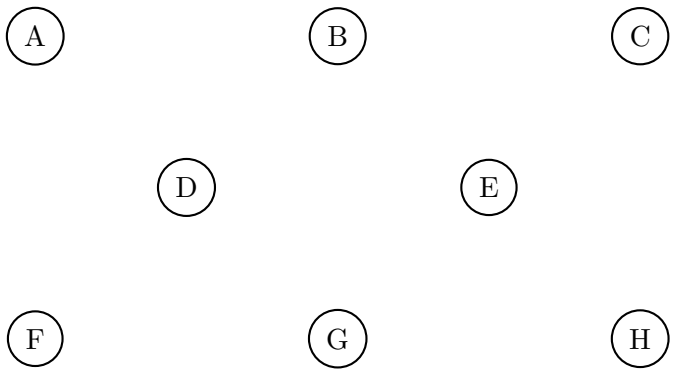

Abbildung 10: Small Internet topology with 8 ASes.

(ii) In practice, it is not only interesting to know whether two ASes are directly connected, but also to know the type of relationship they have between each other. In the following, we ask you to infer the type of the connections present in Figure 12.

D

Dst	Path
A	A; E A
B	A B; E A B
E	E; A E
G	G; E H G
H	E H

F

Dst	Path
B	B
C	C
G	H G
H	H

G

Dst	Path
B	H F B
C	H F C
D	D; H E D
E	H E
F	H F
H	H

Abbildung 11: Routing Tables of AS D, F, and G.

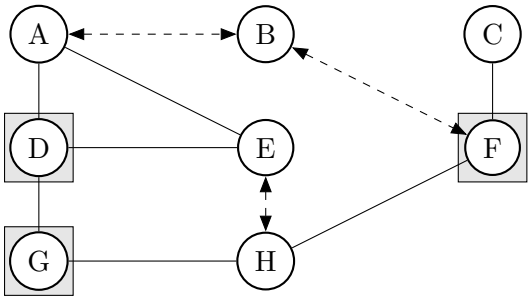

Abbildung 12: Another small Internet topology with 8 ASes.

The edges (A-B), (B-F), and (E-H) are peer/peer relationships (depicted by a dashed line). In addition, we provide you with the routing table of AS D, F, and G in Figure 11. Each AS also applies the default selection and exportation BGP policies based on their customers, peers and providers.

Start by identifying the type of the relationships of the edges (D-E), and (G-H). Complete the table by filling in the type of relationship (cust/prov, prov/cust or peer/peer) and provide a short justification. Pay attention to correctly specify the relationships. For example, X-Y, cust/prov means that X is the customer and Y the provider.

(4 Points)

Link	**Relationship**	**Justification**
D-E		
G-H		

(iii) Continue by identifying the relationships of the remaining edges. Note that for one edge the provided information is not enough to exactly determine its relationship. For that edge, use a question mark as relationship and explain why it is not possible to determine the precise relationship. (8 Points)

Link	**Relationship**	**Justification**
A-*D*		
A-*E*		
C-*F*		
D-*G*		
F-*H*		

leave blank

Task 4: Reliable Transport **29 Points**

leave blank

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 ☐ Using TCP to transfer the same file multiple times over the Internet will always generate the same amount of packets.

true ☐ false ☐ Consider a TCP connection in which the sender has received an ACK and removed the corresponding packet from its sender buffer. Afterwards, the sender can always send a new packet.

true ☐ false ☐ The receiving window protects the receiver buffer from overflowing whereas the congestion window protects the network devices from overflowing.

true ☐ false ☐ A network operator shows you a trace of packets belonging to one TCP connection captured over an Ethernet link. The packets have the following sequence numbers: 584, 1034, 5023, 5414. The operator claims that he did not miss any packets in between the observed ones. Is this statement true or false?

true ☐ false ☐ Similarly to the three-way handshake, TCP also needs at least three packets to properly terminate a (bi-directional) connection.

leave blank

b) Go-Back-N **(7 Points)**

Assume you have a Go-Back-N (GBN) sender and receiver. The receiver acknowledges *each* data segment with a cumulative ACK which indicates the next expected data segment. Furthermore, it saves out-of-order segments in a buffer. The sender and receiver buffer can contain four segments each. The timeout period is much larger than the time required for the sender to transmit four segments in a row. No improvement such as Selective Acknowledgment (SACK) or Selective Repeat is used.

(i) The sender wants to transmit 10 data segments (0,...,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. (2 Points)

Number of segments in the best case: _____

Possible segment that was lost: _____

Number of segments in the worst case: _____

Possible segment that was lost: _____

(ii) Once again, the sender wants to transmit 10 data segments (0,..., 9) to the receiver. This time, assume that exactly one ACK is lost. How many segments has the sender to transmit in the best (resp. worst) case and which corresponding ACK was lost?
(2 Points)

Number of segments in the best case: _____

Possible ACK that was lost: _____

Number of segments in the worst case: _____

Possible ACK that was lost: _____

(iii) 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)? (3 Points)

leave blank

c) Selective Acknowledgment (9 Points)

ACK: 112	b1_start: 115	b1_len: 2
padding	b2_start: 118	b2_len: 3
padding	b3_start: 0	b3_len: 1
padding	b4_start: 3	b4_len: 3

Abbildung 13: Observed SACK packet.

To debug problems with your Internet connection, you run Wireshark on an interface and observe the Selective Acknowledgment (SACK) packet in Figure 13. One SACK packet can contain at most four SACK blocks. Example: the first block (b1_start: 115, b1_len: 2) indicates that packets 115 and 116 were correctly received.

Based on the observed SACK packet, answer the following questions. If possible, answer with a precise value. Otherwise, indicate a range of values, e.g. ≥ 4. Include an explanation for each answer: a number alone does not give you full points.

(i) What is the size (in number of packets) of the sender window? (2 Points)

Sender window size: _____

Explanation: _____

(ii) How big is the current congestion control window (in number of packets)? (1 Point)

Congestion control window size: _____

Explanation: _____

(iii) How many packets were lost? (2 Points)

Number of lost packets: _____

Explanation: _____

(iv) What is the number of bits used for the sequence number? (You can assume that the maximal sequence number is at least twice as large as the sender window.) (2 Points)

Number of bits: _____

Explanation: _____

(v) Assume (for this question only) that the oldest packet in the sender buffer has a sequence number of 110 and that the sender and congestion window sizes did not change. How many *new* packets can the sender transmit after receiving the SACK packet in Figure 13? (2 Points)

Number of packets: _____

Explanation: _____

d) Congestion Control (8 Points)

leave blank

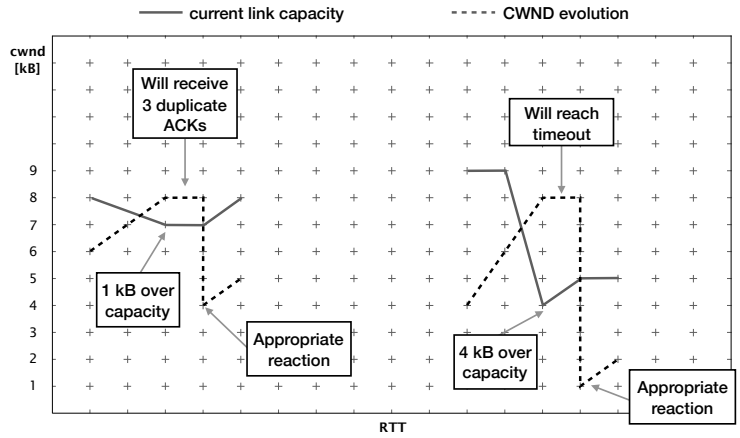

Abbildung 14: Reaction of the CWND (dashed line) if the current link capacity (continuous line) is exceed by at most 2 kB (left) or by more than 2 kB (right).

In this task, you will draw the Congestion Window (CWND) evolution in reaction to the available capacity of a link in a network. The CWND follows the well-known TCP congestion control algorithm using slow-start. Whenever the CWND value exceeds the current link capacity, the CWND algorithm will react in the following way:

1. The current CWND value is kept for the entire next RTT (no increase or decrease);

2a. If the current link capacity was exceeded by at most 2 kBs, the CWND algorithm will observe three duplicate ACKs during the next RTT and will react appropriately (Figure 14 left);

2b. If the current link capacity was exceeded by more than 2 kBs, the CWND algorithm will reach its timeout during the next RTT and will react appropriately (Figure 14 right).

Draw the CWND evolution in Figure 15 in reaction to the link capacity indicated with the continuous line. Start at the bottom left corner (RTT 1, CWND 1 kB) and assume that the CWND corresponds to a flow that just started, e.g. you are in the slow-start phase. You can stop once you reach RTT 22. To help you, a correct portion of the CWND is plotted between RTT 11 and 14 (dashed line).

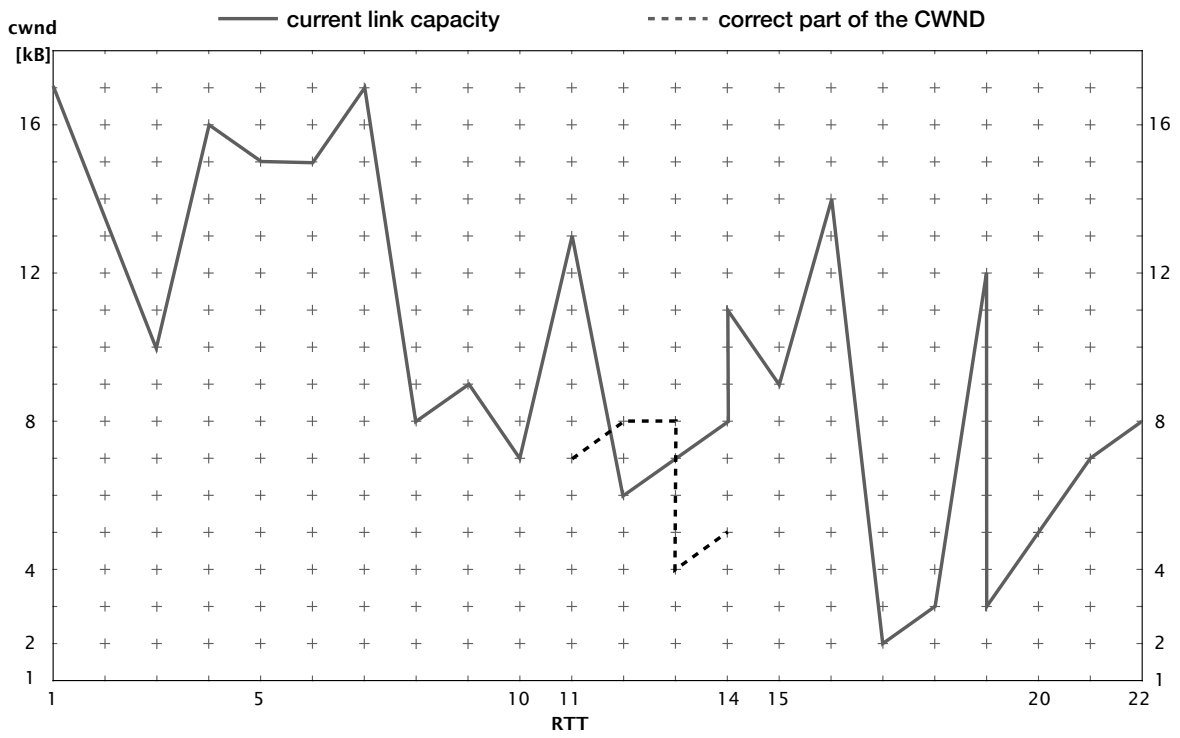

Abbildung 15: Complete the CWND evolution.

If you made a mistake, you can ask for an extra copy of Figure 15. **Important**: If you hand in an extra copy, we will *only* correct the solution on the extra sheet.

leave blank

Task 5: Applications **25 Points**

leave blank

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	
☐	☐	In DNS, "MX" entries specify the mailserver responsible for the domain.
☐	☐	In DNS, a hostname can have multiple "MX" entries with different IP addresses.
☐	☐	In DNS, it is possible to map `http://xy.ch` to a different IP address than `https://xy.ch`.
☐	☐	E-Mails can go through only one SMTP server.
☐	☐	`smtp://www.facebook.com:443/images/cat.mp3` is a Uniform Resource Locator (URL).

b) E-Mail analysis **(9 Points)**

leave blank

You received an email with the raw content shown in Figure 16.

```
1  Received: from edge20.ethz.ch (82.130.99.26) by CAS10.d.ethz.ch
2   (172.31.38.210) with SMTP Server (TLS) id 14.3.408.0; Thu, 2 Aug
3   2018 11:17:27 +0200
4  Received: from phil2.ethz.ch (129.132.65.3) by edge20.ethz.ch (82.130.99.26)
5   with SMTP Server id 14.3.408.0; Thu, 2 Aug 2018 11:17:23 +0200
6  Received: from filter.spam.ch ([5.152.185.154]  helo=filter.spam.ch)
7   by phil2.ethz.ch with esmtps  (TLSv1:AES128-SHA:128)   (Exim 4.69)
8   (envelope-from <john.doe@anonymous.ch>)        id  1fl9jO-0004C9-7T
9   for lvanbever@ethz.ch; Thu, 02 Aug 2018 11:17:15 +0200
10 X-Note: This Email was scanned by filter.spam.ch
11 Received: by filter.spam.ch with PIPE id
12  93122453; Thu, 02 Aug 2018 11:17:13 +0200
13 Received: from [10.40.0.131] (HELO smtp.ch.exg7.mailhost.com)  by
14  filter.spam.ch with ESMTPS id 93122443
15  for lvanbever@ethz.ch; Thu, 02 Aug 2018 11:17:10 +0200
16 Received: from exg7.mailhost.local (192.168.40.105) by
17  exg7.mailhost.local (192.168.40.107) with SMTP Server
18  (version=TLS1_2, cipher=TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256_P256) id
19  15.1.1531.3; Thu, 2 Aug 2018 11:17:09 +0200
20 From: Anonymous Student <john.doe@anonymous.ch>
21 To: Laurent Vanbever <lvanbever@ethz.ch>
22 Subject: Exam solutions
23 Date: Thu, 2 Aug 2018 09:17:09 +0000
24 Message-ID: <11A5442F-4D6E-436F-A873-2E3DA3656C06@anonymous.ch>
25 Accept-Language: de-CH, en-US
26 Content-Language: en-US
27 Content-Type: text/plain; charset="us-ascii"
28 Content-ID: <F43C7219ADA84040984B4640587C2B70@fwd7.mailhost.com>
29 Content-Transfer-Encoding: quoted-printable
30 MIME-Version: 1.0
31
32 Hey, can you give me the solutions for the exam?
```

Abbildung 16: Raw content of a received email

(i) According to Figure 16, what are the e-mail addresses of the sender and the receiver of this message? (1 Point)

Sender: _____

Receiver: _____

(ii) List the IP addresses of all servers that have seen this email according to Figure 16 in chronological order starting with the server that saw the email *first*. (2 Points)

_____

_____

_____

_____

(iii) According to the header in Figure 16, the email passed a spam filter (`filter.spam.ch`). Could one of the other servers have added this entry without the email actually passing `filter.spam.ch`? If yes: why and which server(s) could have done it? If no: why not?
(2 Points)

(iv) Which servers (according to Figure 16) could modify the email message ("Hey, ...")? Why?
(1 Point)

(v) Assume you have `telnet` access to an open SMTP server that does not appear in Figure 16 and you want to fake the email shown in Figure 16. That is, your goal is that the receiver of the email in Figure 16 receives the same email again (with the same sender). Which parts of the email in Figure 16 can you replicate in your email and which parts will be different? Use the line numbers in Figure 16 to list parts that are equal or different in your email and briefly explain the reasons why they are equal or different.
(3 Points)

c) Network debugging **(5 Points)**

leave blank

Figure 17 shows the output of `traceroute` from a host within ETH to `google.ch`.

```
1  $ traceroute google.ch
2  traceroute to google.ch (172.217.168.67), 64 hops max, 52 byte packets
3   1  82.130.102.1  1.038 ms  1.072 ms  0.672 ms
4   2  10.10.0.41  0.963 ms  0.824 ms  0.747 ms
5   3  10.1.11.129  0.985 ms  1.018 ms  0.867 ms
6   4  192.33.92.170  0.887 ms  0.826 ms  0.871 ms
7   5  192.33.92.11  0.833 ms  0.982 ms  0.975 ms
8   6  130.59.38.110  0.766 ms  1.002 ms  1.006 ms
9   7  72.14.195.4  4.369 ms  19.821 ms  1.224 ms
10  8  74.125.243.129  2.667 ms
11     74.125.243.113  1.805 ms
12     74.125.243.129  2.624 ms
13  9  64.233.175.167  2.562 ms
14     172.253.50.5  2.675 ms  2.696 ms
15 10  172.217.168.67  1.075 ms  1.088 ms  1.121 ms
```

Abbildung 17: `traceroute` output

(i) What is the meaning of the time measurements reported by `traceroute` (e.g. "1.038 ms" in line 3 in Figure 17)? (1 Point)

(ii) How many IP packets did the host need to send to obtain the results shown in Figure 17? Explain how you obtained the number. (1 Point)

(iii) How many ICMP packets did the device with IP address 192.33.92.11 (line 7 in Figure 17) generate because of this `traceroute` execution? Explain how you obtained the number. (1 Point)

(iv) Compare line 9 with the lines 10–12 of the output in Figure 17 and explain what these lines tell you about the paths between the host and `google.ch` (2 Points)

leave blank

d) NAT **(6 Points)**

Consider the network topology in Figure 18. Alice has multiple PCs at home (10.0.0.11–13) which share a single public IP address (1.2.3.4) via a NAT device. Further, she operates a surveillance camera server which is directly connected to the Internet with a public IP address (5.6.7.8). The camera transmits the live video signal as a stream of UDP packets with source port 1000 to a configurable destination IP address and port.

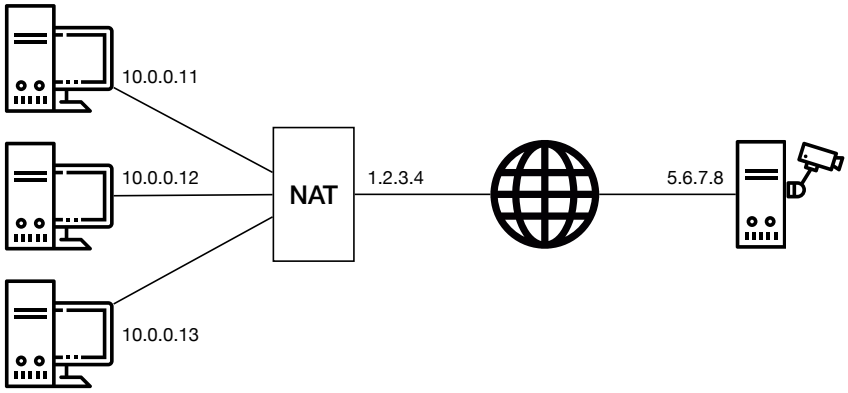

Abbildung 18: Alice operates three PCs and one camera server

(i) Alice wants to receive the live video stream on one of her PCs and thus configures the camera to send the video signal to IP 10.0.0.11 and port 1234. However, she does not receive it on her PC. Why? Where is this traffic sent to? (1 Point)

(ii) Now Alice configures the camera to send the video signal to IP 1.2.3.4 and port 1234. But she still does not receive it on any of her PCs. Why? Where is this traffic sent to? (1 Point)

(iii) What can Alice do such that she receives the video signal at her PC with IP address 10.0.0.11 and at port 1234 assuming that she *cannot* modify the configuration of the NAT? Describe step-by-step what she can do if she has the following possibilities:

- send one single UDP packet with arbitrary source and destination addresses and ports from each of her PCs;
- observe the received packets at each of her PCs and the camera server;
- specify the destination IP address and port for the video signal.

(4 Points)