ETH zürich
Institut für Technische Informatik und Kommunikationsnetze

Department ITET
August 2021

Prof. Dr. Laurent Vanbever
R. Birkner, C. Busse-Grawitz, T. Bühler, T. Holterbach, R. Yang

Exam: Communication Networks

27 August 2021, 14:30–17:00, Room HIL F 61

General remarks:

- Write your **name** and your **ETH student number** below on this front page and **sign it**.
- Put your **legitimation card** on the top right corner of your desk. Make sure that the side containing your name and **student number** is visible.
- Check that you have received **all task sheets** (Pages **1 − 31**).
- Do **not separate** the **task sheets** as we collect the exams **only after you left** the room.

- 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, use the three extra sheets at the **end of the exam**. Indicate the **task** in the corresponding field.

- **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, **place the exam face up on the top left corner** of your desk. Then collect all your belongings and **exit the room** according to the given instructions.

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	/30	
Intra-domain routing	/24	
Inter-domain routing	/39	
Reliable transport	/35	
Applications	/22	
Total	/150	

leave blank

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

a) Warm-up **(5 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	
☐	☐	Removing a single link in a spanning tree disconnects the spanning tree.
☐	☐	Since hosts of different VLANs are isolated, the spanning trees computed for each VLAN cannot overlap (share links).
☐	☐	Hierarchical addressing enables to add new hosts without changing existing or adding new forwarding rules.
☐	☐	Each host needs to have a default gateway to know how to forward packets towards destinations outside of the local network.
☐	☐	Hosts behind a NAT can only get private IP addresses, but not public ones.

b) STP: Spanning Tree Protocol **(7 Points)**

leave blank

Consider the layer-2 network composed of 6 switches depicted in Figure 1. For redundancy reasons, the network exhibits cycles and each switch therefore runs the Spanning Tree Protocol (STP). All links have a cost of 1. When equal cost paths to the root are encountered, switches break the tie based on the sender ID (lower is better).

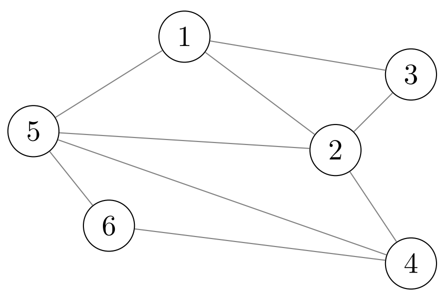

Figure 1: A layer-2 network with 6 Ethernet switches.

(i) In the figure below, cross all the links that end up deactivated in the final state, once all the switches have converged on the final spanning tree. (2 Points)

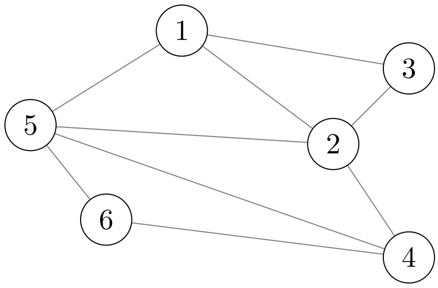

(ii) After convergence, how does the Bridge Protocol Data Unit (BPDU) message switch 6 sends to switch 4 look like? (1 Point)

BPDU(____ , ____ , ____)

(iii) List one benefit and one drawback of the Spanning Tree Protocol (STP). (2 Points)

Benefit: ______________________________

____

Drawback: ______________________________

____

(iv) Switch 5 is reaching the end of its life and is failing. Whenever it sends a BPDU message, it resets the hop count to zero (without changing anything else). Cross all the link(s) in the figure below that end(s) up deactivated in the final state after switch 5's failure. (2 Points)

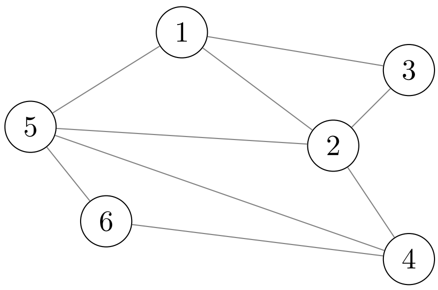

c) MAC-Learning **(8 Points)**

leave blank

Consider the Local Area Network (LAN) made up of 4 Ethernet switches in Figure 2. Several hosts (A, B, C, D, E) are connected to the switches. The MAC tables of all switches are still empty.

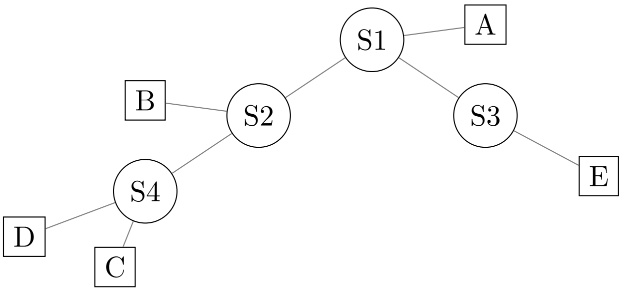

Figure 2: A layer-2 network with 4 Ethernet switches interconnecting 5 hosts.

(i) Host A sends a packet to host B. List below all the hosts that will receive the packet. In addition, fill in the MAC tables of all switches with the learned information. (3 Points)

Hosts receiving the packet: ______________________________

S1 MAC-Table	
dst	next hop

S2 MAC-Table	
dst	next hop

S3 MAC-Table	
dst	next hop

S4 MAC-Table	
dst	next hop

(ii) Host C sends a packet to host A. Again, list all the hosts that receive the packet and update the MAC tables with the learned information. (2 Points)

Hosts receiving the packet: ______________________________

S1 MAC-Table	
dst	next hop

S2 MAC-Table	
dst	next hop

S3 MAC-Table	
dst	next hop

S4 MAC-Table	
dst	next hop

(iii) After some time, the switches have full MAC-tables (i.e., they have an entry for each host in the network). Host B wants to hijack all the packets destined to host A. By only sending packets, how can host B manipulate the switches in the network to receive all that traffic? How many "manipulation" packets are minimally necessary and to which addresses does host B have to send them? Explain your approach, state the required number of manipulated packets, and list the source and destination addresses of all manipulated packets.

Note: The hosts are not aware of the other hosts and do not know the network's topology.

(3 Points)

Explain your approach: ____________________

Manipulated packets (use the format *[src, dst]*): ____________________

d) Internet Protocol (IP) (10 Points)

leave blank

(i) Given the prefix `234.100.15.128/26`, compute the number of addressable hosts, the prefix mask, the network address, the first and last host address and the broadcast address. (3 Points)

Number of addressable hosts: __________ Prefix mask: __________

Network address: __________ Broadcast address: __________

First host address: __________ Last host address: __________

Space for notes (not graded):

(ii) In the lecture, we have seen that one can reduce the size of a forwarding table by combining and filtering forwarding entries. Assume you have a forwarding table with 16 entries for /20 prefixes that fully cover a /16 prefix. The router has two interfaces. One half of the forwarding entries has interface 1 and the other half has interface 2 as output interface.

Consider different allocations of forwarding entries to the output interfaces. In the best case, how many forwarding entries do you need minimally while maintaining the same behavior? What about the worst case? For both cases, provide an example of how the forwarding entries need to be allocated to the output interfaces. (4 Points)

Best case example: ______________________________

____

Number of rules in the *best* case: ______________________________

Worst case example: ______________________________

____

Number of rules in the *worst* case: ______________________________

(iii) The Internet Protocol (IPv4) allows any routers along the path to fragment a packet if the outgoing link's Maximum Transmission Unit (MTU) is smaller than the total packet size. State two reasons why the recomposition of the fragmented packets happens at the end host and not on any router on the path. (3 Points)

Reason 1: ______________________________

____

____

Reason 2: ______________________________

____

____

leave blank

Task 2: Intra-domain routing **24 Points**

a) Warm-up (4 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 ☐ For a topology consisting of five nodes in a ring, link-state protocols converge faster than distance-vector protocols upon link failures.

true ☐ false ☐ In a distance-vector protocol with poisoned reverse, a node's advertisement always contains its minimum known cost for all destinations.

true ☐ false ☐ If a network running OSPF has fully converged and there is no failure, then a link that is *not* part of any shortest path does *not* see any packets.

true ☐ false ☐ For any positively weighted graph, taking the square root of all its link weights does not change the all-pairs shortest paths computed by Dijkstra's algorithm.

b) Applying (with) knowledge (4 Points)

leave blank

Run two iterations of Dijkstra and Bellman-Ford algorithm for the given network.

(i) Fill in the results of iteration 2 and 3 of Dijkstra's algorithm considering A as source node. (2 Points)

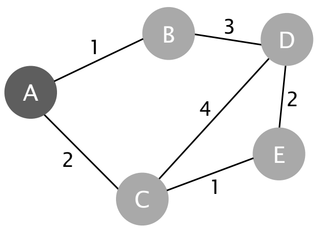

		$D(\cdot)$				
Iteration	Node Set S	A	B	C	D	E
1	A	0	1	2	∞	∞
2						
3						

(ii) Fill in the best distance each node knows **towards A** at the time step 2 and 3 of the Bellman-Ford algorithm. Each message takes exactly 1 time step to traverse a link. Each node can send arbitrarily many messages simultaneously.

(2 Points)

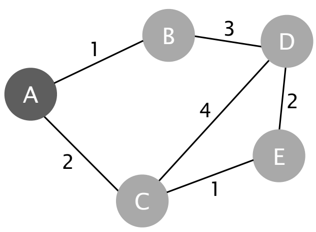

Time step	A	B	C	D	E
1	0	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$
2					
3					

c) Creative maintenance **(4 Points)**

leave blank

Consider the OSPF network depicted in Figure 3. For network maintenance you need to take down the link between node A and B and would like to move all the traffic away from it. Unfortunately, you **cannot access node A** anymore.

Instead, to move traffic away, you can perform two types of actions:

1. **add exactly one *unidirectional* link with a weight of your choice which does not start from or end at A;**
2. **reconfigure the unidirectional weight of any existing link except for the ones starting from A**. By unidirectional, we mean that you can set different weights to each direction of a link (if needed).

(i) List below the unidirectional link you would add together with the *smallest* possible number of unidirectional weight changes you could perform such that no shortest path contains the link between A and B. List each weight change as (X, Y, k) indicating that you change the weight of the link from X to Y to k. (4 Points)

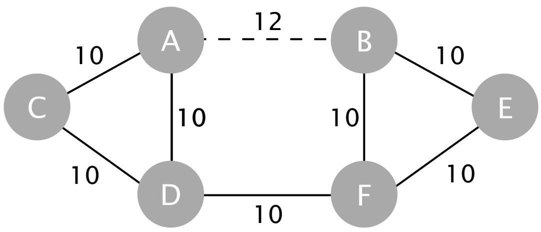

Figure 3: Weighted OSPF network in which the link between A and B needs to be taken down.

Adding link from ______ to ______ with weight ______

Unidirectional weight change(s):

____

____

____

____

____

d) Beyond Dijkstra and Bellman-Ford **(12 Points)**

leave blank

Consider a fully-centralized version of OSPF. Instead of each node running Dijkstra, each node sends their network view to one pre-selected "leader" router using statically-configured routes. The leader router then runs the *all-pairs* shortest path algorithm and pushes the results to all other nodes, again using statically-configured routes.

(i) Assume your network is running the new centralized routing algorithm. What is the biggest-possible impact on routing if: *(i)* one router goes down; and *(ii)* one link goes down? Explain what concretely could happen and the resulting impact. (2 Points)

Biggest impact following a router failure:

Biggest impact following a link failure:

(ii) Assume the all-pairs shortest path algorithm the leader node runs has a complexity of $O(n^3)$, where n is the number of nodes. Assume also that OSPF has a complexity of $O(nm + n^2 \log(n))$, where m is the number of edges.

Which algorithm scales better on: *(i)* full-mesh topologies; and *(ii)* tree topologies? *Briefly explain your reasoning.* (3 Points)

Most scalable algorithm for full-mesh topologies:

Most scalable algorithm for tree topologies:

(iii) Assume you have a graph G for which you know the shortest path between any pair of nodes. The shortest path between any node i and j in G is given by $\texttt{sp}(i, j, G)$, and its cost by $\texttt{cost\_sp}(i, j, G)$. Unlike in the previous question, we now assume that link weights are symmetrical: each link direction has the same weight.

Assume we add a new node x to G and connect it to some of the existing node(s). We denote the augmented graph with G'. Because we add a new node and edge(s), the shortest paths in G' might differ from the corresponding ones in G. (This is the case when going via the new node is shorter.) We therefore need to re-compute $\texttt{sp}(i, j, G')$ and $\texttt{cost\_sp}(i, j, G')$ for any pair of $i, j \in G'$.

The goal of this question is to figure out a way to do so while reusing as much as possible the paths/costs we know from G. You suspect that there is a way to only compute the shortest paths/cost from x (the new node we added) to all the other nodes in G' and update the rest of the shortest paths/costs by adapting the ones already computed in G.

To help you answer, consider any two nodes i and j in the original graph G. In G', the shortest path between i and j either does or does *not* go via the new node x.

(7 Points)

Consider first that the new shortest path between i and j does *not* go via the new node. Explain how you could compute $\texttt{cost\_sp}(i, j, G')$ based on $\texttt{cost\_sp}(i, j, G)$.

Consider now that the new shortest path between i and j *does* go via the new node. Explain how you could compute $\texttt{cost\_sp}(i, j, G')$ based on $\texttt{cost\_sp}(i, j, G)$ and the shortest paths/cost from x to all the other nodes in G'.

leave blank

Task 3: Inter-domain routing **39 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	
☐	☐	Just because a router receives a WITHDRAW message for a prefix P it does not necessarily need to change the forwarding entry for prefix P.
☐	☐	A router in an AS without any policies/route-maps receives two routes for the same prefix. It always prefers the route with the shorter AS path length.
☐	☐	A BGP route-map configured on the "in" direction for an eBGP session is applied to all incoming TCP packets received from this neighbor.
☐	☐	A router can set the local preference of outgoing routes to a high value to make sure that its neighbors in other ASes (eBGP peers) prefer these routes.
☐	☐	One reason why Resource Public Key Infrastructure (RPKI) is not yet deployed in each AS is the requirement that each router has to perform and support cryptographic operations in order to validate the origin of a route.
☐	☐	A router receives the following two BGP updates for the same prefix: Update 1 from neighbor AS 20: local pref 100, AS path [20 30 50], MED 10 Update 2 from neighbor AS 30: local pref 100, AS path [30 30 50], MED 20 The BGP decision process will always prefer message 1.

b) BGP and load-balancing **(10 Points)**

leave blank

Figure 4 shows an AS topology with 4 ASes. Arrows point from providers to customers and dashed double-headed arrows indicate peer connections. The ASes use the normal provider, peer, customer policies. AS 1 has two providers P1 and P2 and owns the prefix 14.122.160.0/19. Assume that each AS only consists of a *single* router.

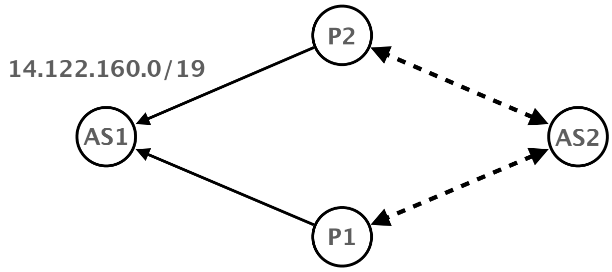

Figure 4: A network with 4 ASes and two paths between AS 1 and AS 2.

(i) Assume that AS 1 advertises 14.122.160.0/19 to both providers (P1 and P2). Using the standard BGP protocol and algorithm discussed in the lecture, how will traffic from AS 2 towards 14.122.160.0/19 reach AS 1? Will it be load balanced over P1 and P2 or rather only reach AS 1 over one provider? Explain. (2 Points)

(ii) AS 1 wants to have more control over the incoming traffic and starts to perform traffic engineering. Which prefix(es) does AS 1 have to advertise to P1 and P2 such that the following two conditions hold at the same time?

- A part of AS 2's traffic towards 14.122.160.0/19 enters AS 1 over P1 while the remaining traffic enters over P2.
- In case one of the two paths between AS 1 and AS 2 fails, AS 2 should still be able to reach all destinations in 14.122.160.0/19. (2 Points)

Prefix(es) advertised to P1: ________

Prefix(es) advertised to P2: ________

Provider P1 gets a new customer (AS 3) which hosts a popular video platform called myVideo. Figure 5 shows the extended topology. Unsurprisingly, AS 1 now starts to receive a huge amount of traffic over the link with provider P1 as a lot of AS 1's hosts start to watch videos on myVideo.

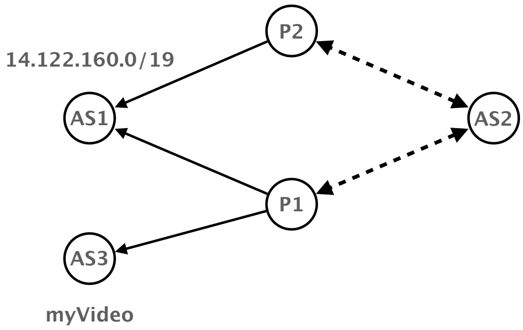

Figure 5: P1 has a new customer AS 3 which hosts a video platform.

(iii) Is AS 1 also able to load balance the traffic coming from myVideo over its two providers P1 and P2, e.g., using the same technique as in the previous question? If it is possible, explain which prefixes AS 1 has to advertise, otherwise explain why the myVideo traffic cannot be load balanced. (2 Points)

(iv) To simplify the load-balancing implementation, you can assume that you can change the type of a *single, existing* link in the given topology (e.g., a provider-customer link is now a peer-to-peer link). Which link do you have to change such that the existing load balancing solution from question (ii) also affects traffic coming from myVideo? Clearly indicate the new roles (customer, peer or provider) of the two ASes connected to the link for which you changed the type. (2 Points)

(v) In the real Internet, a single AS can obviously not just arbitrarily change the type of a link between ASes. Explain an approach/technique that ASes use in today's Internet to have more control and reduced costs when faced with huge amounts of traffic, e.g., from a video streaming platform such as Netflix. (2 Points)

c) Partial RPKI deployment **(9 Points)**

leave blank

Figure 6 shows an AS topology with 4 ASes. Arrows point from providers to customers and dashed double-headed arrows indicate peer connections. The ASes use the normal provider, peer, customer policies.

Unfortunately, only two of the four ASes (namely AS 1 and AS 4) have RPKI deployed. That means these ASes correctly register their prefixes in RPKI and validate all incoming routes. If a route is *invalid* according to RPKI it is dropped and not further considered. A registered prefix in RPKI protects the entire IP space of this prefix, e.g., a registered /8 prefix also covers a /16 prefix which is part of the /8.

AS 4 announces the prefix 20.0.0.0/16 to all its neighbors. AS 3 performs malicious activities and is hijacking a part of AS 4's prefix space, namely 20.0.128.0/17 with an AS path starting with AS 3.

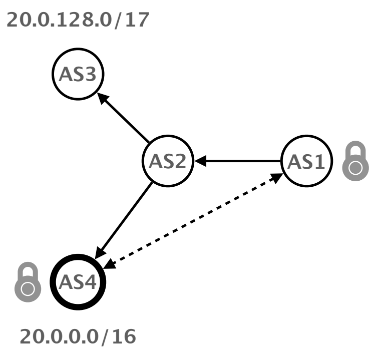

Figure 6: A network with 4 ASes and partial RPKI deployment (AS1 and AS4).

(i) A host in AS 1 sends traffic to 20.0.5.23. What path do the packets take (list the ASes)? (1 Point)

Path towards 20.0.5.23: ________________

(ii) Looking again at AS 1. What is the forwarding table of AS 1? I.e., what is the BGP next hop for each prefix it has knowledge of? *Hint: you might not need all the rows in the table below.* (2 Points)

prefix	**next hop** (AS number)

(iii) Hosts in AS 1 complain that they no longer can reach a popular video streaming website with IP 20.0.162.30. Explain what the problem is. Why can the hosts no longer reach this IP address even though AS 1 is using RPKI which should protect against basic BGP hijacks? (3 Points)

____

____

____

(iv) Explain one possible solution that AS 1 could perform (without any help from other ASes) to make sure that its hosts can once again reach 20.0.162.30. Are there any drawbacks of your solution? (3 Points)

Possible solution: ______________________________

____

____

Drawbacks (if any): ______________________________

____

____

d) BGP update messages **(14 Points)**

leave blank

You just started as a network operator at a new company which owns its own AS. To understand the network better, you look at all the *incoming* BGP messages on a *single* router X. You observe the messages in the table below and you know the following additional points:

- Router X has only iBGP sessions, no eBGP ones.
- The BGP update messages are shown in chronological order starting with the oldest message.
- All routers in your AS are connected in an iBGP full-mesh.
- Additional BGP attributes which are not shown in the table (e.g., MED value) are not important.
- {} indicates that the specific message has an empty AS path.
- The rightmost AS number in the AS path originates the announcement.
- Routers which have an eBGP session use next-hop-self.
- There are no malicious activities going on. For example, other ASes *do not* modify their AS paths, i.e., you can trust the shown information.

type	**prefix**	**next hop**	**local pref.**	**AS path**
UPDATE	37.55.128.0/20	123.200.0.4	10	50
UPDATE	5.3.128.0/24	123.200.0.6	200	120
UDPATE	47.33.2.0/24	123.200.0.4	50	70 44 99 30
UPDATE	207.126.0.0/16	123.200.0.2	10	50 632 153
UPDATE	99.176.77.0/24	123.200.0.3	50	60 943 163 23
UPDATE	123.10.4.0/24	123.200.0.5	100	{}
WITHDRAW	83.53.4.0/24	123.200.0.6	200	120 77 1664 423
UPDATE	207.126.0.0/16	123.200.0.4	10	50 632 153
UPDATE	37.55.128.0/21	123.200.0.2	10	50
UDPATE	47.33.2.0/24	123.200.0.3	50	60 643 30
UPDATE	99.176.77.0/24	123.200.0.4	120	70 163 23
WITHDRAW	47.33.2.0/24	123.200.0.3	50	60 643 30

(i) How many *additional* routers can you identify in your AS based on the observed BGP messages on router X? Explain how you found your answer. (2 Points)

Number of other routers: ______________________

Explanation: ______________________

____

____

(ii) How many eBGP sessions does your network have based on the observed BGP messages on router X? Explain how you found your answer. (2 Points)

Number of eBGP sessions: ______________________

Explanation: ______________________

____

____

(iii) Assume that initially router X had an empty forwarding table. How does the forwarding table of router X looks like after it received all the shown BGP messages? You can directly answer in the table below. If you do not know the next hop for a given prefix, indicate it with a question mark (?). *Hint: you might not need all the rows in the table.* (7 Points)

prefix	**next hop**

(iv) From the observed BGP messages, can you identify a case of traffic engineering? More precisely, does either your own AS or one of your neighbors perform traffic engineering such that the ingress or egress traffic follows a specific configuration? Indicate which BGP messages show traffic engineering and explain how you detected it. (3 Points)

leave blank

Task 4: Reliable transport **35 Points**

a) Warm-up (4 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	
☐	☐	The size of the TCP receiving window never changes throughout the duration of a connection.
☐	☐	In the slow start phase, the size of the congestion window (which starts from 1) will be 4 after 3 RTTs.
☐	☐	A client having an ongoing TCP connection to a server (IP 1.2.3.4, port 80) is able to start a UDP flow towards the same server and port (IP 1.2.3.4, port 80).
☐	☐	DNS queries are often done over UDP because it is fast and provides lower overhead compared to TCP.

b) Fight for bandwidth (13 Points)

leave blank

Consider the network depicted in the figure below, where R2 performs *per-flow* load balancing to R3, R4 and R5 for the packets destined to the destination server. Recall that *per-flow* load balancing means that packets of a particular flow will always follow the same (randomly selected) path.

In this exercise, we will always assume that the hosts as well as the server have unlimited resources to process and send packets to their directly connected router. Each link connecting two routers in the network is full-duplex (i.e., it allows to simultaneously send and receive data without affecting each other) with symmetrical bandwidth as labeled in the figure. The hosts do **not** send any additional traffic than the one explicitly mentioned in the questions.

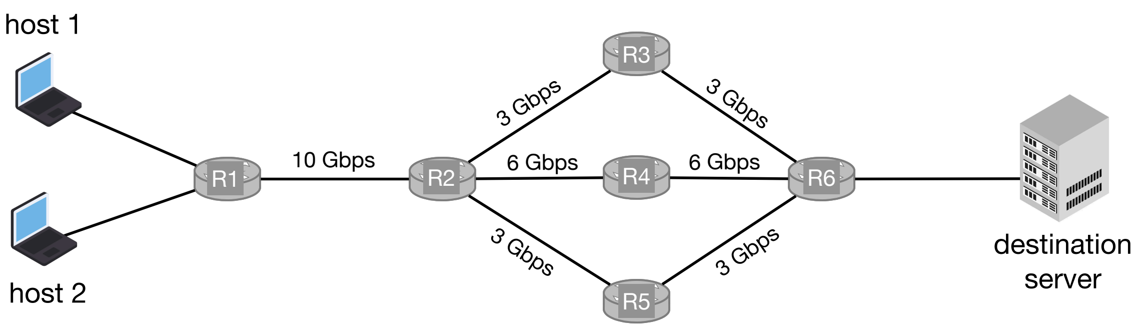

(i) Host 1 continuously sends more than 10 Gbps of data to the server using a UDP flow. The path used by the UDP flow is R1-R2-R4-R6. How does the bitrate measured on the link R1-R2 evolve over time? Pick the correct scenario among the four presented below. (2 Points)

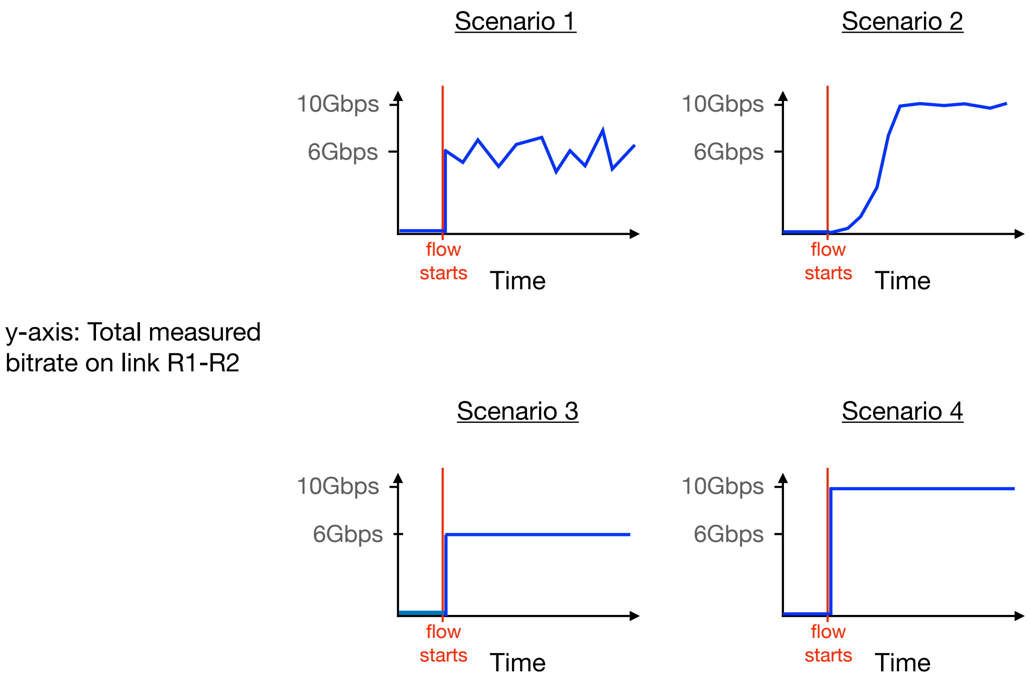

Scenario: ____

(ii) Host 1 has three large files to send to the destination server. It decides to send the three files in parallel using three distinct TCP connections. What would be the maximum and minimum achievable total throughput between host 1 and the server? Note that there is no other traffic in the network than the three TCP connections. (2 Points)

Maximum total throughput: ________________________________

Minimum total throughput: ________________________________

(iii) Assume now that host 1 is sending data using three distinct TCP connections with the maximum throughput indicated in **(ii)**. Meanwhile, host 2 starts to send one file to the destination server using TCP. What is the maximum achievable throughput for host 2? Which link is the bottleneck? (3 Points)

Maximum throughput for host 2: ________________________________

Bottleneck link: ______________________________________

(iv) Assume that the network operator has configured R2 to perform *per-packet* load balancing. This means R2 forwards packets destined to the server to either R3, R4 or R5, following a round-robin fashion and independently of which flow the packets belong to.

Host 1 continuously sends data at 12 Gbps to the server using a UDP flow. The UDP packets all have the same size, and there is no other traffic in the network. What is the expected throughput measured on the link between R6 and the destination server? (2 Points)

(v) The network operator decides to configure R6 to perform *per-packet* load balancing towards R3, R4 and R5. All the other routers perform *per-flow* load balancing, including R2. There is a gray failure happening on the link R6-R4, which means a small portion of the ACKs that traverse this link will be dropped randomly.

Assume host 1 and host 2 send data to the server with a reliable transport protocol which uses cumulative acknowledgements. Is it possible in this scenario that host 1 and 2 will observe a performance degradation? Explain your answer. (2 Points)

(vi) Now host 1 and 2 use a different transport protocol which uses independent acknowledgements instead (*i.e.,* every data packet is acknowledged independently, with one dedicated acknowledgement packet). Under the same scenario with the gray failure described above, is it possible that host 1 and 2 observe a performance degradation? Explain your answer. (2 Points)

c) GBN protocol **(14 Points)**

leave blank

Figure 7 on the next page shows the beginning of a Go-Back-N (GBN) time-sequence diagram. Here is a **non-exhaustive** list of implementation choices made for the GBN sender and receiver implementation. **Read them carefully.**

- The sender and receiver window have a size of 4 packets;
- 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 [A9, A10, A10, A10], it will immediately retransmit the data segment D10;
- 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 6 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.

(i) **Your task:** Draw the successful transmission of 7 data segments (D0 to D6) if:

- the third data segment (D2) is lost **the first two times it is transmitted** (the first time it is lost, is already indicated in the diagram);
- the sixth ACK (A6) is lost **as indicated in the diagram**.

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.

Note: We provide a backup copy of Figure 7 at the end of the exam. Please indicate clearly which diagram contains your final solution. (10 Points)

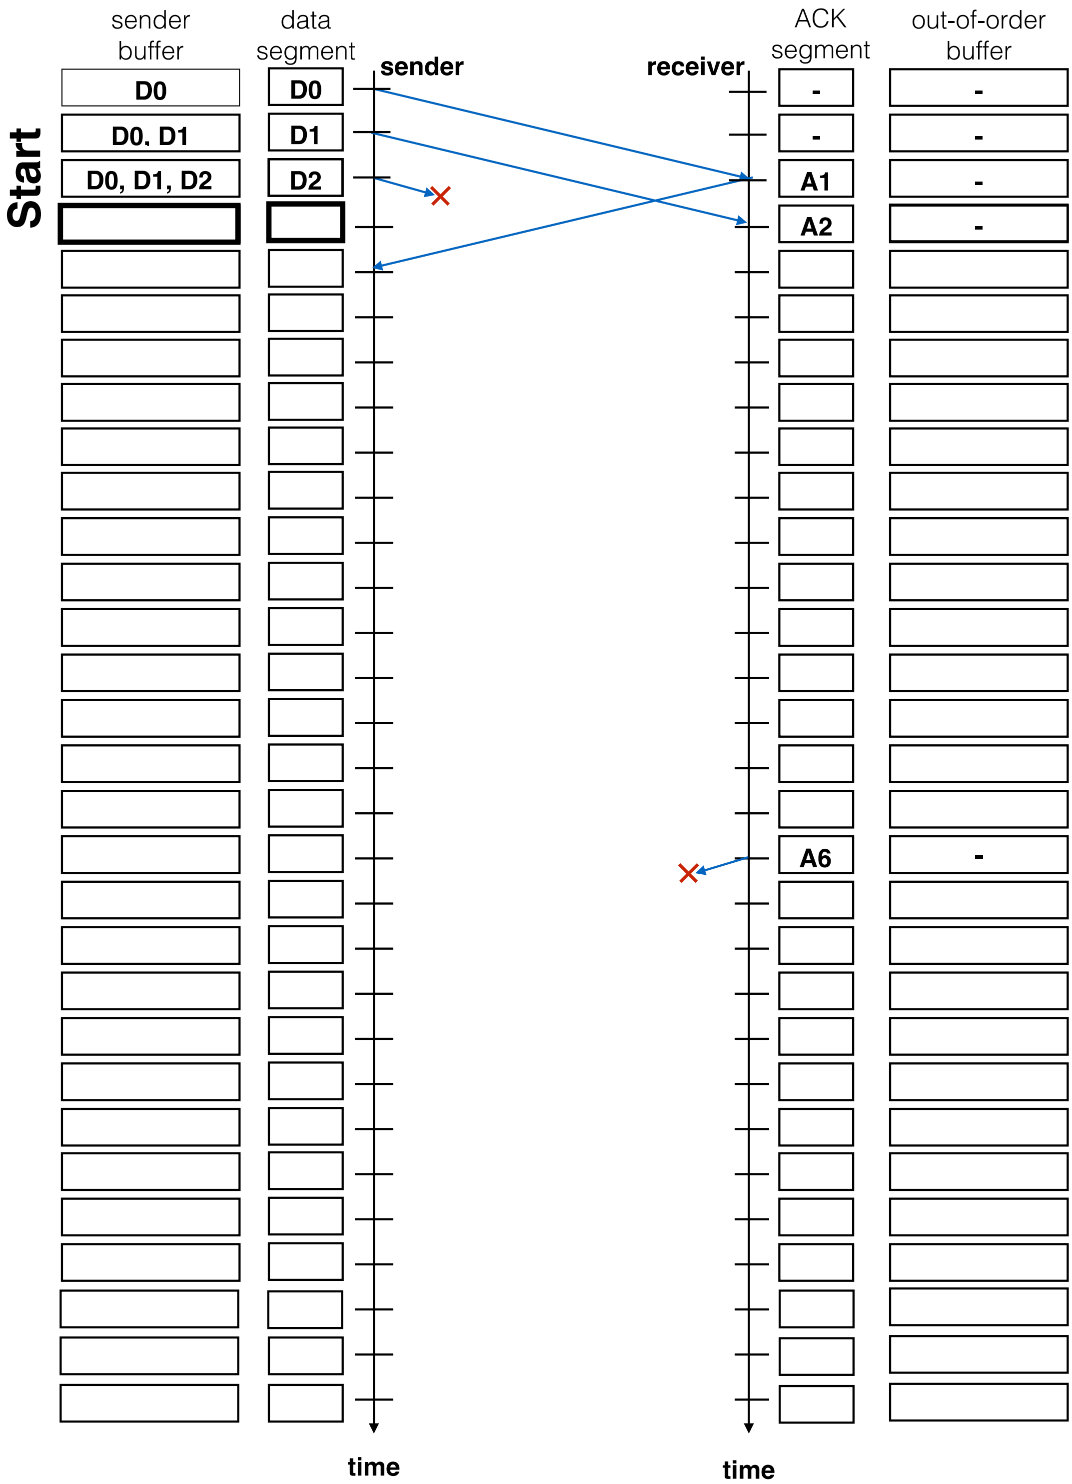

Figure 7: Time-sequence diagram of a GBN protocol with Fast Retransmit and cumulative ACKs.

(ii) Assume you have a network with packet loss, but no reordering. The sender and receiver use a window size of 4 packets and the maximum sequence number is much larger than 4. The sender starts to transmit data to the receiver. The first data segment has sequence number 0. Assume that the receiver **is about to** send an ACK with sequence number 4 to the sender. At this moment, what is the smallest sequence number among all the packets that are still in the sender buffer? Write down that sequence number both under the best and worst network condition. (4 Points)

Smallest sequence number for the best network condition: ________________

Smallest sequence number for the worst network condition: ________________

d) Congestion control **(4 Points)**

leave blank

In this question, we consider a misbehaving TCP receiver. Upon receiving a data segment which contains B bytes, the TCP receiver divides the resulting acknowledgment into M separate acknowledgments (where $M \leq B$) that it sends back to the sender. Each acknowledgement covers one of M distinct pieces of the received data. For example, if the receiver receives a data segment containing bytes 1000 to 2000 and $M = 2$, it will send 2 ACKs with sequence number 1501 and 2001.

We now evaluate how the congestion window size increases if a TCP sender receives ACKs from this misbehaving TCP receiver. You can assume that the sender is constantly in the slow start phase and will increase the congestion window by one whenever it receives an ACK acknowledging bytes in the current sender window.

(i) If $M = 4$, what is the congestion window size after **1** RTT assuming its initial size is set to 1? (1 Point)

Congestion window size: ________________

(ii) What is the congestion window size after **n** RTTs when the receiver divides each acknowledgment into **M** distinct ones? Assume that the initial congestion window size is once again 1 and use **M** and **n** to express the final congestion window size. (3 Points)

Congestion window size: ________________

leave blank

Task 5: Applications 22 Points

leave blank

a) Warm-up (10 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 ☐ Without DNS root servers, no DNS names could ever be resolved.

true ☐ false ☐ Setting short Time-To-Live (TTL) values on A or AAAA records increases the query load on the root name servers.

true ☐ false ☐ Different DNS resolvers might learn different IP addresses for the same DNS name (e.g., `comm-net.ethz.ch`).

true ☐ false ☐ Assume you type in `ping comm-net.ethz.ch` in a terminal and that your local DNS resolver has a cached entry (with a positive TTL) for one of the authoritative servers of `ethz.ch`. There is no other entry in the cache. Your resolver will not contact any other name server but the cached one.

true ☐ false ☐ The authoritative DNS resolver of `ethz.ch` can return a different IP address for `www.ethz.ch` depending on whether you visit `http://www.ethz.ch` or `https://www.ethz.ch` in your browser.

true ☐ false ☐ Forward and reverse HTTP proxies act in turn as an HTTP client and as an HTTP server.

true ☐ false ☐ HTTP servers can track client's requests when they use persistent HTTP connections.

true ☐ false ☐ It is possible to run an HTTP server on a random TCP port (i.e., not port 80 or port 443), but then the clients need to indicate the port explicitly in the URL in order to establish a connection to the said server.

true ☐ false ☐ A successful e-mail exchange between two distinct e-mail addresses will necessarily see more than one SMTP server.

true ☐ false ☐ The Sender Policy Framework (SPF) uses the DNS to store the IP prefixes of the hosts allowed to send an email from a particular domain.

b) DNS Setup **(12 Points)**

leave blank

Assume you are the operator for the DNS sub-domain `ee.ethz.ch`. The responsible of the Networked Systems Group (NSG) reaches out to you as they want to create and independently maintain their own DNS sub-domain `nsg.ee.ethz.ch`. Among others, they want to create DNS names for:

- their group webpage `www.nsg.ee.ethz.ch`, pointing at `129.132.30.1`;
- a project webpage `fun.nsg.ee.ethz.ch`, also pointing at `129.132.30.1`;
- their chat room `chat.nsg.ee.ethz.ch`, pointing at `129.132.30.2`; and
- a mail server `mail.nsg.ee.ethz.ch`, pointing at `129.132.30.3` and which should receive any email sent to a `@nsg.ee.ethz.ch` email address.

Unfortunately, the responsible of NSG does not seem to understand much about DNS. He needs your help to set this up.

You start by setting up two DNS name servers `ns1.nsg.ee.ethz.ch` and `ns2.nsg.ee.ethz.ch` which you host on `129.132.20.1` and `129.132.20.2`, respectively.

(i) Indicate the resource records that these two DNS servers should store. For each record, indicate its corresponding name, type, and value.

Note you might not need all 6 records. (4 Points)

Record #1. Name: ____________ Type: ______ Value: ____________

Record #2. Name: ____________ Type: ______ Value: ____________

Record #3. Name: ____________ Type: ______ Value: ____________

Record #4. Name: ____________ Type: ______ Value: ____________

Record #5. Name: ____________ Type: ______ Value: ____________

Record #6. Name: ____________ Type: ______ Value: ____________

(ii) Indicate the resource records that need to be added (if any) to the name servers responsible for `ee.ethz.ch`. Briefly explain the usage of these extra records or why no such extra record is needed.

Note you might not need all 4 records. (2 Points)

Record #1. Name: ____________ Type: ______ Value: ____________

Record #2. Name: ____________ Type: ______ Value: ____________

Record #3. Name: ____________ Type: ______ Value: ____________

Record #4. Name: ____________ Type: ______ Value: ____________

Explanation: ____

____

(iii) Indicate the resource records that need to be added (if any) to the name servers responsible for `ethz.ch`. Briefly explain the usage of these extra records or why no such extra record is needed.

Note you might not need all 4 records. (2 Points)

Record #1. Name: ______________ Type: ____ Value: ______________

Record #2. Name: ______________ Type: ____ Value: ______________

Record #3. Name: ______________ Type: ____ Value: ______________

Record #4. Name: ______________ Type: ____ Value: ______________

Explanation: ______________

(iv) While setting up the records, you observe that NSG wants to host two webservers (for `www` and `fun`) on the same IP address (`129.132.30.1`). Is that possible? If so, briefly explain how this works. If not, briefly explain why. (2 Points)

(v) It looks like you underestimated NSG's popularity! It appears the two name servers `ns1.nsg.ee.ethz.ch` and `ns2.nsg.ee.ethz.ch` are completely overloaded with (legitimate) DNS requests.

Explain two *distinct* techniques you could use to scale the system further. (2 Points)

Technique 1: ______________

Technique 2: ______________

Extra copy of Figure 7

In case you made a mistake you can use this copy of the GBN diagram for question *4 c) (i)*. Clearly indicate which diagram contains your final solution.

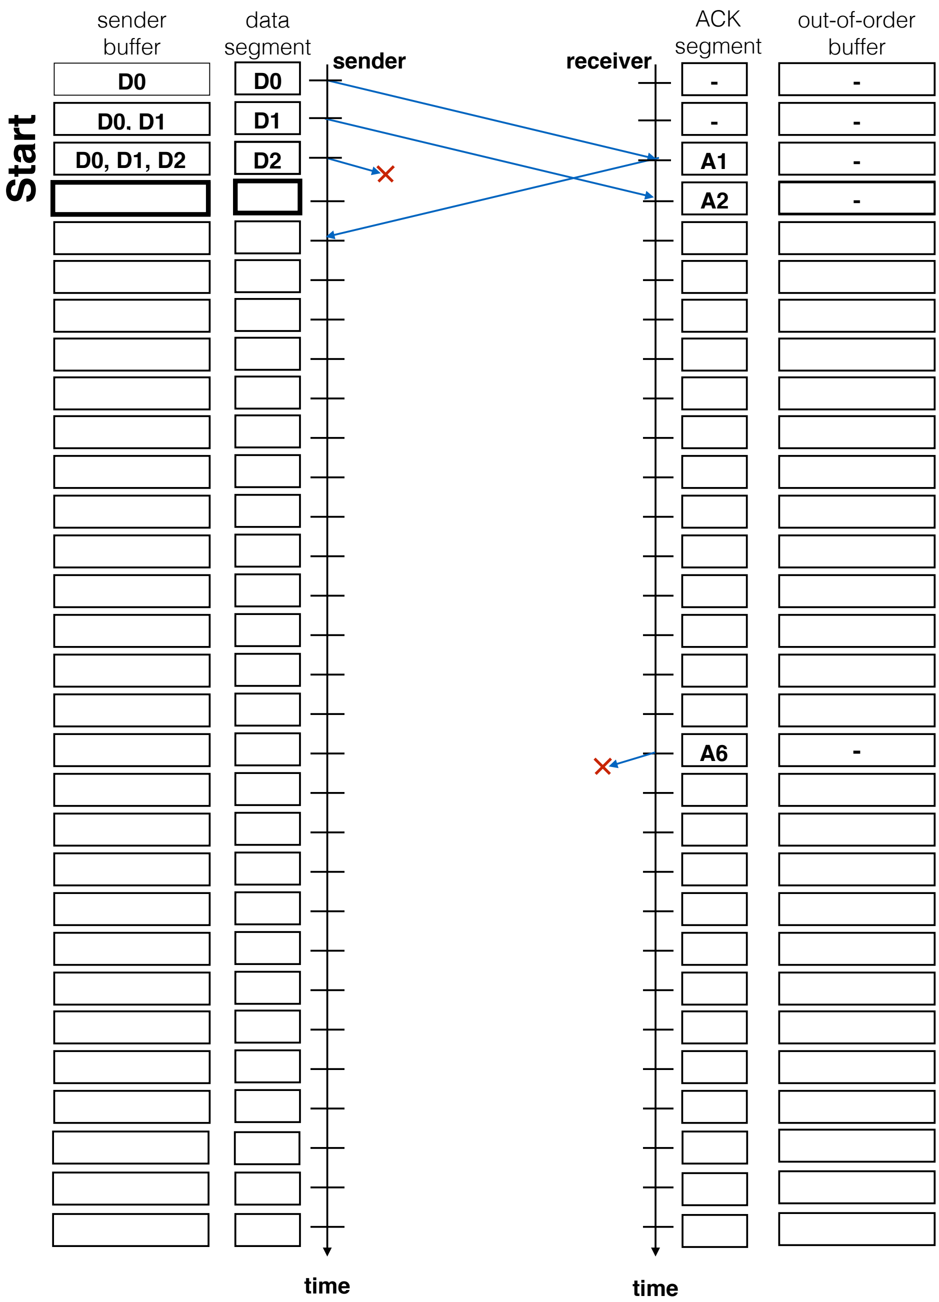

Figure 8: Copy of the time-sequence diagram in task *4 c) (i)*.

Extra Sheet 1

In case you need more space, use the following pages. Make sure to always indicate the task to which the answer belongs (e.g., *3 d) (ii)*).

Task: ______________

Task: ______________

Extra Sheet 2

Task: ________________

Task: ________________

Extra Sheet 3

Task: ______________

Task: ______________