ETH zürich

Institut für Technische Informatik und Kommunikationsnetze

Department ITET
February 2022

Prof. Dr. Laurent Vanbever
C. Busse-Grawitz, T. Bühler

Exam: Communication Networks

17 February 2022, 13:30–16:00, Room HIL E 6

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	/25	
Inter-domain routing	/38	
Reliable transport	/36	
Applications	/21	
Total	/150	

leave blank

Task 1: Ethernet & IP 30 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 □ Consider two hosts A and B located in distinct VLANs. A may receive some of B's broadcasted Ethernet frames during the flooding step of the MAC learning algorithm run by the switches.

true □ false □ Unlike plain L2 switches which do not require any configuration (they are plug-and-play), VLAN-enabled L2 switches must be configured to support multiple VLANs.

true □ false □ 192.33.88.0/255.255.240.0 contains 192.33.95.0/255.255.255.0

true □ false □ An IPv4 NAT can support more than 2^{16} concurrent connections to the same external IP and port.

true □ false □ IPv4 and IPv6 are incompatible with each other.

true □ false □ P1 is an IPv4 packet with n options and P2 is an IPv6 packet with m options. If $m > n$, then a router always needs more time to process P2 than P1.

b) Address Resolution Protocol (ARP) (6 Points)

leave blank

Consider the small network consisting of two hosts (A and B) and a router R, which is connected to the Internet as shown in Figure 2. The default gateway of both hosts is the router.

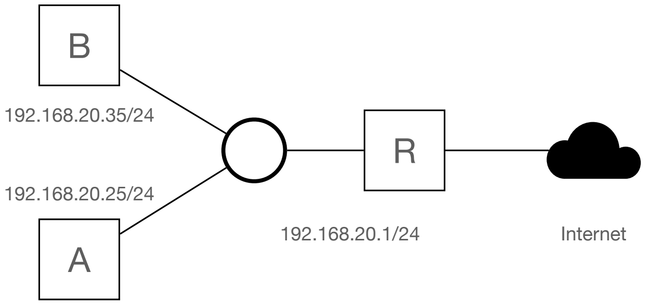

Figure 1: Two hosts connected through a router to the Internet.

(i) Explain what the Address Resolution Protocol (ARP) is used for. (1 Point)

(ii) All the ARP tables are empty. Host A (192.168.20.25) is trying to send a packet to host B (192.168.20.35). Does it need to send an ARP request? If yes for which address? If no, why not? (1 Point)

(iii) Now, host A (192.168.20.25) wants to send a packet to 8.8.8.8. Does host A need to send an ARP request? If yes, for which address? If no, why not? (1 Point)

(iv) Now, host A wants to send a packet to 1.1.1.1. Again, does host A need to send an ARP request? If yes for which address? If no, why not? (1 Point)

(v) Curious, you start monitoring the ARP traffic. You realize that the ARP replies of the default gateway change between two MAC addresses. The network operator explains you that they are using a technique called Proxy ARP. Basically, they use a virtual IP for the default gateway that actually represents two different (physical) routers to which the MAC addresses belong that you observed. Explain one possible reason for using such a setup. (2 Points)

c) Best-case vs worst-case scenario (8 Points)

leave blank

Consider the following layer-2 network in which switches run the distributed Spanning-Tree Protocol (STP), as seen in the lecture, and where all links have a cost of 1. This network interconnects two hosts, host 1 and host 2, and an IP router which acts as default gateway for the hosts. In case a switch receives two BPDUs from different neighbors with equal cost, it selects the neighbor with the lower identifier.

In the following, we ask you to indicate the *best* and *worst* possible position(s) for the root switch distinguishing between two cases:

1. host 1 and host 2 are located in the same IP subnet, e.g. 192.168.1.0/24.
2. host 1 and host 2 are located in distinct IP subnets, e.g. 192.168.1.0/24 for host 1 and 192.168.2.0/24 for host 2.

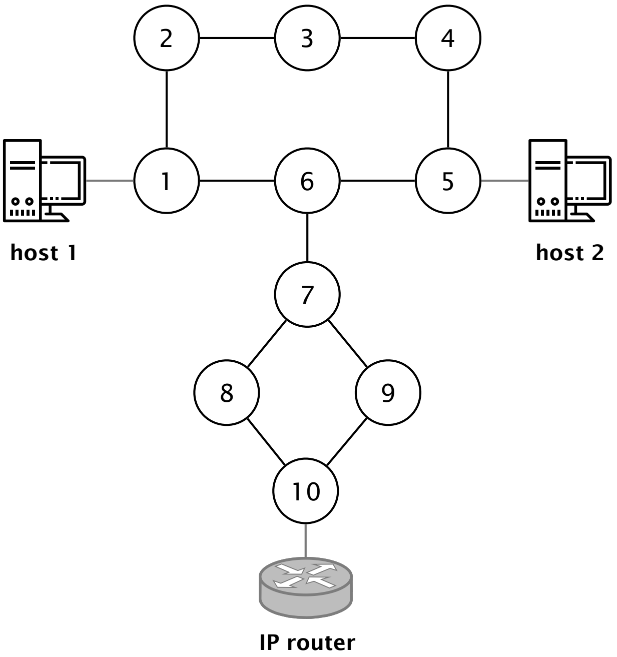

Figure 2: Two hosts and an IP router connected through a layer-2 network.

For each case, indicate the identifier(s) of *all* the possible best (worst) positions alongside with the number of inter-switch links that will be crossed when host 1 communicates with host 2. (You can consider that links connecting hosts/router to switches have a cost of 0.)

Case 1:

All the possible **best** positions for a root switch when the hosts are in the **same** subnet:

Number of inter-switch links crossed when host 1 communicates with host 2:

Case 2:

All the possible **worst** positions for a root switch when the hosts are in the **same** subnet:

____

____

Number of inter-switch links crossed when host 1 communicates with host 2: __________

Case 3:

All the possible **best** positions for a root switch when the hosts are in **distinct** subnets:

____

____

Number of inter-switch links crossed when host 1 communicates with host 2: __________

Case 4:

All the possible **worst** positions for a root switch when the hosts are in **distinct** subnets:

____

____

Number of inter-switch links crossed when host 1 communicates with host 2: __________

d) Compressing forwarding tables (10 Points)

leave blank

Consider an IP router with a forwarding table composed of the following 12 entries. You can assume that any IP packet this router will ever receive will be matched by exactly one of these entries.

ID	prefix	next hop
1	148.187.192.0/21	0
2	148.187.193.0/24	0
3	148.187.194.0/24	0
4	148.187.196.0/22	1
5	148.187.196.0/24	0
6	148.187.197.0/24	1
7	148.187.198.0/24	1
8	148.187.199.0/24	1
9	148.187.204.0/22	0
10	148.187.204.0/24	0
11	148.187.205.0/24	0
12	148.187.206.0/24	1

(i) Write down the identifier(s) (between 1 and 12) of the rules you can remove from the forwarding table *without* changing its forwarding behavior. (7 Points)

Identifiers of the rules removed: ____________________

____

____

(ii) Assume now you can add forwarding rules *in addition to* remove them. What change(s) (addition(s), removal(s)) would you propose to reduce the forwarding table even further? For each rule you add (if any), indicate the prefix in addition to the next hop. For each rule you remove (if any), indicate its identifier in the table above.
Your final forwarding table should contain the minimal number of rules to achieve the same forwarding behavior for the prefix space covered in the table above. However, your solution can also cover additional prefix space. (3 Points)

Rules added : ____________________

____

____

Identifiers of the rules removed: ____________________

____

Task 2: Intra-domain routing 25 Points

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 □ Link-state protocols rely on the Bellman-Ford algorithm.

true □ false □ Decreasing the time between OSPF `HELLO`s always speeds up the detection of a failed link as failed.

true □ false □ The network in Figure 3 runs a distance-vector protocol *with* poisoned reverse. Link A changes its weight from w_1 to w_2. The convergence time is *completely* independent of the weight difference $w_2 - w_1$.

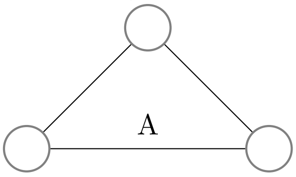

Figure 3: Network that runs a distance-vector protocol.

true □ false □ If one increases the weight of a link that is currently *not* part of any shortest path, this will not change the all-pairs shortest paths computed by Dijkstra.

true □ false □ For distance-vector protocols, nodes need to store more state than for link-state protocols.

true □ false □ Assume a graph G. In its set of all-pairs shortest paths, (at least) one path has the maximum number of hops. This number of hops corresponds to the maximum number of iterations of the Dijkstra algorithm on that graph G (that is, how often the node set is updated).

b) Vanishing links (9 Points)

In this task, you analyze the same topology in three different failure scenarios. The network runs OSPF to route packets.

Use **Figure 4** to fill in *all* your solutions. Optionally, use Figure 5 to sketch ideas.

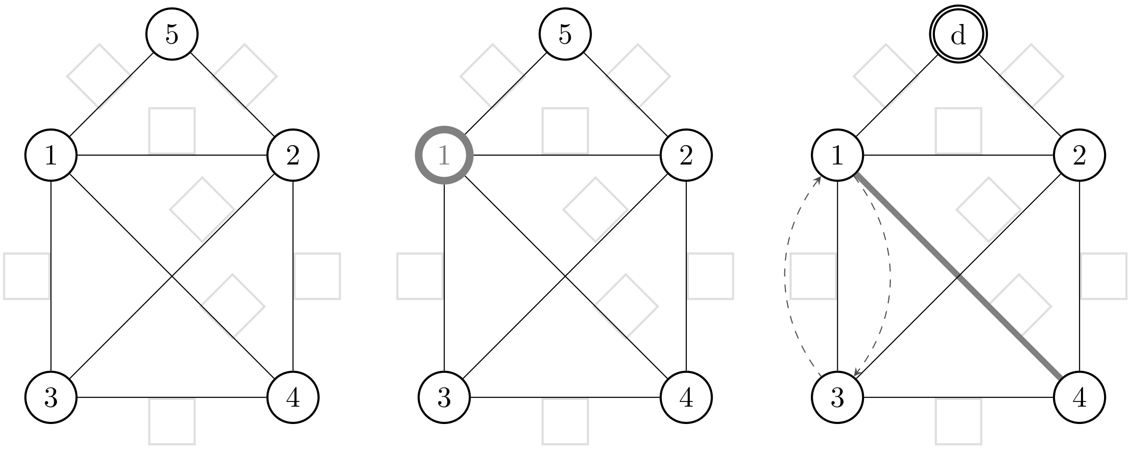

No blackholing, question (i) Node failure, question (ii) Create loops, question (iii)

Figure 4

(i) Enter a weight selection (in the light gray squares) that results in the maximum possible number of simultaneous link failures such that no packets are lost due to temporary blackholes during the convergence process. Cross the edges with an "X" that could fail simultaneously for this weight selection. (3 Points)

(ii) Fill in weights such that no transient forwarding loops exist if node 1 (in bold) fails. (2 Points)

(iii) Now, only consider the traffic towards node **d**. Find weights such that there *is* a transient forwarding loop between node 1 and 3 if the link between 1 and 4 fails (in bold). These weights should not result in any other forwarding loops. (4 Points)

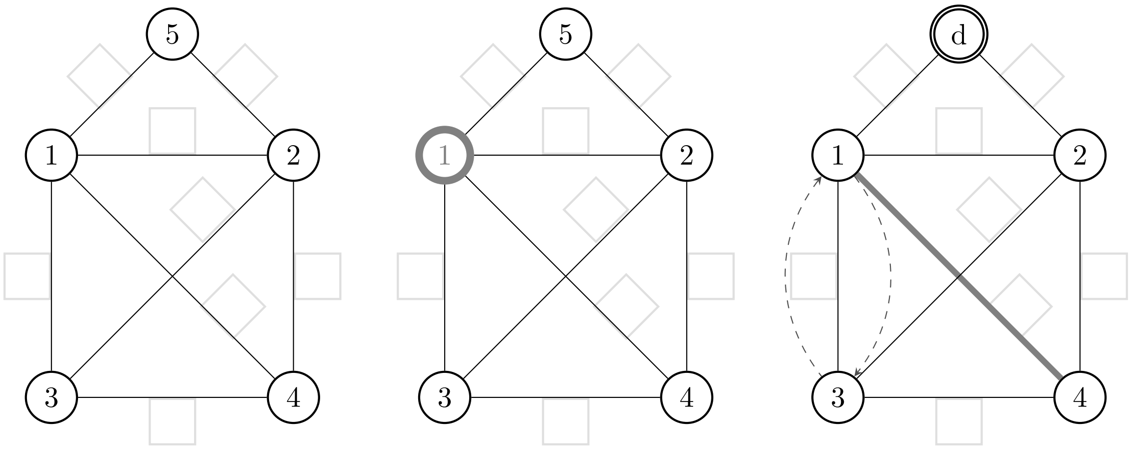

Figure 5: **You can use this figure for sketching. Solutions in this figure will NOT be corrected.**

c) Shortest paths, revisited (10 Points)

leave blank

In this task, you analyze possible modifications to Dijkstra (the algorithm that OSPF uses).

(i) Consider that the total cost of a path is no longer the addition of the per-link costs, but the multiplication ($w = w_1 \cdot w_2 \cdot w_3 \ldots \cdot w_n$ instead of $w = w_1 + w_2 + w_3 \ldots + w_n$). Describe how one would need to modify the given Dijkstra algorithm below to find the optimal path (i.e., with the lowest total cost) in such a way. Assume that all weights are bigger than 1. (Only consider this modification for this concrete question, the following questions again assume an additive distance.) (2 Points)

```
S = {u}
for all nodes v:
    if (v is adjacent to u):
        D(v) = c(u, v)
    else:
        D(v) = infinity

while not all nodes in S:
    add w with the smallest D(w) to S
    update D(v) for all adjacent v not in S:
        D(v) = min{D(v), D(w) + c(w,v)}
```

(ii) You are annoyed that a single link outage triggers a Dijkstra recomputation for the *entire* network. A friend suggests that if the link between node A and B goes down, you actually only need to recompute the shortest path between A and B. A and B can then pretend that there was no outage, since they now have an alternative to the link that went down. Therefore, all other nodes that were relying on the A-B connection can also continue to use it. Explain what scenario could result in an enormous packet loss if you apply this idea. (Assume that the network remains connected after the outage.)

(3 Points)

(iii) Your company has recently bought a (non-stationary) satellite to connect the nodes A, B and C (Figure 6). You model the link weights with the distance. Since the satellite is constantly moving, this also causes the link weights to constantly change. Hence, you cannot simply recompute Dijkstra for every link weight change. However, you realize that for this concrete topology, one only needs to recompute Dijkstra if the sum of some satellite links crosses certain thresholds (assuming that all non-satellite links keep their weights). Fill these conditions in the following template. Note that the satellite link weights stay between (including) 1 and 10.

Hint: First, ignore the satellite links, and draw the Dijkstra output for that case for source nodes A, B, C respectively. (You can e.g., use Figure 7 for this.) Then, consider how the satellite links can affect these paths.

(5 Points)

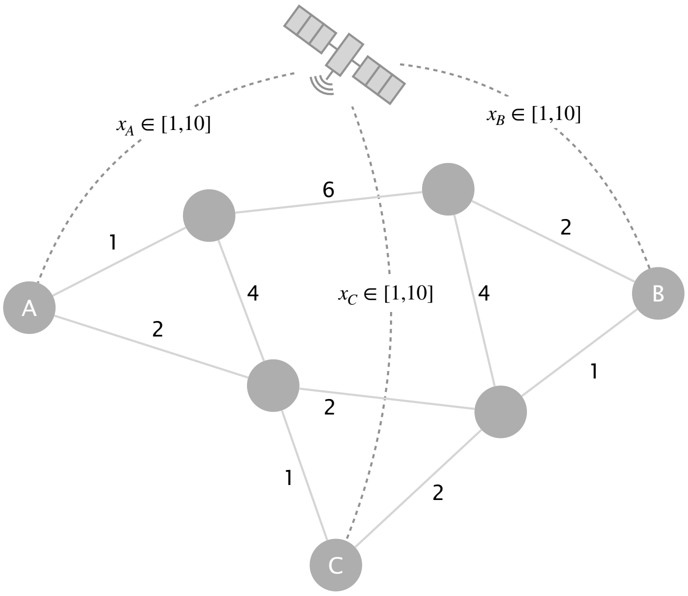

	sum of weights				*threshold*
	□	+	□	crosses	□
	□	+	□	crosses	□
	□	+	□	crosses	□
e.g.	x_A	+	x_B	crosses	42

Figure 6: Network with satellite links (dashed)

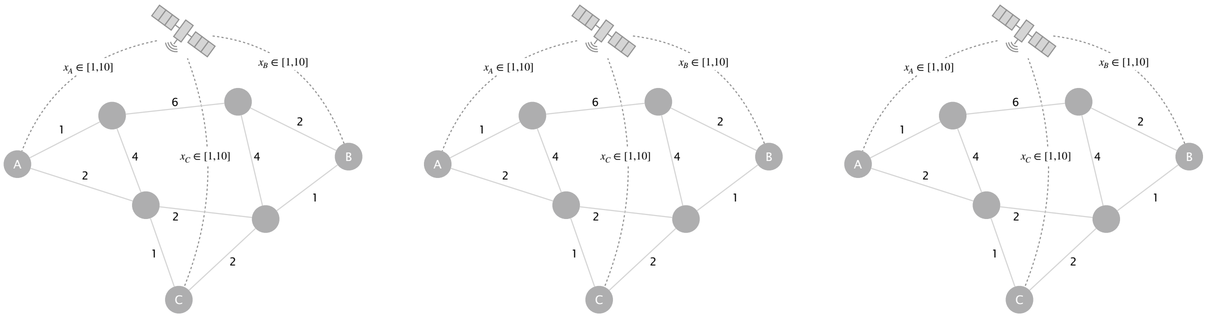

(a) (b) (c)

Figure 7: **You can use this figure for sketching. Solutions in this figure will NOT be corrected.**

leave blank

Task 3: Inter-domain routing 38 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 □ A Tier 1 AS can be directly connected to a Tier 3 AS.

true □ false □ Given BGP's loop-prevention mechanism (it looks for its own AS number in the announced AS path), it is impossible that forwarded packets will end up in a forwarding loop.

true □ false □ In the lecture we have seen the concept of "asymmetric routing", i.e. the forwarding path from A to B can be different than the path from B to A. However, inside one AS which is shared by both paths, traffic from A to B will always use the same path as traffic from B to A.

true □ false □ Your home "router" which you receive from your ISP normally only has a single default route rather than a route for every single BGP prefix in the Internet.

true □ false □ A router with empty forwarding table and no route-maps receives the following two BGP updates:

(1) for 10.0.0.0/8 from neighbor AS 10: local pref 100, AS path [10 50]
(2) for 10.0.0.0/24 from neighbor AS 20: local pref 100, AS path [20 70 39 50]

Afterwards, traffic towards 10.0.0.35 will be forwarded to neighbor AS 10.

b) BGP Attributes (13 Points)

leave blank

(i) Consider the AS in Figure 8. It has three border routers (A, B and C) and two internal routers (D, E). The routers are connected through an iBGP full-mesh (not shown in the figure). OSPF is used internally with the given link weights (one is missing). Each border router receives multiple BGP advertisements (belonging to four prefixes $P1 \ldots P4$) from its eBGP neighbors with the depicted attributes (AS path and *unknown* local-pref).

Below the figure you find the forwarding tables after the routers processed all advertisements. Each table shows the selected egress router (A, B or C) for the four prefixes.

Find working local preference values for each BGP advertisement as well as one possible OSPF link weight for the link between router A and D such that you end up with the given forwarding tables. Write your solutions directly into Figure 8. (9 Points)

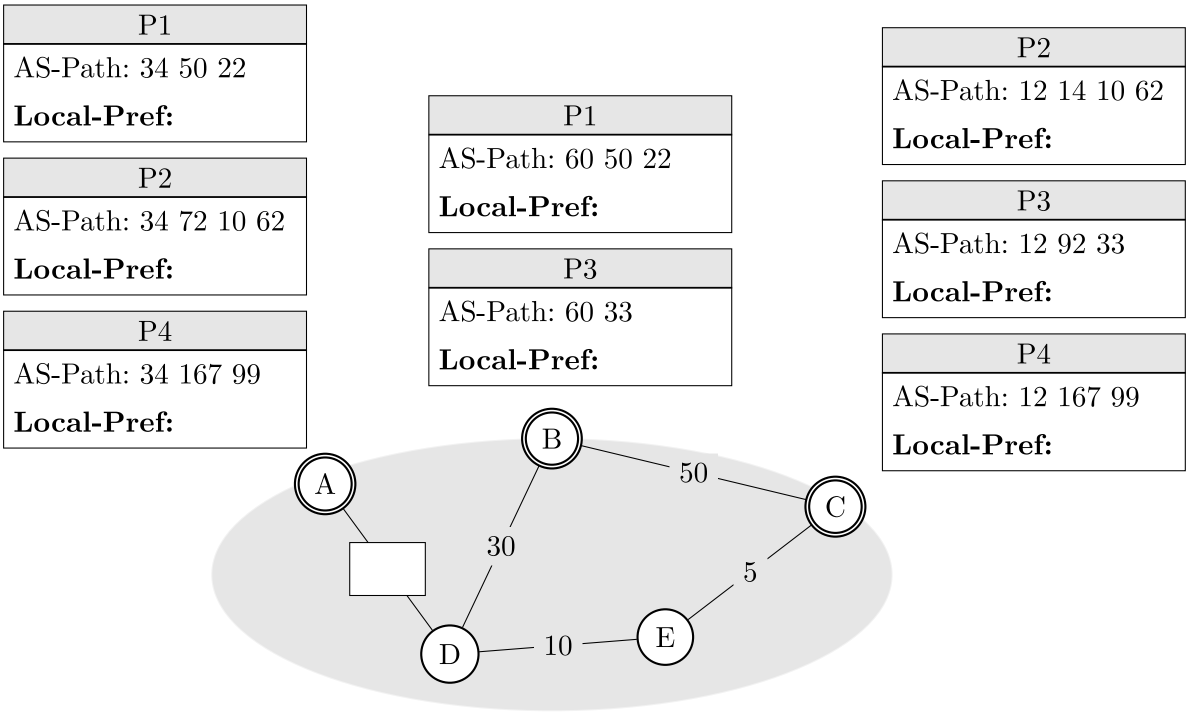

P1

AS-Path: 34 50 22

Local-Pref:

P2

AS-Path: 34 72 10 62

Local-Pref:

P4

AS-Path: 34 167 99

Local-Pref:

P1

AS-Path: 60 50 22

Local-Pref:

P3

AS-Path: 60 33

Local-Pref:

P2

AS-Path: 12 14 10 62

Local-Pref:

P3

AS-Path: 12 92 33

Local-Pref:

P4

AS-Path: 12 167 99

Local-Pref:

Figure 8: A simple BGP network forming an iBGP full-mesh.

A

Prefix	Egress
P1	A
P2	A
P3	C
P4	A

B

Prefix	Egress
P1	B
P2	A
P3	C
P4	C

C

Prefix	Egress
P1	A
P2	A
P3	C
P4	C

D

Prefix	Egress
P1	A
P2	A
P3	C
P4	C

E

Prefix	Egress
P1	A
P2	A
P3	C
P4	C

(ii) The MED attribute can be used to control the inbound traffic. Explain in which scenarios the MED is applicable and mention one disadvantage of MED. (2 Points)

Scenario: ____________________________________

Disadvantage: ____________________________________

(iii) Explain one other technique (which does *not* use the MED attribute) an AS can use to control over which border router it receives inbound traffic. Are there any disadvantages of your technique? (2 Points)

Technique: ____________________________________

Disadvantage: ____________________________________

c) Path Analysis (10 Points)

leave blank

Consider the Internet topology with 8 nodes shown in Figure 10. The topolgy consists of one IXP (D) and 7 other ASes. Each AS consists of a single BGP router and uses the default selection and exportation BGP policies based on their customers, peers and providers.

Remember that we use community values to tell the IXP to which peers it should forward a given announcement. A community value of D:X would mean that IXP D should announce the corresponding route to peer X.

(i) You are given the full routing table of the IXP in Figure 9. The table shows the announced prefix with corresponding AS path and BGP community value. From the information given in the routing table, identify all ASes which are directly connected and draw corresponding links into Figure 10. (4 Points)

D

Prefix	AS path	BGP Community
C	C	D:B D:E
F	E F	D:B D:C
H	E F H	D:C
G	C G	D:B D:E
E	E	D:B D:C
A	B A	D:C D:E

Figure 9: *Full* routing table of IXP D.

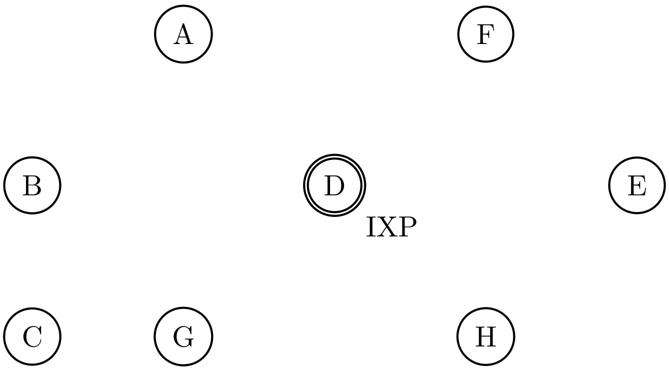

Figure 10: Small Internet topology with 8 nodes. Node D is an IXP.
Draw all links which you can infer from the routing table above.

(ii) Assume now that every node in Figure 10 advertises its own prefix. From the routing table given in Figure 9 identify *two* cases of traffic engineering. I.e., prefixes which are not forwarded equally or not reachable from certain nodes. For each case, explain how you detected it. (3 Points)

Case 1:

Case 2:

(iii) Note that this task is unrelated to the network shown in Figure 10. Assume that you receive a route for a prefix P with AS path ABC. You then perform a traceroute towards a destination in P. Looking at the traceroute output you realize that your traffic took the AS path $ABXC$. Explain two reasons why this could happen, i.e. the actual forwarding path does not match with the announced path in the BGP route. (3 Points)

Reason 1: ____________________

Reason 2: ____________________

d) BGP Hijacks (10 Points)

leave blank

Consider the Internet topology consisting of 9 Autonomous Systems (ASes) in Figure 11. Single-headed plain arrows point from providers to their customers (AS C is the provider of AS H), while double-headed dashed arrows connect peers (AS B and AS C 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.

AS H is the origin of prefix 45.7.0.0/20 and advertises it to its neighbors. Independently of what the external advertisements are, AS H ***always*** prefers its internal route to reach any IP destination in 45.7.0.0/20.

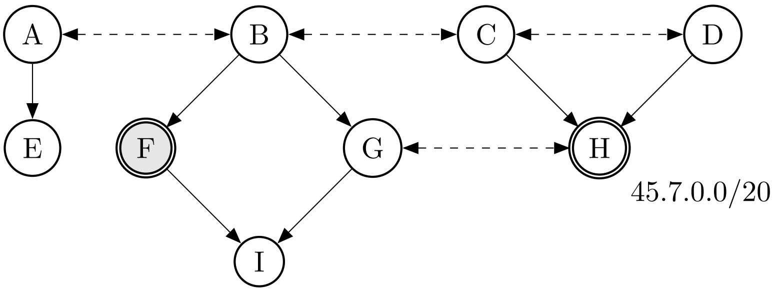

Figure 11: An Internet topology with 9 ASes. AS F aims at hijacking traffic destined to AS H.

(i) AS F performs a more-specific hijack attack, e.g. with the prefix 45.7.0.0/24. Under this attack scenario, would AS D send its traffic to the correct origin (AS H) or would it prefer the hijacker (AS F)? If it sends the traffic to H explain why the hijacked prefix did not reach AS D. Otherwise explain which path the traffic would take from D to the hijacker F. (2 Points)

(ii) AS F now performs a same-prefix attack and advertises 45.7.0.0/20 to *all* its neighbors. From which ASes is AS F able to attract traffic? (2 Points)

(iii) AS F realizes that the attack is very easy to detect as it drops all the affected traffic. To improve on that, design an attack with the following two properties:

- AS F is able to attract *all* the traffic towards 45.7.0.0/20 from *at least four* ASes;
- AS F has a valid return path available over which it can forward all the attracted traffic back to AS H.

Try to solve this question in the following three steps. (6 Points)

First, decide which neighbor you want to use for your return path (AS B or AS I). You can assume that AS F would use a static route to forward all attracted traffic to this neighbor. List all the ASes that belong to your return path:

Second, try to attract *all* traffic towards 45.7.0.0/20 from at least four ASes. Keep in mind that these four ASes should not belong to your return path which you selected in the previous step. Indicate which prefix(es) AS F advertises to which neighbor in order to achieve that:

Finally, make sure that the ASes on your return path do not get influenced by the advertisements from the previous step. I.e., they should still prefer the route from AS H. Describe how you have to modify (hint: BGP attributes) the advertisements from the previous step to achieve that:

○ leave blank

Task 4: Reliable transport 36 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	
□	□	In a network with *no* packet loss, the UDP protocol would provide all features of a reliable transport protocol.
□	□	In a reliable transport protocol *without* cumulative ACKs, the sender receives at least one ACK for each transmitted data packet.
□	□	Host A sends a single TCP flow to destination D while host B sends three TCP flows. All *four* flows share a link. The TCP congestion control algorithm ensures that the available bandwidth is equally distributed between host A and B.
□	□	A host sends data to a server S using a TCP connection with source port 33333. Later the host once again wants to send data to S. It needs to use the same source port (33333) as otherwise S cannot know that the data belongs to the same host.
□	□	Assume that the TCP congestion control algorithm would divide the current CWND by four instead of dividing it by two after receiving duplicated ACKs. This would still result in an algorithm which provides fairness.

b) Congestion Control (15 Points)

□ leave blank

(i) The congestion control algorithm discussed in the lecture detects congestion based on packet loss (i.e., duplicated ACKs or timeouts). However, loss is not the only signal which we could use to detect congestion. Describe one other possible signal which a congestion control algorithm could use to detect that a link/the network is congested. How does your signal change under congestion? (3 Points)

Signal: ______________________________

____

____

Reaction to congestion: ______________________________

____

____

(ii) Besides the congestion window (CWND), TCP also uses a receiving window (RWND). What is the purpose of the receiving window? How does TCP react if the CWND is bigger than the RWND? (2 Points)

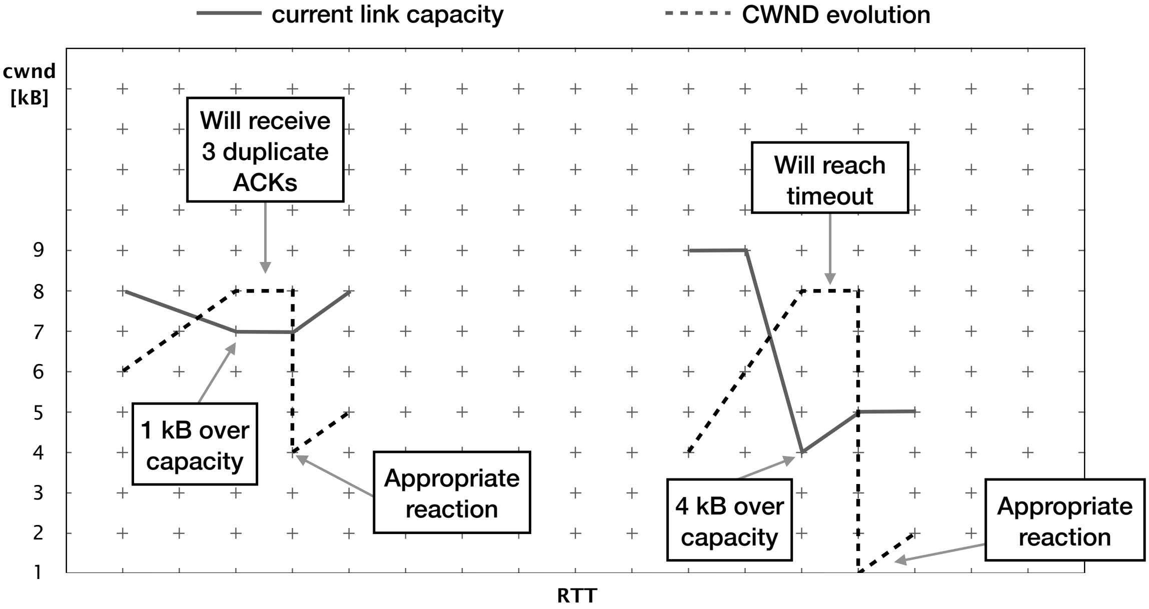

Figure 12: Reaction of the CWND (dashed line) if the current link capacity (continuous line) is exceed by at most 2 kB (left) or by more than 2 kB (right).

(iii) 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 12 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 12 right).

Draw the CWND evolution in Figure 13 in reaction to the link capacity indicated with the continuous line. We observe an ongoing flow and the first step of the CWND is indicated with the dashed line on the left side of the figure. Continue from RTT 2 (CWND at 10 kB) and assume that we are currently *not* in a slow-start phase. You can stop once you reach RTT 22. (10 Points)

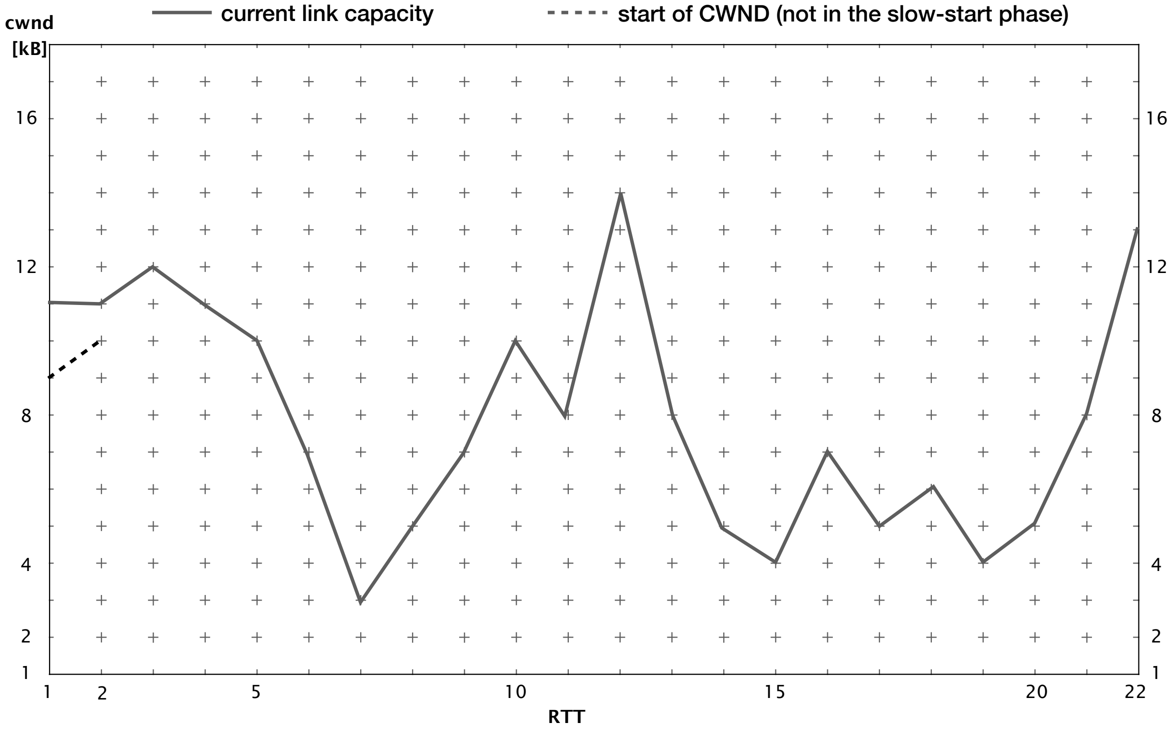

Figure 13: Complete the CWND evolution.

If you made a mistake, you can use the extra copy of Figure 13 at the end of the exam. **Important**: Please indicate your final version clearly.

c) **GBN** **(16 Points)**

leave blank

In this question we consider a GBN protocol which uses Selective Repeat. Figure 14 (on the next page) shows an ongoing transmission between a sender and a receiver. Here is a *non-exhaustive* list of the choices we made about the protocol implementation details:

- The sender saves all the transmitted but so far not acknowledged data segments in a sender buffer which can contain up to 4 segments;
- The receiver answers with an ACK for *every* received data segment. The ACK number always points to the next expected in-order data segment and cumulatively acknowledges all the previous data segments. Out-of-order segments are saved in a buffer and delivered (removed from the buffer) as soon as the missing packets arrive;
- Lost data segments are detected based on duplicate ACKs received by the sender. If the sender receives 3 duplicate ACKs, the potentially missing segment is immediately retransmitted;
- After a timeout is reached (indicated in the figure), the sender retransmits all data segments which are currently in its sender buffer. The retransmission happens in order (considering the sequence number), one packet per step;
- The used sequence number is exchanged using a header field which contains 5 bits. That means we can have sequence numbers from 0 up to 31. Afterwards, the sequence number wraps around and starts again at 0. For example: . . . , 30, 31, 0, 1, . . .

(i) Fill in the missing parts for the sender (sender buffer and transmitted data segment) and receiver (out-of-order buffer and transmitted ACK) directly in Figure 14. We observe an ongoing connection using the previously described GBN protocol. Note that some fields are already filled in to help you. For the ACK segments, use the following syntax: AN acknowledges the correct reception of all the data segments up to and including $N-1$. If there are no segments in the sender buffer or the out-of-order buffer, indicate it clearly by crossing the corresponding buffer. (12 Points)

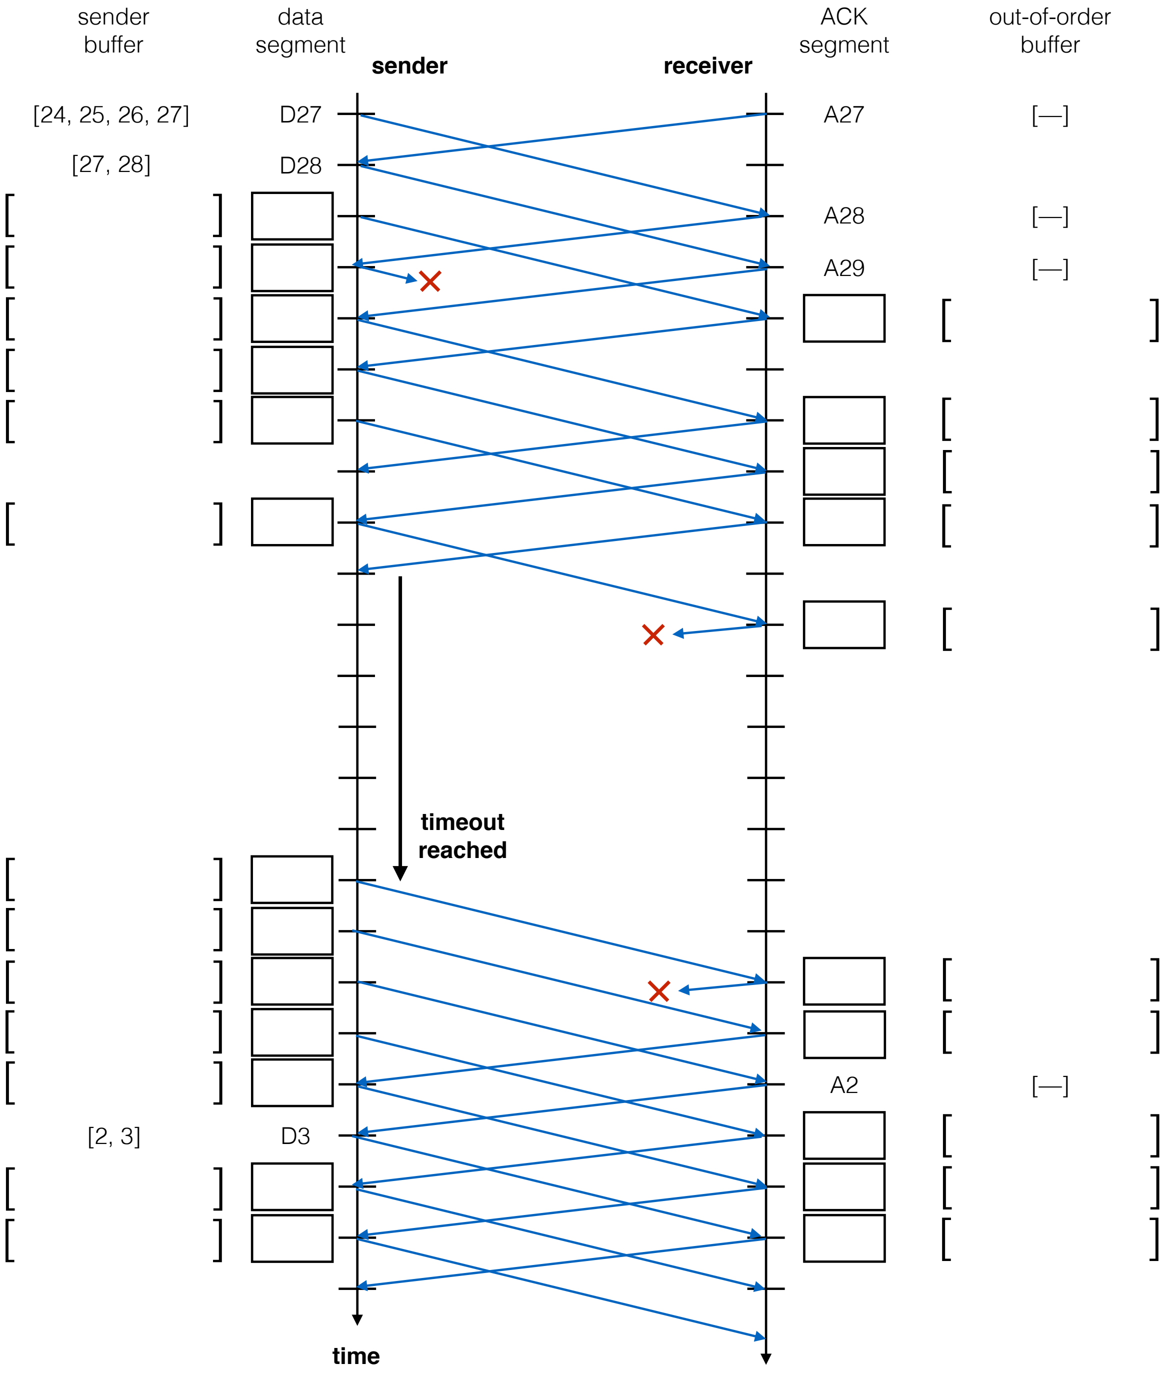

Figure 14: Incomplete time-sequence diagram of an ongoing connection using the previously described GBN protocol.

If you made a mistake, you can use the extra copy of Figure 14 at the end of the exam. **Important**: Please indicate your final version clearly.

(ii) As we have seen in the GBN implementation described before, the sequence number wraps around at some point. That means the sender and the receiver will eventually reuse the same sequence numbers again. Explain why this is not a problem in practice and how sender and receiver know which sequence numbers are currently expected. Note that this question is not connected to Figure 14 and can be solved independently.

(2 Points)

(iii) To further reduce the amount of unnecessary retransmissions we can use Selective Acknowledgement (SACK). A SACK header often contains a normal ACK number which points to the next expected in-order segment. In addition, the header also contains blocks of already received out-of-order segments. Explain why the sender can *not* immediately remove the already received out-of-order packets from its sender window therefore allowing to transmit new data. (2 Points)

○ leave blank

Task 5: Applications 21 Points

□ leave blank

a) Warm-up (9 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 □ DNS is a centralized service that translates human-readable names into IP addresses.

true □ false □ www.www.nsg.ethz.ch is a valid DNS name.

true □ false □ All top-level-domain DNS servers for .ch know the IP address of www.ethz.ch.

true □ false □ A DNS root server will receive more DNS requests if it stops caching name-to-IP mappings.

true □ false □ Two users, A and B, accessing the same URL at exactly the same time, might see different web pages as a result.

true □ false □ You open a web browser and type in a URL. Your computer will necessarily generate at least one DNS request.

true □ false □ Persistent HTTP improves throughput over vanilla HTTP by using more connections.

true □ false □ Using BGP Anycast as load-balancing technique, one can ensure HTTP requests will go to the least loaded server.

true □ false □ The Mail User Agent relies on DNS to learn the IP address of the Mail Transmission Agent to send mails to.

b) One ETH domain, two DNS configurations (12 Points)

leave blank

While preparing for the comm-net exam, you decide to take a look at the respective DNS setups of ETHZ and EPFL and notice some interesting difference:

ETH's DNS setup:

```
;; QUESTION SECTION:
;www.ethz.ch.   IN   A

;; ANSWER SECTION:
www.ethz.ch.   48   IN   A   129.132.19.216
```

EPFL's DNS setup:

```
;; QUESTION SECTION:
;www.epfl.ch.   IN   A

;; ANSWER SECTION:
www.epfl.ch.   42838   IN   CNAME   www.epfl.ch.cdn.cloudflare.net.
www.epfl.ch.cdn.cloudflare.net.   76 IN   A   104.20.228.42
www.epfl.ch.cdn.cloudflare.net.   76 IN   A   104.20.229.42
www.epfl.ch.cdn.cloudflare.net.   76 IN   A   172.67.2.106
```

(i) You first consider EPFL's DNS setup and notice that they are using Cloudflare as Content Delivery Network (CDN). Briefly explain what a CDN is along with two reasons ***other than load-balancing*** explaining why EPFL might be using one. (3 Points)

What is a CDN?

First reason to use a CDN:

Second reason to use a CDN:

(ii) Why is EPFL using a CNAME record? Explain briefly. (1 Point)

(iii) How does EPFL load balance its Web requests? What's the advantage of load-balancing requests that way? (2 Points)

How does EPFL load-balance?

Why is it interesting to do so?

(iv) You then consider the ETHZ DNS setup. While ETHZ does *not* use a CDN to host its website, you know for a fact that there is more than one server hosting the website. Explain two distinct techniques ETHZ might be using to load balance incoming web requests onto multiple servers despite listing a single IP address in the DNS. Make sure to explain how each technique works (simply mentioning the technique is *not* enough). For each technique, also mention one advantage/disadvantage.

(6 Points)

Technique 1:

Advantage of Technique 1:

Disadvantage of Technique 1:

Technique 2:

Advantage of Technique 2:

Disadvantage of Technique 2:

Extra copy of Figure 13

In case you made a mistake you can use this copy of the CWND evolution for question *4 b) (iii)*. Clearly indicate which figure contains your final solution.

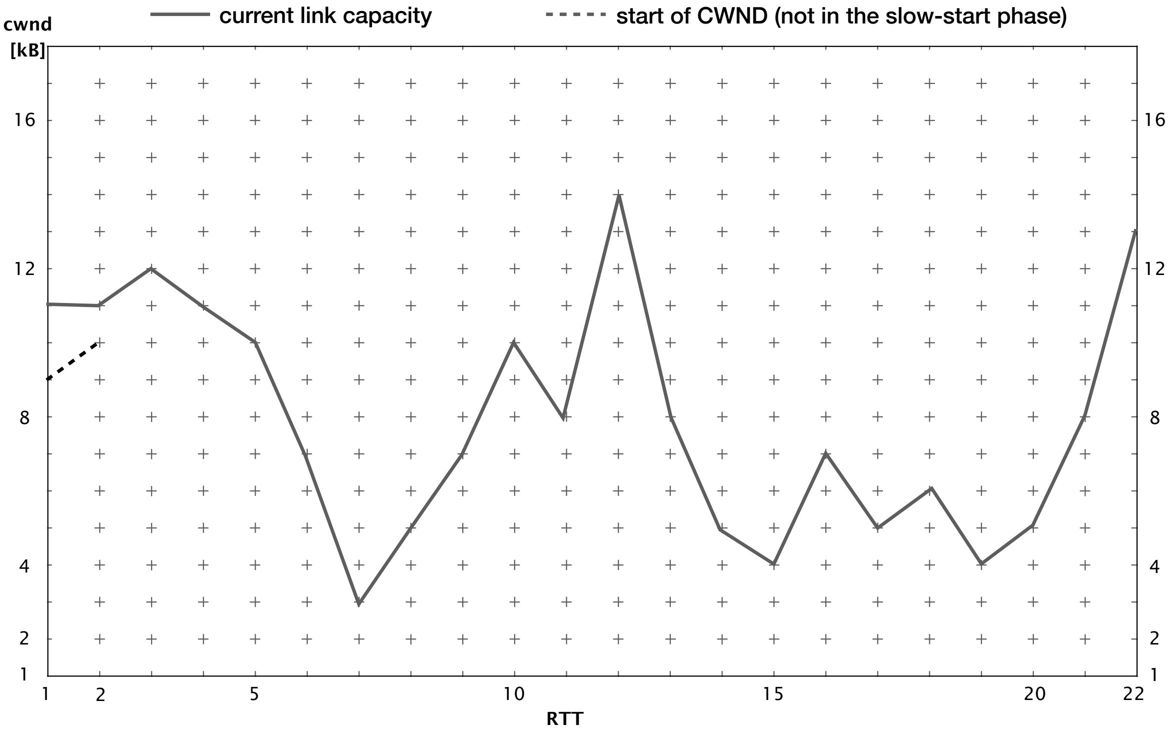

Figure 15: Copy of the CWND evolution in task *4 b) (iii)*.

Extra copy of Figure 14

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.

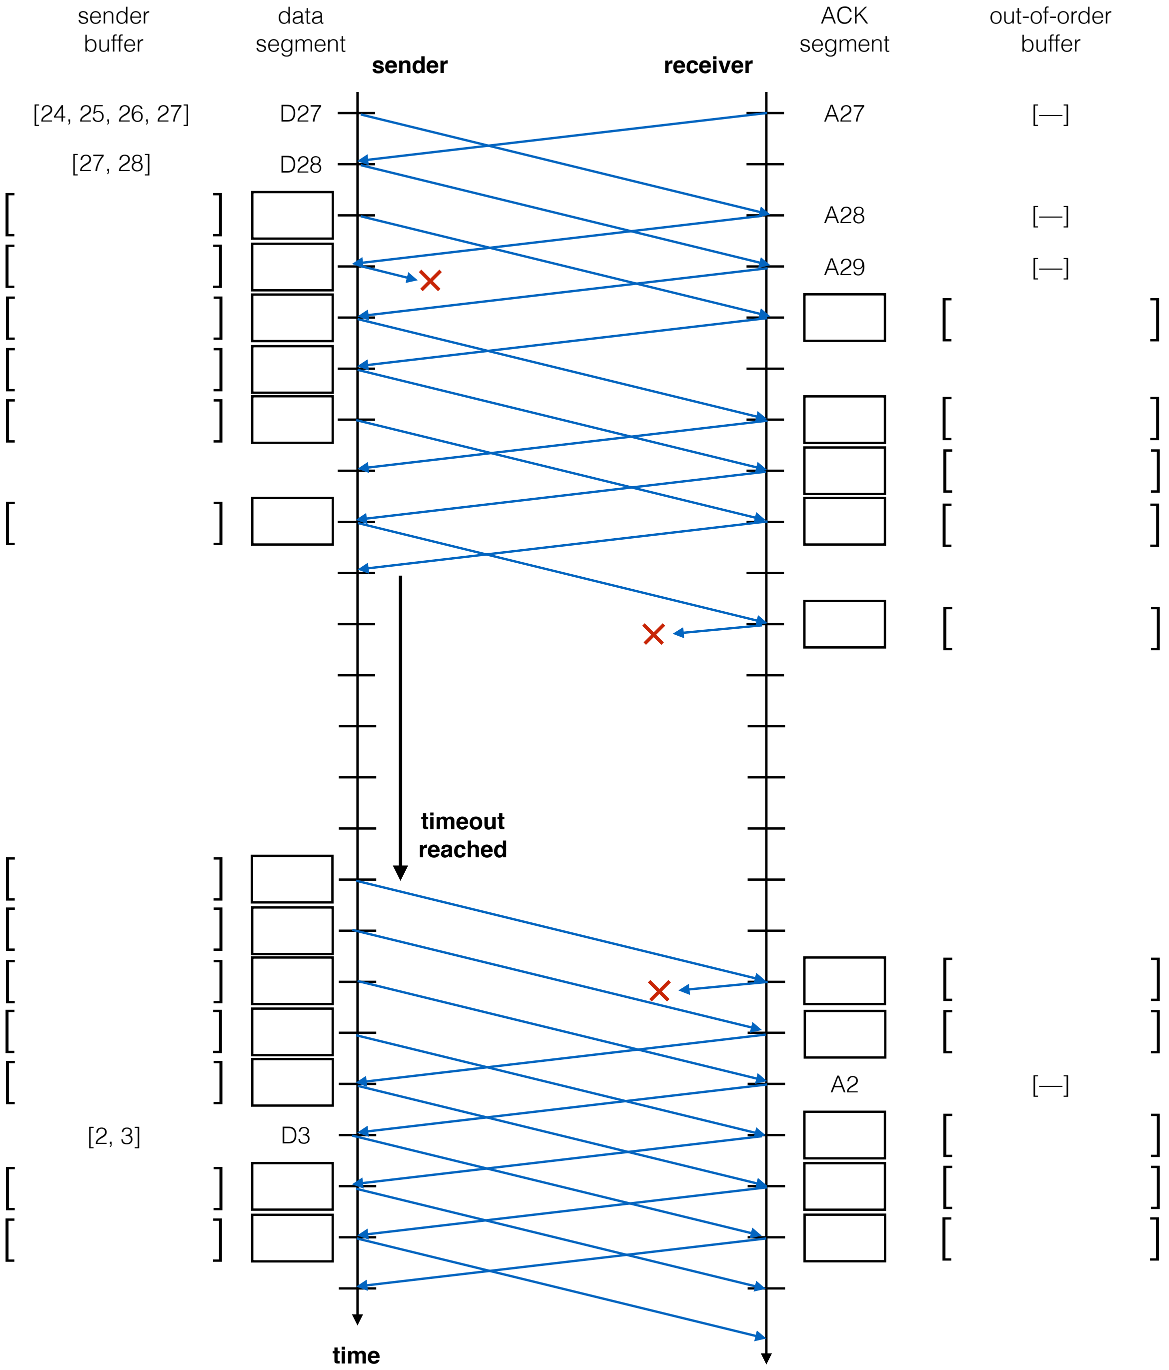

Figure 16: Copy of the GBN diagram in task *4 c) (i)*.

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: ______________