

Communication Networks

Prof. Laurent Vanbever

Solution: Exercise 99 – Additional Practice

Link State and Distance Vector

99.1 Warm-up Questions (Exam Question 2016)

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

- a) Consider a positively weighted graph G . Applying the Bellman-Ford (used by distance-vector protocols) or Dijkstra (used by link-state protocols) algorithm on G would lead to the same forwarding state.

Solution: False. In general, both solve the shortest-path problem. However, whenever multiple shortest paths exist, the two algorithms can decide differently based on their tie-breaking criteria.

- b) Link-state protocols (such as OSPF) are guaranteed to compute loop-free forwarding state as long as the link-state databases are consistent on all routers.

Solution: True. However, they can experience transient loops while it isn't the case.

- c) Link-state protocols (such as OSPF) require routers to maintain less state than distance-vector protocols (such as RIP).

Solution: False. Link-state protocols require routers to maintain the entire topology in memory (Link-State database). Distance-vector protocols only need to maintain the costs to reach each prefix.

d) Poisoned reverse solves the problem of count-to-infinity.

Solution: False. The problem is still there it is mitigated by having a small infinity value.

e) Consider a positively weighted graph G . Multiplying all link weights by 2 would change the all-pairs shortest paths computed by the Dijkstra algorithm on G .


Solution: False. Multiplying by a constant factor keeps the ranking between the paths constant.

f) Consider a positively weighted graph G . Adding 1 to all link weights would change the all-pairs shortest paths computed by the Dijkstra algorithm on G .


Solution: True. Longer paths will see a bigger increase than shorter ones.

99.2 A very long downtime (Exam Question 2019)

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



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



- a) Explain why the downtimes experienced by the three prefixes differ.

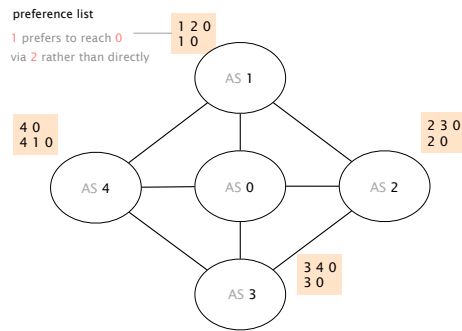
Solution: The router updates the entry in its forwarding one by one, and since there are more than 500k entries to update, it takes quite some time. Hence, it takes for each entry a different time until it is updated.

- b) Describe two solutions to speed up the convergence for these three prefixes.

Solution: One option is to prioritize the three prefixes by updating the data plane for these three prefixes before all the others. Another option is to use a hierarchical forwarding table to speed up the overall convergence process. To this end, one would use two tables: a first table that maps all of the 500k prefixes to a much smaller number of virtual next-hops and then a second table that maps the virtual next-hops to the actual output port. In this case, one does not have to update all the entries for the 500k prefixes, but only change the output port of a few virtual next-hops.

BGP Routing and Security

99.3 Convergence



Does this network ever converge?

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


Considering that only AS 0 originates prefixes, does that BGP network have a unique, stable solution?

- If yes, indicate the path that each AS selects in the stable solution.
- If not, describe an example of oscillation. For instance, by describing a sequence of messages that repeats itself.

Solution: This BGP network does have a unique, stable solution in which:

- AS 1 selects [1,2,0] (preferred path);
- AS 2 selects [2, 0];
- AS 3 selects [3, 4, 0] (preferred path);
- AS 4 selects [4, 0] (preferred path).

99.4 BGP Security (Exam Question 2020)



An Internet topology of 9 ASes in which AS I announces a prefix and AS G tries to hijack it.

Consider the Internet topology consisting of 9 Autonomous Systems (ASes) in the Figure above. 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 *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 (in alphabetical order; e.g., a route from AS *A* is preferred over AS *B* in case of a tie).

AS *I* is the origin of prefix 33.33.0.0/16 and advertises it to its neighbors. Independently of what the external advertisements are, AS *I* **always** prefers its internal route to reach any IP destination in 33.33.0.0/16.

- a) AS *G* wants to hijack the traffic going to AS *I* for 33.33.0.0/16. It starts advertising the exact same prefix with itself, AS *G*, as origin. From which ASes is it able to hijack the traffic?

Solution: We can hijack traffic from the following ASes: *A, B, D, E, H*.

- b) The ASes notice the hijack and, as a counter-measure, deploy Resource Public Key Infrastructure (RPKI) Internet-wide. After that, from which ASes is the attacker able to hijack the traffic by still advertising the exact same prefix with itself as origin?

Solution: AS *G* is not able to attract any traffic anymore as all the ASes can tell that the route is invalid.

- c) RPKI has a flaw. What is the problem of RPKI? How can AS *G* hijack the prefix 33.33.0.0/16 despite RPKI? From which ASes is AS *G* able to hijack the traffic?

Solution: RPKI only checks whether the origin AS indeed owns the advertised prefix, but it has no way of validating that the announcement actually originated at the correct AS and has not just been added to the beginning of the AS path.

AS *G* can simply add a fake origin to its announcements and advertise 33.33.0.0/16 with an AS path of *G, I*.

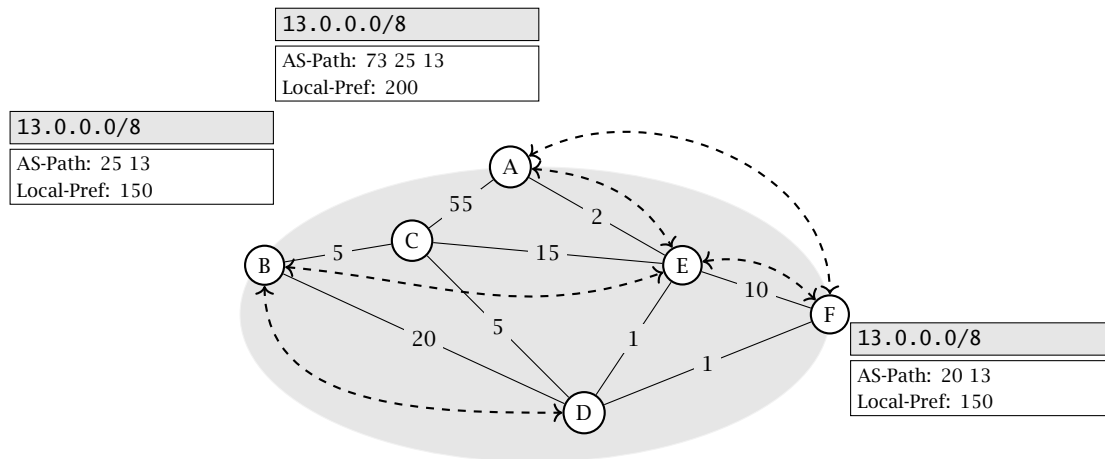
By doing so, it can only attract the traffic from ASes *A, D*. Since AS *G* prepends the true origin to the AS path, the AS path length of its announcements increases by one. Hence, some ASes prefer the true announcement over the hijack because of that (e.g., AS *B*).

- d) In response, the ASes switch to BGPsec (Secure BGP). Explain what security it provides and how AS *E* can detect that the announcement from AS *G* has a forged AS path.

Solution: BGPsec provides origin and path authentication through cryptographic signatures. It allows to verify the entire path.

In BGPsec, every AS signs the message with its private key and includes the AS number of the next AS on the path. When a router receives an announcement, it can check each signature on the path using the public keys of the ASes. It then quickly sees whether the path is valid or not. In this case, AS *E* could detect that the origin (i.e., the very first hop on the AS path) is not correctly signed.

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



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

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

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

Solution:

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

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

Solution:

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


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

Solution:

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

99.6 Traffic Engineering

Assume that ETH has only one prefix: 82.130.64.0/21. As depicted on the left, the ETH network is connected to three providers (Swisscom, Deutsche Telekom and Switch) and the providers are interconnected with each other. The contract with Swisscom is the cheapest one (indicated by the dollar symbols). For this reason, ETH wants to receive all the incoming traffic over the Swisscom link and therefore announces its prefix only to Swisscom.



ETH is connected to three providers with different costs.

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

Solution:

Not a good solution. If the link fails, ETH will no longer receive any traffic. ETH is no longer reachable from other networks.

- b) To improve the connectivity in case of a link failure between ETH and Swisscom, ETH wants to optimize its announcements. Write down the prefixes which ETH announces to Swisscom, Deutsche Telekom and Switch. During normal operation (no link failure) ETH should still receive all incoming traffic over the Swisscom link.

Solution:

To Swisscom: 82.130.64.0/22 and 82.130.68.0/22 (other splits are also possible)

To Deutsche Telekom: 82.130.64.0/21

To Switch: 82.130.64.0/21

- c) After further investigations, ETH decides that only traffic towards 82.130.68.0/23 has to be received over the Swisscom link. All the other traffic can enter over any of the providers. Which prefixes do you have to announce to achieve this traffic distribution?

Solution:

To Swisscom: 82.130.68.0/23 and 82.130.64.0/21

To Deutsche Telekom: 82.130.64.0/21

To Switch: 82.130.64.0/21

Reliable Transport

99.7 TCP Warm-up (Exam Question 2019)

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

- a) In contrast to the GBN protocol used in the project, TCP's sequence number often increases by more than one between two consecutive data packets.

Solution: True. The TCP sequence number correlates to the current byte in the transmitted byte stream.

- b) A client having an ongoing TCP connection to a server (IP 1.2.3.4, port 80) is not able to start a second TCP connection towards 1.2.3.4:80.

Solution: False. The source port of the two flows can be different to distinguish the two connections.


- c) A TCP packet with a value of 9 in its header length field (HdrLen, sometimes also called data offset) indicates 16 bytes of TCP options.

Solution: True. HdrLen of 9 indicates 36 bytes. The default TCP header size is 20 bytes (no options) so we have 16 bytes for TCP options.

- d) Consider an ongoing TCP flow. Whenever the congestion window increases, the sender can transmit additional data segments.

Solution: False. It also depends on the sender/receiver window sizes. It could e.g., be that the receiver currently cannot handle an additional data segment.

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



e) Flow f was in the slow start phase exactly twice.

Solution: True. Once before point A and once after point B.

f) Flow f experiences at least one packet loss between time A and B.

Solution: False. The duplicated ACKs could also be due to packet reordering or delays (no packet loss required).

g) In the future, f will never be able to experience a higher CWND size than B.

Solution: False. The CWND could reach higher values in the future, for example if other ongoing flows end (more bandwidth available).

h) Consider another flow f_2 starting at exactly the same time as f and traversing the exact same path, then f_2 would experience the exact same CWND evolution.

Solution: False. For example the packets of the second flow could be dropped/delayed at different points in time which would result in a different CWND evolution.

IP, DHCP, ARP

99.8 IPv6 Calculations

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

- a) You are the operator of an enterprise network. Your ISP is giving you a /96 subnet 2001::/96. How many addresses do you have available? Is that a reasonable subnet size compared with the currently available IPv4 addresses?

Solution: You have $2^{128-96} - 1 = 4'294'967'295$ addresses available (no broadcast addresses in IPv6). This is as if you had all IPv4 addresses available exclusively for your enterprise network.

- b) In your enterprise network, each host machine is identified by a unique ID starting from 1. Now that you have a lot of IPv6 addresses available, you decide to give each host a unique IP address. The host with ID 1 gets the first IPv6 address in your subnet (2001::1), the host with ID 2 the second IP address and so on. Complete the following table:


Solution:

IPv6 address	host ID
2001::5	5
2001::E	14
2001::3A5	933
2001::FFFF:FFFF	4 294 967 295
2001::3:0	196 608

99.9 Putting everything together

Consider the network on the left composed of three Ethernet segments separated by two intermediate routers (A and B). In this network, the server's interface along with the routers' interfaces are configured with static IP addresses. While clients connected to the 3.0.0.0/24 Ethernet segment obtain an IP address via DHCP.

Assuming that the client has just started, with a perfectly empty state, precisely describe all packets that are generated when the command "ping 1.0.0.2" is issued (until the server answers back). Among others, your answer *must* include the content of the Layer 2 and Layer 3 headers.



Describe everything that happens to the packets sent between the client and the server

Solution:

- *src_mac: client_mac, dst_mac: broadcast*
src_ip: 0.0.0.0, dst_ip: 255.255.255.255
payload: DHCP discovery
- *src_mac: dhcp_server_mac, dst_mac: broadcast*
src_ip: 3.0.0.1, dst_ip: 255.255.255.255
payload: DHCP offer for 3.0.0.3
Note: In addition to the IP address of the client (arbitrarily picked by the DHCP server), the DHCP offer also contains the gateway (3.0.0.2) to use. The DHCP offer can be sent either in broadcast or unicast (both answers accepted). Broadcast is often used as some network stacks will not accept "unicasted" frames unless an IP address is configured first.
- *src_mac: client_mac, dst_mac: broadcast*
payload: ARP request: Who has 3.0.0.2? Tell 3.0.0.3.
- *src_mac: routerB_right_mac, dst_mac: client_mac*
payload: ARP reply: 3.0.0.2 is at routerB_right_mac
- *src_mac: client_mac, dst_mac: routerB_right_mac*
src_ip: 3.0.0.3, dst_ip: 1.0.0.2
payload: ICMP echo request
- ... (rest of the packets are omitted)
- *src_mac: server_mac, dst_mac: routerA_left_mac*
src_ip: 1.0.0.2, dst_ip: 3.0.0.3
payload: ICMP echo reply

99.10 Detective work

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

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

- a) Can you identify all the hosts that are part of the local network?

Solution: The local hosts are all the sources and destinations that do not have to go through the default gateway (e.g., their MAC address is not replaced by the MAC address of the router):

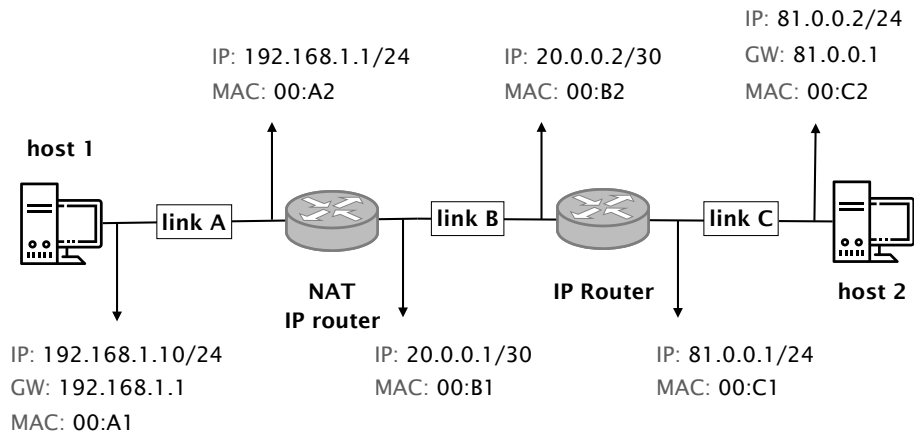
- 65.222.11.1
- 65.222.8.2
- 65.222.9.99
- 65.222.15.254
- 65.222.13.255

- b) Can you reconstruct the IP subnet used to address the hosts within that local network?

Solution: First, we should note, that the router MAC address is only used for IP sources or destinations outside the local subnet (router is used as gateway) or for packets from/towards the router. With this in mind, we can identify the lowest subnet address from the packets (65.222.11.1 -> 65.222.8.1) and (65.222.11.1 -> 65.222.8.2) as 65.222.8.1. Furthermore, we can infer that 65.222.15.254 still belongs to the local subnet (65.222.8.2 -> 65.222.15.254) but 65.222.16.1 is a destination outside of the network (65.222.11.1 -> 65.222.16.1). We can therefore identify the used subnet as 65.222.8.0/21.

99.11 Changing addresses (Exam Question 2019)

Consider the network depicted in the Figure below which is composed of two hosts along with two routers, one of which acts as Network Address Translator (NAT). Host 1 is located in a private subnet (192.168.1.0/24) and uses 192.168.1.1 as gateway, while host 2 is located in a public subnet (81.0.0.0/24) and uses 81.0.0.1 as gateway. The Figure below also depicts the MAC address of each of the 6 interfaces connected at either end of the three links. The NAT/router performs address translation between the private and the public subnets, translating traffic originating from private IPs to its public one (here, 20.0.0.1), and vice-versa.



A network topology relying on Network Address Translation.

- a) Consider that host 1 tries to open a TCP connection with host 2 on port 80 using 1337 as (random) source port. Write down a possible sequence of packet headers observed at each link for the first two packets (i.e., the SYN sent by host 1, and the SYN/ACK sent by host 2). Fill in the table below to answer. Assume that hosts and routers have the required MAC addresses in their ARP table.

Solution:

	src MAC	dst MAC	src IP	dst IP	src TCP port	dst TCP port
link A	00:A1	00:A2	192.168.1.10	81.0.0.2	1337	80
link B	00:B1	00:B2	20.0.0.1	81.0.0.2	rand	80
link C	00:C1	00:C2	20.0.0.1	81.0.0.2	rand	80
link C	00:C2	00:C1	81.0.0.2	20.0.0.1	80	rand
link B	00:B2	00:B1	81.0.0.2	20.0.0.1	80	rand
link A	00:A2	00:A1	81.0.0.2	192.168.1.10	80	1337

- b) Could host 2 initiate a TCP connection to host 1? Briefly explain why/why not.

Solution: That is not possible as host 1 is behind a NAT. Packets that reach the public IP of the NAT will not be forwarded to host 1 as there is no corresponding NAT rule available.

DNS

99.12 Curious students

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

Solution: You can simply use `dig` to issue a query for any external website you're interested in and observe the response time. If the external website has been visited recently, the response time should be close to immediate (as the local DNS server is located close by). If not, the response time will be slower as the local DNS server would have to initiate a new remote query.

You could also look at the TTL value returned by the local server and compare it to the TTL you get when querying directly the authoritative DNS server for that domain. If the TTL returned by the local server is lower than the one from the authoritative DNS server, you know that the entry has been cached by the local server and hence, someone has visited the website recently.

99.13 Multiple answers

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

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

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

Solution: (i) Resiliency: Should the first IP addresses not be reachable, you can try with the second one without having to do another DNS request.

(ii) Load of the authoritative DNS server: By giving multiple answers, you have a higher chance that one of the answers works for the entire lifetime of the reply (defined TTL value). Your DNS server will therefore get less frequent requests compared to a DNS server which only answers with one entry and a low TTL value.

(iii) Load balancing: You can use it as one way to do load balancing.

99.14 DNS basics (Exam Question 2019)

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

Note: You can assume the following:

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

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

Solution:

A	Local DNS resolver installed on the client
B	DNS server for <code>ch</code> .
C	DNS server for <code>ethz.ch</code> .
D	DNS server for <code>mail.ethz.ch</code> .
E	DNS server for <code>www.mail.ethz.ch</code> .
F	DNS server for
G	DNS server for

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

Solution:

1.	source: X, destination: A, query: <code>www.mail.ethz.ch</code> .
2.	source: A, destination: B, query: <code>www.mail.ethz.ch</code>
3.	source: B, destination: C, query: <code>www.mail.ethz.ch</code>
4.	source: C, destination: D, query: <code>www.mail.ethz.ch</code>
5.	source: D, destination: E, query: <code>www.mail.ethz.ch</code>
6.	
7.	
8.	
9.	
10.	

- c) Immediately after you determined the IPv4 address for `http://www.mail.ethz.ch/login`, you want to determine the IPv4 address for `http://www.ethz.ch`. List all DNS queries (including their source and destination) for a *recursive* resolver to determine the IPv4 address hosting `http://www.ethz.ch` *assuming that*

all the responses from above are still in the caches. You can use the letters A-G from the list above to specify the DNS servers and X to for the client. *Hint:* The correct answer requires 10 lines or less.

Solution:

1.	source: X, destination: A, query: www.ethz.ch
2.	source: A, destination: B, query: www.ethz.ch
3.	source: B, destination: C, query: www.ethz.ch
4.	
5.	
6.	
7.	
8.	
9.	
10.	

HTTP (and Web)

99.15 Loading a website (Exam Question 2017)

The website `www.your-shop.ch` consists of the following elements:

- HTML `www.your-shop.ch/index.html`
- Stylesheet `www.your-shop.ch/style.css`
- image `www.your-shop.ch/logo.png`
- image `images.your-shop.ch/product.jpg`
- Facebook like button `cdn.facebook.com/like.png`
- Facebook “tracking” code `www.facebook.com/track.js`
- Google “tracking” code `www.google.com/track.js`

- a) Assuming that your host is configured to use a local recursive DNS server in your network and all caches are empty. List all the DNS queries that your host sends to this DNS server when you open up `https://www.your-shop.ch/` in your favorite browser.

Solution: The host will send one query per domain:

- `www.your-shop.ch`
- `images.your-shop.ch`
- `cdn.facebook.com`
- `www.facebook.com`
- `www.google.com`

- b) After loading the website, you send an email to `contact@your-shop.ch` via a mail server that uses the same DNS server as your host. Does the local recursive DNS server need to run additional queries to other DNS servers if it has all the replies from the queries in the previous task in its cache? Explain why or why not.

Solution: Yes. The queries in the previous task asked for A-entries. The mail server needs an MX-entry for `your-shop.ch`, which was not queried before and is therefore not in the cache.

- c) How many TCP connections would an unoptimized browser (also referred to as “naive” in the lecture) open to load `https://www.your-shop.ch/?` Briefly explain your answer.

Solution: It will require 7 TCP connections: one connection per object.

- d) During your holidays in Australia, you realize that the Facebook like button loads much faster than the logo of the shop even though both images have the same size. Can you explain the reason for this and why you do not observe this behavior in Switzerland?

Solution: The Facebook button is hosted in a CDN and can be loaded fast from the Australian server of the CDN. The shop logo is not loaded from a CDN but from the Shop’s webserver (which could be in Switzerland).

- e) One hour later (still in Australia), you open the shop’s website again. This time, the logo of the shop and the Facebook button appear at the same time. Explain two distinct reasons that would justify this behavior.

Solution: The logo could have been moved to a CDN or it is cached by the browser.