Chair of Network Architectures and Services
School of Computation, Information and Technology
Technical University of Munich

Esolution
Place student sticker here

Note:

- During the attendance check a sticker containing a unique code will be put on this exam.
- This code contains a unique number that associates this exam with your registration number.
- This number is printed both next to the code and to the signature field in the attendance check list.

Advanced Computer Networking

Exam:	IN2097 / Retake	**Date:**	Wednesday 9th April, 2025
Examiner:	Prof. Dr.-Ing. Georg Carle	**Time:**	11:00 – 12:15

Working instructions

- This exam consists of **16 pages** with a total of **6 problems**.
  Please make sure now that you received a complete copy of the exam.
- The total amount of achievable credits in this exam is 75 credits.
- Detaching pages from the exam is prohibited.
- Allowed resources:
  - the provided cheat sheet
  - one **analog dictionary** English ↔ native language
- Subproblems marked by * can be solved without results of previous subproblems.
- **Answers are only accepted if the solution approach is documented.** Give a reason for each answer unless explicitly stated otherwise in the respective subproblem.
- Do not write with red or green colors nor use pencils.
- Physically turn off all electronic devices, put them into your bag and close the bag.

Left room from _____ to _____ / Early submission at _____

Problem 1 Quiz (10 credits)

The following questions cover multiple topics and can be solved independently of each other. **For each multiple choice question, exactly one answer is correct.**

Multiple Choice

Mark correct answers with a cross ☒
To undo a cross, completely fill out the answer option ■
To re-mark an option, use a human-readable marking ×■

a)* What is the type of the following identifier or address? `3f:2e:10:af:7f:ee:81:f0`

- ☐ MAC address
- ☐ IPv4 address
- ☒ IPv6 address
- ☐ None of the listed options
- ☐ IP identifier
- ☐ Port number

b)* Consider an IPv4 packet larger than the next link's MTU arriving at a router. The DF flag **is** set. What happens?

- ☐ The packet is fragmented
- ☐ The truncated packet is forwarded
- ☐ The packet is segmented
- ☒ The packet is dropped

c)* We calculated a delay bound using network calculus. Mark the correct answer.

- ☒ We will likely never observe this value in reality
- ☐ We will definitely observe this value in reality
- ☐ We will most likely observe this value in reality
- ☐ We will certainly never observe this value in reality

d)* Which of the following options is the definition of **replicability** according to ACM?

- ☐ *Same* team executes experiment using *different* setup
- ☐ *Same* team executes experiment using *same* setup
- ☒ *Different* team executes experiment using *different* setup
- ☐ *Different* team executes experiment using *same* setup

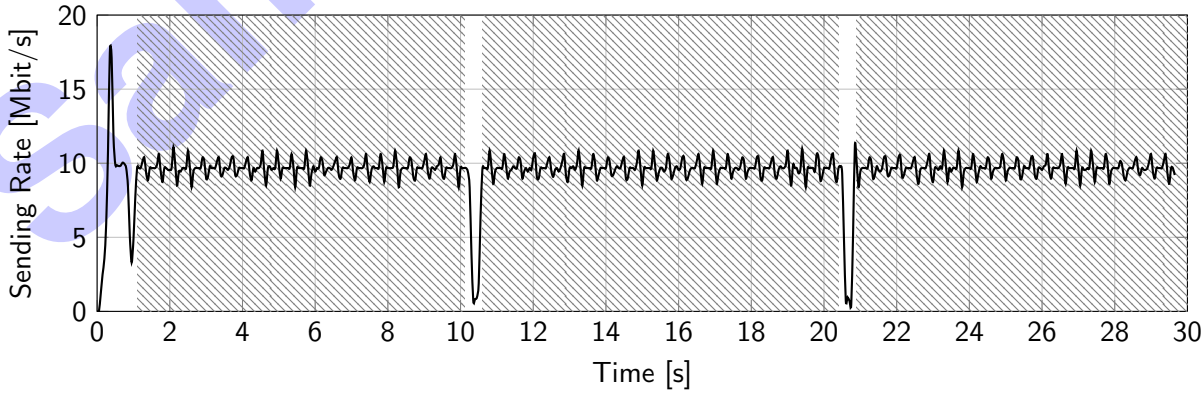

Figure 1.1: Sending rate of a BBR flow with a BBR phase marked by a grey background line pattern

e)* What phase of BBR is highlighted in Figure 1.1?

- ☐ Probe RTT
- ☐ Drain
- ☐ Startup
- ☒ Probe BW

f)* What element of P4 has a `default_action`?

- [ ] Header
- [ ] Control block
- [x] Table
- [ ] Parser

g)* What is the highest, theoretically possible VLAN ID?

- [ ] 8192
- [x] `0xFFF`
- [ ] 2^{12}
- [ ] `0x8FF`

h)* Which match type is non-mandatory according to the P4 standard?

- [x] range
- [ ] ternary
- [ ] exact
- [ ] lpm

i)* Which of these protocols uses only Layer 2 to create tunnels?

- [ ] VXLAN
- [ ] MASQUE
- [x] 801.2Q (Q-in-Q)
- [ ] IPsec

j)* What is the outcome of the STP?

- [ ] Maximum spanning tree
- [ ] Minimum spanning tree
- [x] Shortest path tree
- [ ] Merkle tree

Problem 2 Packets, Bits and Bytes (10 credits)

0 1 2

a)* List all protocols of the network stack (horizontally) for an already established HTTP/3 connection with `fail.measr.net` (see Figure 2.1) with all commonly contained **protocol headers**. Start on Layer 2 using Ethernet.

L2	L3	L4		L7	
Ethernet	IPv6	UDP	QUIC Short Header	HTTP/3 Header	HTTP/3 Payload

Consider the following IPv6 address and prefix length: `2001:4ca0:2001:10:ae1f:6bff:fe7c:896b/60`

0 1 2

b)* Assume this IPv6 address was assigned using SLAAC without privacy extension. What was the MAC address used to generate this IP address? List and explain your steps.

`2001:4ca0:2001:10:ae1f:6bff:fe7c:896b/60`

1. Extract the bytes originating from the IP address:
   `2001:4ca0:2001:10:ae1f:6bff:fe7c:896b` ⟶
2. Flip the LG bit: `0xae1f6b7c896b` ⟶ `ac:1f:6b:7c:89:6b`

0 1

c)* Mark the network and host part of the given IPv6 address.

`2001:4ca0:2001:0010:ae1f:6bff:fe7c:896b/60`

Network part | Host part

Students participating in a networking lecture were tasked with the exercise below. One of the students' answers is shown in Figure 2.2. Unfortunately, it contains several errors.

Consider the network topology displayed in Figure 2.1. Host H1 wants to connect to a webserver listening on `fail.measr.net` using HTTP/3. Assume that all caches are initially empty and that the IP address and gateway configurations were performed manually.

Fill out the given headers of the first Neighbor Solicitation sent when accessing this webserver using the provided template.

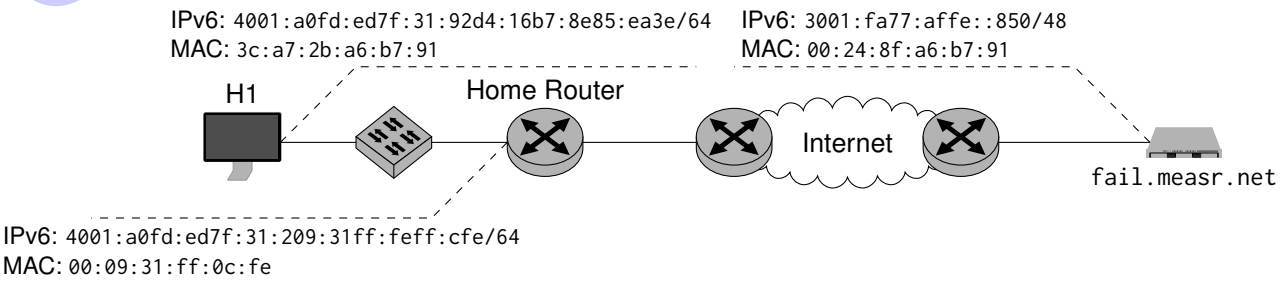

Figure 2.1: Topology of the considered question

d)* Find and correct **5** errors using the following table.[1] Swapped fields count as one error.

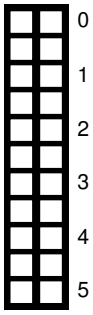

Field(s)	Description of the error	Corrected Value/Action
Source/Destination MAC Addresses	swapped	swap both fields
IPv6 header: Payload Length	Payload not 72 bytes long	$32_{(10)}$
IPv6 header: Next Header	Should not be ICMPv4 but v6	`0x3a`
IPv6 header: Hop Limit	Larger than field size allows	Any natural number $0 < x \leq 255_{(10)}$
IPv6 header: Source Address	IP should be of H1	`4001:a0fd:ed7f:31:92d4:16b7:-8e85:ea3e`
ICMPv6: Type	Should not be router solicitation but neighbor solicitation	$135_{(10)}$

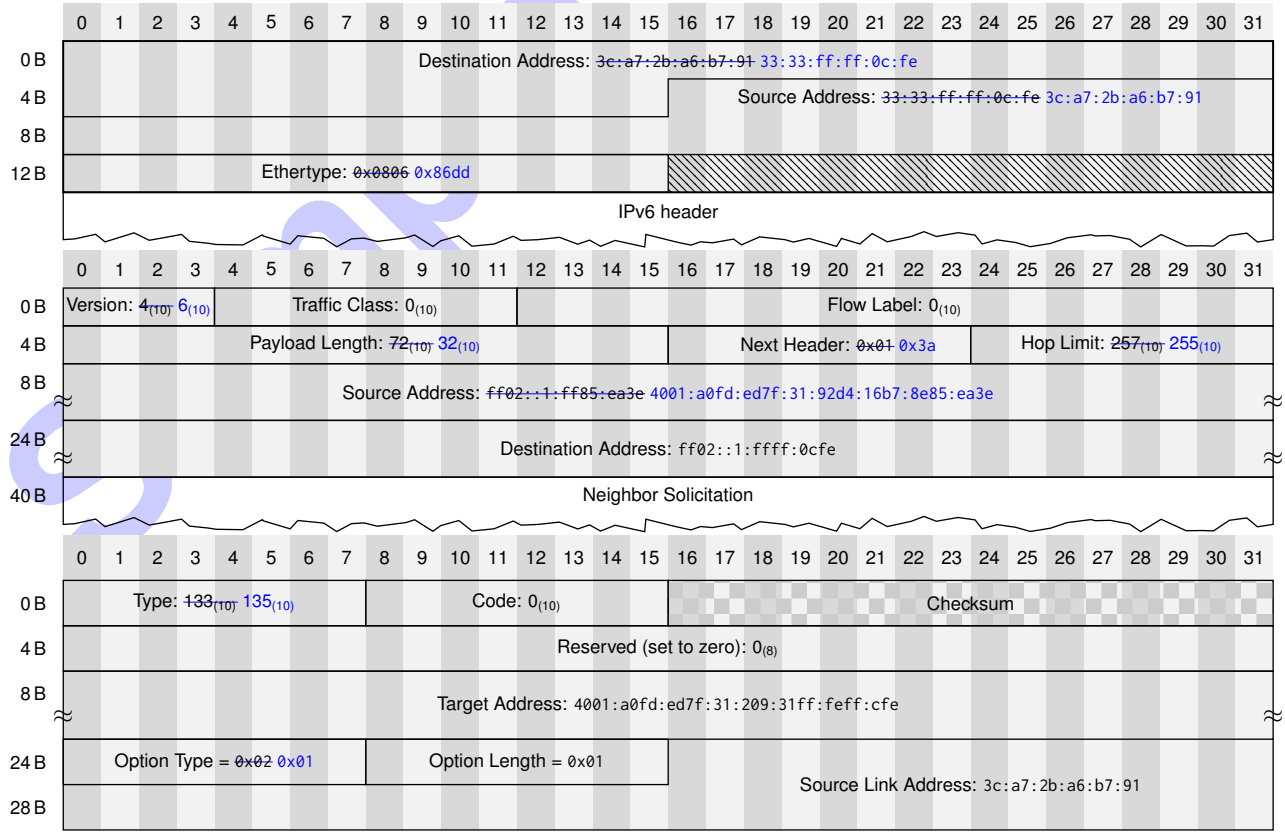

Figure 2.2: The student's answer

[1] There are **more than 5 errors** in the student's submission.

Problem 3 Traceroute (13.5 credits)

This problem investigates the tool traceroute in a software-defined network with a topology given in Figure 3.1.

Traceroute can utilize different protocols, e.g., ICMP, UDP, or TCP. When using the UDP protocol, typically, Port 443 is used as the destination port.

0 1 2

a)* Give a brief explanation why this port number is preferred over other ports, e.g., Port 5000, when using UDP for `traceroute`.

- Port 443 is the well-known port used for QUIC
- Port 5000 is not a well-known protocol, i.e., any protocol can use it
- Port 443 is often allowed even for firewalled networks whereas other protocols are not

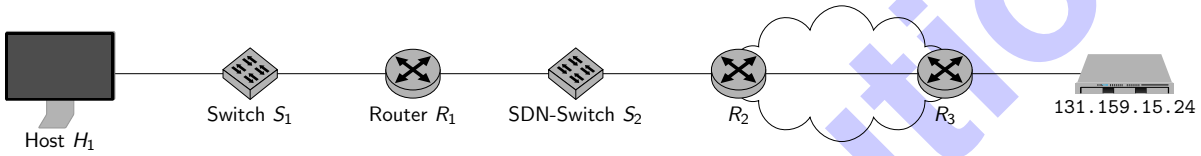

Figure 3.1: Network topology

Switch S_1 is a regular switch without SDN capability. Switch S_2 is an SDN switch, where the default rules were deleted. After that, the rules specified in Listing 1 were installed. If no rules are available for a packet, S_2 is configured to drop the packet.

```
1 ovs-ofctl add-flow S2 dl_type=0x0806,actions=flood
2 ovs-ofctl add-flow S2 dl_type=0x0800,nw_proto=0x11,actions=flood
```

Listing 1: OpenFlow rules installed on S_2

Remark: `nw_proto` specifies the entry of the protocol field on the network layer protocol. The cheatsheet contains the values for different protocols.

0 1 2

b)* Explain the rules specified in Listing 1.

Argument	Explanation
`add-flow S2`	installs rule on switch *S2*
`dl_type=0x0806`	rule is executed for packets with Ethertype 0x0806 (ARP)
`dl_type=0x0800`	rule is executed for packets with Ethertype 0x0800 (IPv4)
`nw_proto=0x11`	rule is executed for UDP (protocol type 0x11)
`actions=flood`	packets are forwarded on every switch port except the receiving one

Host H_1 successfully loaded a website located on the server with the address `131.159.15.24`. To investigate the path the packet travels through the network, the tool `traceroute` is used.

Hints for traceroute:

- IP addresses can be abbreviated by `devicename.ip`, e.g., H_1.ip (Routers can have several interfaces ip1 and ip2, use one of them consistently throughout one traceroute run)
- Response times can be chosen arbitrarily.
- Traceroute generates one query for every node and queries up to 7 nodes.
- `-P` parameter allows the specification of the protocol (e.g., `icmp`, `udp`, and `tcp`).

c)* Create a sample output for the command `traceroute -P icmp 131.159.15.24` executed on Host H_1.

0 1 2

- 0: [R_1.ip] 10ms
- 1: *
- 2: *
- ...
- 7: *

d) Describe your output for `traceroute -P icmp 131.159.15.24`, explaining the involved protocols and flow rules (see Listing 1) for requests and responses.

0 1 2

- Requests and replies use ICMP
- There is no rule for ICMP therefore ICMP cannot pass S_2
- R_1 is reached by ICMP and can answer
- No packets travel beyond R_1

e) Create an output for the command `traceroute -P udp 131.159.15.24` executed on Host H_1.

0 1 2

- 0: [R_1.ip] 10ms
- 1: *
- ...
- 7: *

f) Describe your output for `traceroute -P udp 131.159.15.24`, explaining the involved protocols and flow rules (see Listing 1) for requests and responses.

0 1 2

- Requests use UDP, replies ICMP
- There is no rule for ICMP, therefore, replies cannot pass S_2
- R_1 is reached by ICMP and can answer

g) Create an OpenFlow rule, which fixes the problem observed in Subproblem c). Do not allow more protocols than absolutely necessary to fix the problem.

0 1

```
ovs-ofctl add-flow S2 dl_type=0x0800,nw_proto=1,actions=flood
```
This rule allows ICMP messages to pass S_2

Problem 4 Internet Measurement: DNS (12 credits)

We want to perform an analysis of open resolvers and identify different types of resolvers reachable on the Internet. Nawrocki et al.[2] identified the following types of open resolvers:

Recursive Resolver Resolver that accepts recursive DNS requests, resolves them iteratively and returns the result directly.

Recursive Forwarder They forward received DNS requests to a recursive resolver and return the result received from it to the requesting client.

Transparent Forwarder Operates like a forwarder but spoofs the source IP address to be the one of the client. The recursive resolver returns the response directly to the client.

We want to conduct a large-scale IPv4 and IPv6 scanning campaign reaching as many hosts on the Internet as reasonably possible in order to identify the types of reachable open resolvers.

0 1

a)* Which IPv4 addresses do we want to scan? Where can we get a list of those or how can we generate one on the fly?

As we want to scan all reachable IPv4 addresses, we can either use a multiplicative group with a prime larger than the IPv4 address space or use a RIB dump of an AS edge router or iterate over whole IP address space and exclude non-routeable addresses.

0 1

b)* What IPv6 addresses can we scan reasonably? Where can we get a list of those or how can we generate one on the fly?

As the IPv6 address space is extremely large, full enumeration is not possible, even when using an edge router RIB dump.
Therefore, we have to rely on IPv6 hitlists as IPv6 address source.

0 ½

c)* Name an Internet measurement tool that can be used for this large-scale DNS scan?

zmap, MassDNS or zdns

Nawrocki et al. propose to use source-specific responses from a custom authoritative name server implementation (mirroring name server). This implementation will reply with two `A` (or `AAAA`) records, one being a static sanity-check IP and the other being the request's source IP[3]. Thereby, we can determine the entity which performed the actual resolution and can infer the resolver type using this information.

Assume this mirroring name server is deployed at 138.246.253.25 and 2001:4ca0:108:42::25, which should correspond to `mirror-ns.measr.net`. Our scanners are configured to resolve for `query.mirror.measr.net`.

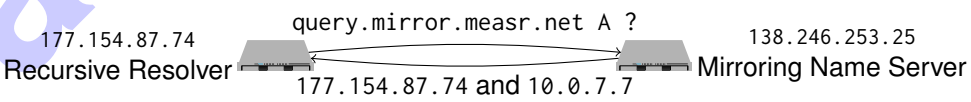

Figure 4.1: Exemplary query and answer of the mirroring name server with sanity-check IP `10.0.7.7`

0 1

d) Our scan server queried 177.154.9.9 for `query.mirror.measr.net` and got the answer 177.154.87.74 and 10.0.7.7 back from IP 177.154.9.9. What type of open resolver is 177.154.9.9?

Recursive forwarder, since the query was resolved by 177.154.87.74 and not 177.154.9.9 and the result was returned by 177.154.9.9.

[2] M. Nawrocki, M. Koch, T. C. Schmidt, and M. Wählisch. 2021. Transparent Forwarders: An Unnoticed Component of the Open DNS Infrastructure. In Proceedings of CoNEXT '21. ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3485983.3494872

[3] Example: When IP A queries the mirroring name server for a specific domain, the response contains IP A and a static IP S. See Figure 4.1 for a specific example.

e)* Give all resource records (RRs) necessary on the authoritative name server of `measr.net` for the working deployment of the mirroring name server. Use the format like the example RR given for `info.measr.net`.

0 1 2 3

Name	Type	Value
info.measr.net	TXT	Managed by TUM-I8
mirror-ns.measr.net	A	138.246.253.25
mirror-ns.measr.net	AAAA	2001:4ca0:108:42::25
mirror.measr.net	NS	mirror-ns.measr.net

We now consider the scan of the transparent forwarder using the recursive resolver (Resolver) as upstream.

f) Draw the path of resolution for `query.mirror.measr.net` starting from the scan server. Label each arrow sequentially starting with 1. Assume that all caches are empty initially.

0 1 2 3 4 5

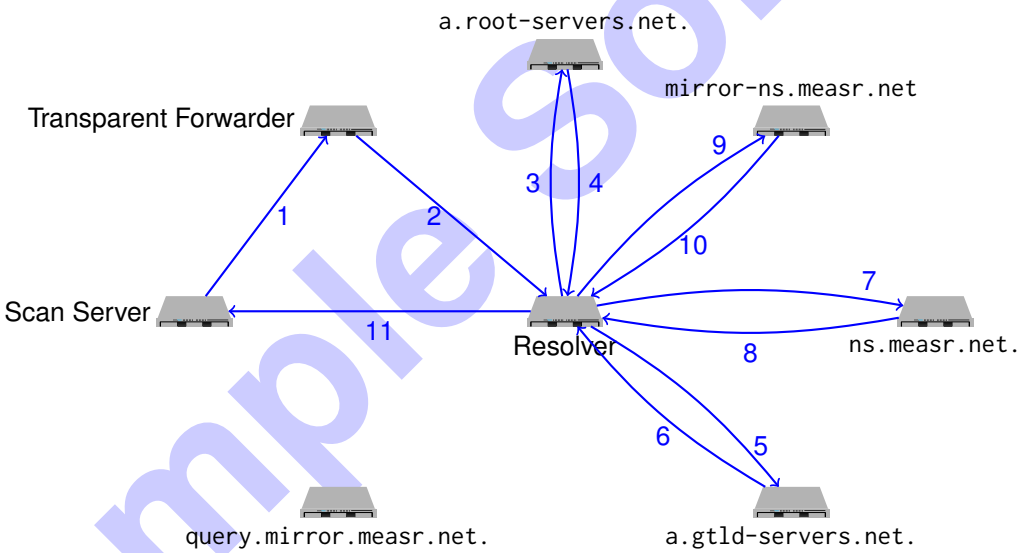

Zone	Authoritative Name Server
.	a.root-servers.net
net.	a.gtld-servers.net.
measr.net.	ns.measr.net.

Problem 5 Network Calculus (9 credits)

This problem investigates delay bounds in networks using Network Calculus.

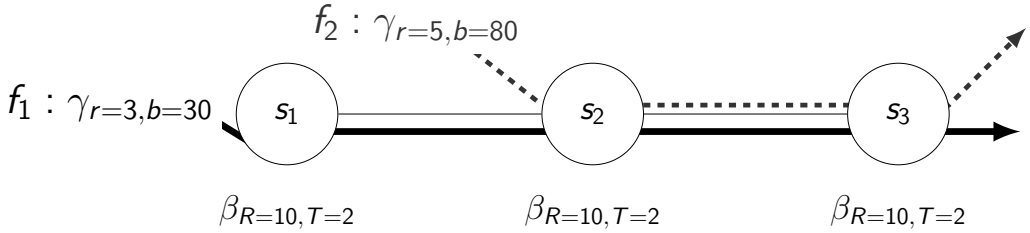

Figure 5.1: Topology with server- and flow specifications

Consider the topology in Figure 5.1. Assume each server employs strict priority queuing. Flow f_1 has the lowest priority while Flow f_2 has the highest priority. Flow f_1 traverses three servers, Flow f_2 traverses two servers.
We are interested in calculating the end-to-end delay bound of **Flow f_1** using the Separate Flow Analysis.
Hint: $\beta^{l.o.} = \beta_{R-r,\frac{b+R\cdot T}{R-r}}$ and $\alpha^\star = \gamma_{r,b+r\cdot T}$

0 1 2 3 4

a)* Perform the first step of the Separate Flow Analysis.

- Left-over service curve at s_1: $\beta_{s_1}^{l.o.1} = \beta_{10,2}$
- Left-over service curve at s_2: $\beta_{s_2}^{l.o.1} = \left[\beta_{10,2} - \gamma_{5,80}\right]^+ = \beta_{10-5,\frac{80+10\cdot 2}{10-5}} = \beta_{5,20}$
- Output arrival curve of f_2 after s_2: $\alpha^\star = \gamma_{5,80+5\cdot 2} = \gamma_{5,90}$
- Left-over service curve at s_3: $\beta_{s_3}^{l.o.1} = \left[\beta_{10,2} - \alpha^\star\right]^+ = \beta_{10-5,\frac{90+10\cdot 2}{10-5}} = \beta_{5,22}$

0 1

b) Perform the second step of the Separate Flow Analysis.

$$\beta_{e2e}^{l.o.} = \beta_{s_1}^{l.o.} \otimes \beta_{s_2}^{l.o.} \otimes \beta_{s_3}^{l.o.} = \beta_{\min(10,5,5),2+20+22} = \beta_{5,44}$$

0 1

c) Perform the third step of the Separate Flow Analysis.

$$d_{e2e} = T_{e2e} + \frac{b}{R_{e2e}} = 44 + \frac{30}{5} = 50$$

d) Assume the following changes to the scenario in Figure 5.1:

- We add a third Flow f_3 to the topology that only traverses Server s_3
- Flow f_3 is defined as a token-bucket with $r = 1$ and $b = 10$ and the same priority as Flow f_2

0 1 2

Calculate the left-over service curve for Flow f_1 at Server s_3.

- Sum of high priority traffic at s_3: $\gamma_{5,90} + \gamma_{1,10} = \gamma_{6,100}$
- Compute left-over service curve: $\beta_{s_3}^{l.o.1} = \left[\beta_{10,2} - \gamma_{6,100}\right]^+ = \beta_{10-6,\frac{100+10\cdot 2}{10-6}} = \beta_{4,30}$

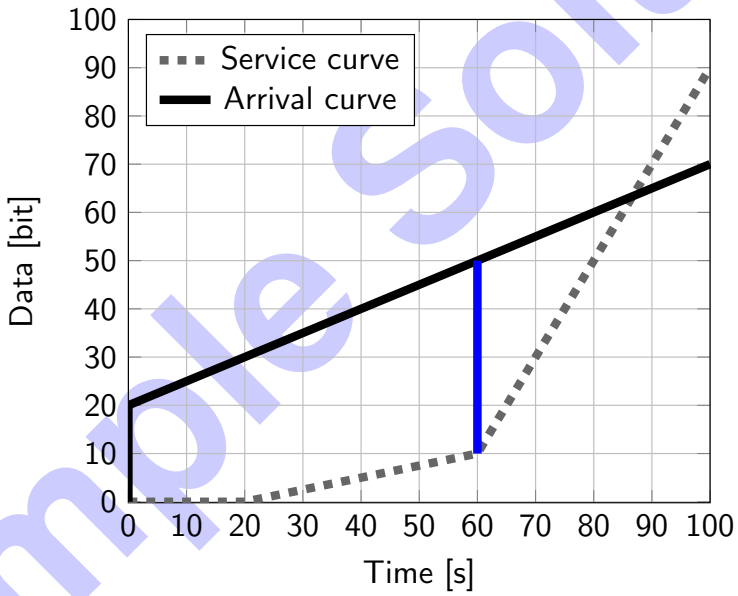

Figure 5.2: Arrival and service curve

e)* Consider the two curves in Figure 5.2. What is the **backlog** bound of a flow with this token-bucket arrival curve traversing a server with this **non-**rate-latency service curve?

0 1

Largest vertical deviation is $50bit - 10bit = 40bit$

Problem 6 ISP Troubles (20.5 credits)

In this exercise, we consider the troubles of a small regional ISP named *ExcellenceISP*. *ExcellenceISP* operates the autonomous system AS 42 and serves local home users and businesses.
ExcellenceISP wants to use an efficient routing table data structure for their routers. They decided to implement DIR-24-8 but need your help to implement the routing table in Table 6.1 as DIR-24-8 data structure.

0 1

a)* Explain what entries and indices in TBL24 (both entry variants) represent, in the DIR-24-8 data structure.

TBL24: Each entry represents a /24 IPv4 prefix with the network part of the prefix network address being the array index. Is the highest bit set, the entry refers to a TBLlong section, otherwise it corresponds to the next-hop ID.

0 1

b)* Explain what entries and indices in TBLlong represent, in the DIR-24-8 data structure.

TBLlong entry: represents the next-hop ID for a single IP address, where the index modulo 256 represents the last byte of the IP address.

Table 6.1: Routing table

№	Prefix	Network IP in Hex.	Interface	Next Hop	Next-Hop ID
①	10.7.1.4/30	0x0a070104	eth3	–	5
②	192.168.23.0/29	0xc0a81700	eth5	–	3
③	10.7.2.0/24	0x0a070200	eth1	–	4
④	192.168.22.0/24	0xc0a81600	eth4	–	0
⑤	10.7.0.0/22	0x0a070000	eth2	192.168.23.1	2
⑥	0.0.0.0/0	0x00000000	eth0	192.168.22.1	1

0 1 2 3 4 5 6

c) Fill all respective fields in the DIR-24-8 data structure, according to the routing table in Figure 6.1. Write all numbers in hexadecimal notation of correct length with leading null hexadecimal characters.

TBL24

Array Index	
0x0a06ff	0x0001
0x0a0700	0x0002
0x0a0701	0x8000
0x0a0702	0x0004
0x0a0703	0x0002
0x0a0704	0x0001
0x0a0705	0x0001
0xc0a815	0x0001
0xc0a816	0x0000
0xc0a817	0x8001
0xc0a818	0x0001

TBLlong

0x000000	0x02
0x000001	0x02
0x000002	0x02
0x000003	0x02
0x000004	0x05
0x000005	0x05
0x000006	0x05
0x000100	0x03
0x000101	0x03
0x000102	0x03
0x000103	0x03
0x000104	0x03
0x000105	0x03
0x000106	0x03

TBLlong

0x000200	
0x000201	
0x000202	
0x000203	
0x000204	
0x000205	
0x000206	

Next-Hop IDs

0	eth4, –
1	eth0, 192.168.22.1
2	eth2, 192.168.23.1
3	eth5, –
4	eth1, –
5	eth3, –
6	

Route ④ is already given.

Most IPv6 prefixes are /48 or bigger. An intern at *ExcellenceISP* proposed a DIR-48-16 algorithm, which works similar to DIR-24-8, but for IPv6 prefixes with 48 bits for the first table and 16 bit for the long table.[4]

d) Is it practical to use such a DIR-48-16 algorithm for IPv6 routing? Explain your decision using an approximation.

0 1

No, as the first table alone would be larger than $2 * 2^{48}\text{B} > 500\,000\,\text{TiB}$, which is way too large for RAM storage, which is necessary for efficient execution.

BGP

For the following subproblems, we use the AS topology in Figure 6.1. Consider the following hints regarding notation and routing policies.

- $\rightarrow$ represents a customer to provider relationship: customer $\rightarrow$ provider
- $-$ represents a peering relationship.
- All ASes apply standard routing behavior. Furthermore, the following policies are applied:
  - For routes with the same prefix, the AS selects the most cost-efficient route.
  - For routes with the same prefix and with an equal traffic cost, the shorter route is selected. If they are equally long, use the next hop with the lower AS number.

e)* Explain, what a Tier-1 provider is. Give all Tier-1 providers of the topology in Figure 6.1.

0 1 2

Providers, which only have peering relationships and no upstream provider.
Here: AS 1, AS 2 and AS3

f)* Draw the AS path of packets going from **AS 42** to **AS 89**.

0 1 2 3

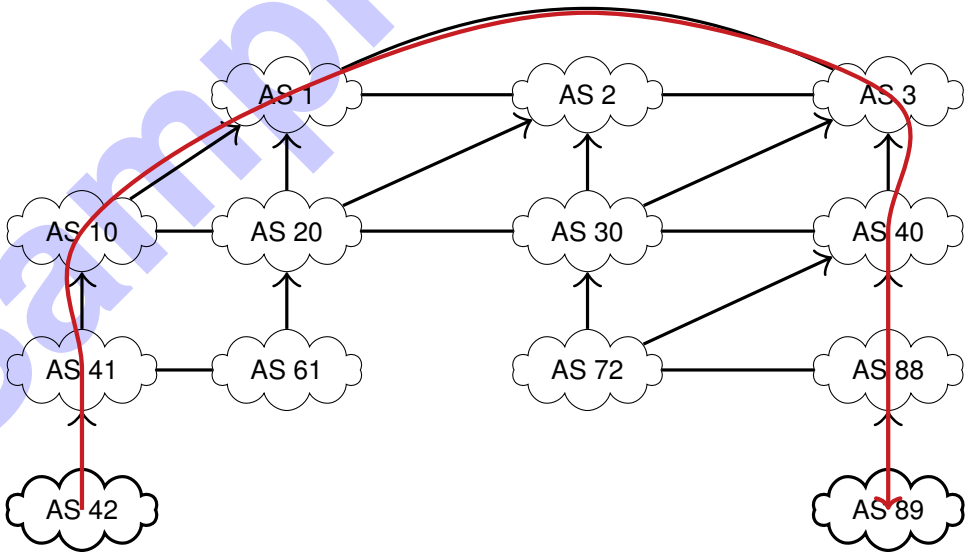

Figure 6.1: AS topology

AS 42 $\leftarrow$ AS 41 $\leftarrow$ AS 10 $\leftarrow$ AS 1 $-$ AS 3 $\rightarrow$ AS 40 $\rightarrow$ AS 88 $\rightarrow$ AS 89

[4] *ExcellenceISP* routers handle only prefixes between /64 and /0.

For the following two subproblems, use Figure 6.2 to support your explanation. Make this clear in each subproblem.

AS 42 possesses the prefix 138.246.0.0/16 and announces it. Due to a configuration error, AS 30 also announces 138.246.0.0/16. As a result of this, customers of AS 42 started complaining that they could not access most websites any more.

0 1 2 3

g)* We now consider a query from the customer 138.246.177.241 (AS 42) to **an IP address** located within **AS 89** and the potential answer. Explain in detail, how query and potentially answer are routed between ASes and why the communication fails.

The messages towards AS 89 will reach it without problem like in subproblem f). The response will instead be routed to AS 30 because of the erroneous announcement. This is the case as from AS 40, the most cost-effective and shortest path to the prefix 138.246.0.0/16 leads to AS 30 using the peering link.

0 1 2

h)* All customers of AS 42 use a recursive resolver located in AS 2. Can customers of AS 42 still access the website hosted at **the domain** `measr.net`, which is located in AS 41? Explain your choice.
Assume that they have never before visited this website.

No, they can not access `measr.net`. This is since the DNS response from the recursive resolver in AS 2 will not arrive at AS 42 but at AS 30, since the path to AS 30's announcement is shorter from AS 2.

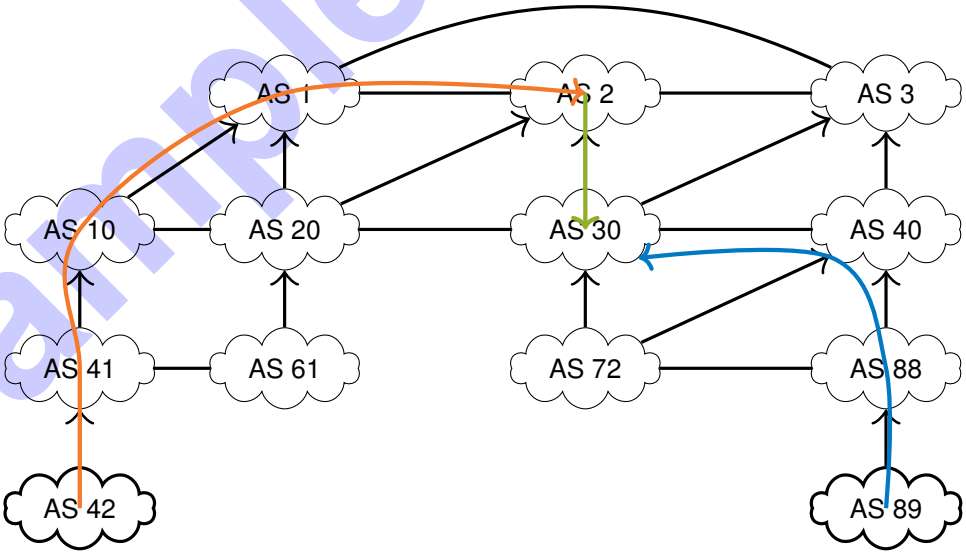

Figure 6.2: AS topology

Array Index	TBL24
0x0a06ff	0x0001
0x0a0700	0x0002
0x0a0701	0x8000
0x0a0702	0x0004
0x0a0703	0x0002
0x0a0704	0x0001
0x0a0705	0x0001
0xc0a815	0x0001
0xc0a816	0x0000
0xc0a817	0x8001
0xc0a818	0x0001

	TBLlong
0x000000	0x02
0x000001	0x02
0x000002	0x02
0x000003	0x02
0x000004	0x05
0x000005	0x05
0x000006	0x05
0x000100	0x03
0x000101	0x03
0x000102	0x03
0x000103	0x03
0x000104	0x03
0x000105	0x03
0x000106	0x03

	TBLlong
0x000200	
0x000201	
0x000202	
0x000203	
0x000204	
0x000205	
0x000206	

	Next-Hop IDs
0	eth4, –
1	eth0, 192.168.22.1
2	eth2, 192.168.23.1
3	eth5, –
4	eth1, –
5	eth3, –
6	

Route ④ is already given.

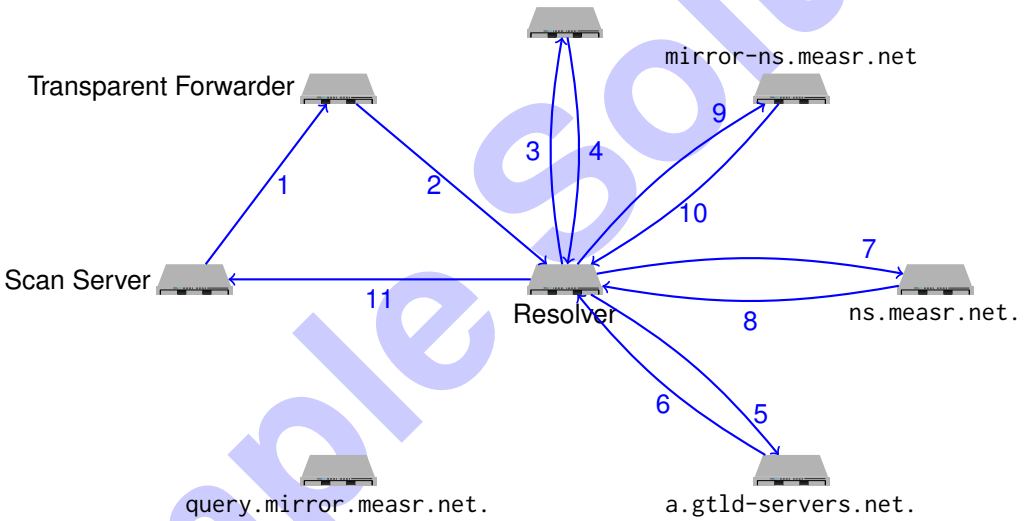

Additional space for solutions–clearly mark the (sub)problem your answers are related to and strike out invalid solutions.