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 / Endterm	**Date:**	Wednesday 18th February, 2026
Examiner:	Prof. Dr.-Ing. Georg Carle	**Time:**	14:00 – 15:15

Working instructions

- 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** distributed together with the exam
  - 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.
- 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.

Left room from _____ to _____ / Early submission at _____

Problem 1 Quiz (14 credits)

The following questions cover multiple topics. **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:ex:81:f0`

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

b)* What does an Ethernet switch do when it receives a frame with a destination MAC address that is not present in its MAC address table?

- ☐ It sends the frame only to the Default Gateway.
- ☐ It asks the Root Bridge for instructions.
- ☐ It drops the frame immediately.
- ☒ It floods the frame out of all ports except the receiving port.

c)* What is the primary benefit of VXLAN over traditional 802.1Q VLANs?

- ☐ Native encryption
- ☒ Scalability up to 16M IDs
- ☐ None of the listed options
- ☐ Reduced MTU size
- ☐ Layer 1 link speed
- ☐ Lower header overhead

d)* Which transport layer protocol is used to encapsulate VXLAN traffic?

- ☐ ICMP
- ☐ SCTP
- ☒ None of the listed options
- ☐ TCP
- ☐ QUIC
- ☐ IGMP

e)* What is the primary difference between the "Control Plane" and the "Data Plane" in an SDN architecture?

- ☐ The Control Plane is local to the switch; the Data Plane is centralized.
- ☐ The Control Plane is implemented in hardware; the Data Plane in software.
- ☒ The Control Plane makes decisions; the Data Plane forwards packets.
- ☐ The Control Plane handles user data; the Data Plane handles signaling.

f)* What is **Network Function Chaining (NFC)** in the context of Network Function Virtualization (NFV)?

- ☐ A technique for load balancing traffic across different CPU cores.
- ☐ Connecting multiple virtual machines using a software switch.
- ☒ Directing traffic through a specific sequence of virtual network functions.
- ☐ A method for encrypting virtual machine traffic.

g)* What is the primary mechanism DPDK's **Poll** Mode Driver (PMD) uses to improve performance?

- ☐ Compressing packets before processing
- ☐ Increasing the number of hardware interrupts
- ☒ Continuously checking the NIC for new packets
- ☐ None of the listed options

h)* Why is the arithmetic mean or median often a misleading metric for assessing user experience in network latency?

- ☐ It overemphasizes the fastest 10 % of packets.
- ☐ It ignores the propagation delay of the physical medium.
- ☐ It is too computationally expensive to calculate in real-time.
- ☒ It hides rare outliers that affect a service.

i)* If a latency measurement reports a 99th percentile of 250 ms, what does this imply?

- ☐ The average latency of the slowest 1 % of users is 250 ms.
- ☐ Only 1 % of the total requests were successful.
- ☐ 99 % of all measurements have a latency of exactly 250 ms.
- ☒ 99 % of all requests were completed in 250 ms or less.

j)* Which of the following options is the definition of **reproducibility** according to ACM?

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

k)* Which of the following statements about the **Probe Bandwidth** phase in TCP BBR is true?

- ☐ The data inflight is reduced to only 4 segments for 200 ms + 1 RTT
- ☐ The data inflight is reduced to 70 % of the previous sending rate for 200 ms + 1 RTT
- ☒ The sending rate is increased by 25 % every eighth RTT for 10 s
- ☐ The sending rate is increased by 1 MSS every eighth RTT for 9 s

l)* What is the goal of flow control? How is the sending rate changed?

0
1

Goal: Avoid overwhelming the receiver's buffer.

Sending rate adjustment: Sending rate is adjusted based on the receiver's advertised window size.

We now consider a bloom filter with the following three hash functions $\forall_{i \in \{1,2,3\}} h_i : \{0,1\}^* \to \{0,1\}^4$: (1) h_1('world') = 2, (2) h_2('world') = 4, (3) h_3('world') = 8.

Index	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Value	0	0	1	0	0	0	0	0	0	0	0	0	1	0	1	1

Figure 1.1: Bloom filter

m)* Given the bloom filter (Fig. 1.1), determine the result of the bloom filter lookup for the element x = 'world'.

- ☐ True: the element was definitely added to the filter beforehand
- ☒ False: the element was definitely not added to the filter beforehand
- ☐ True: the element was likely added to the filter beforehand
- ☐ False: the element was likely not added to the filter beforehand

n)* Given the bloom filter in Figure 1.1, add the element x = 'world' to the bloom filter in the box below. Only fill in changed values.

0
1

Index	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Value before addition	0	0	1	0	0	0	0	0	0	0	0	0	1	0	1	1
Value after addition			1		1				1							

Problem 2 Transport Layer (10 credits)

You are optimizing the network stack for "LittleLimeLever," a real-time video conferencing application. The application currently uses HTTP / 2 over TCP but the performance is often suboptimal. You are thinking about different aspects to improve the user experience.

0 1 a)* When a user walks out of the office, switching from Wi-Fi to 5G, the video call suddenly ends. Shortly explain, why TCP cannot maintain the connection in this scenario.

- TCP connections are identified by a **5-tuple**: (L4 protocol, src IP, src port, dst IP, dst port).
- When the user switches from Wi-Fi to 5G, the **source IP (and source port) change**. The connection is terminated.

0 1 b) Which mechanism in QUIC would fix this issue? Why does it recognize the connection after the network change?

Mechanism: **Connection Migration**

Why it works: It uses **Connection IDs** instead of the 5-tuple to identify the connection.

0 1 c)* Figure 2.1 shows how the RTT of a TCP connection changes with increasing data inflight in an ideal environment. Update Figure 2.2 accordingly.

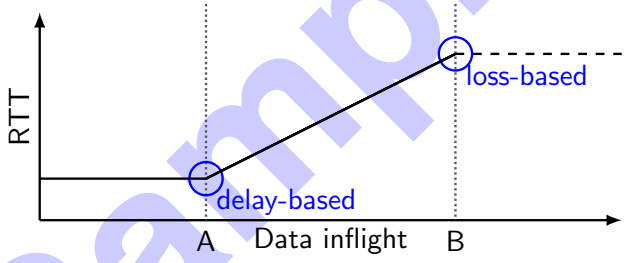

Figure 2.1: RTT

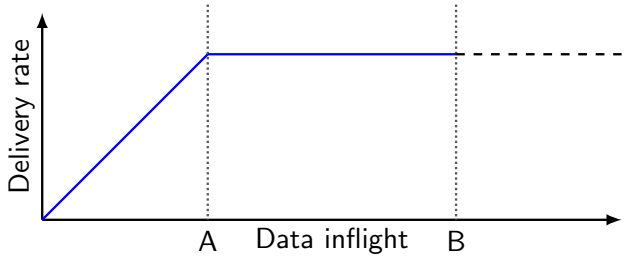

Figure 2.2: Delivery rate

0 1 d)* How large is the amount of inflight data at the labels A and B in Figures 2.1 and 2.2?

A: BDP

B: BDP + bottleneck buffer size

0 1 e)* Mark and label the operation point of loss-based and delay-based congestion control algorithms in Figure 2.1.

f)* Which type of congestion control algorithm (CCA) would you suggest to use for this application: loss-based, delay-based, or model-based? Justify your choice.

0
1

CCA Type: model-based or delay-based

Reason: Loss-based algorithms fill up buffers, causing increased latency.

For real-time applications like video conferencing, we need low latency.

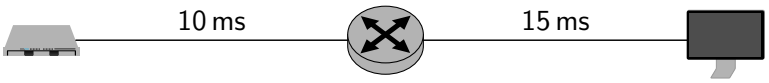

Figure 2.3: Sample network topology with propagation delays.

g)* Client and server are connected via one router (see Figure 2.3). The label on each link indicates the one-way propagation delay. The router limits throughput to 80 Mbit/s in both directions and has a 5 kB buffer per direction. You can assume no other traffic in the network and that the connection was established more than a minute ago. What would the RTT be with TCP BBR?

0
1

RTprop = 2 * (10 ms + 15 ms) = 50 ms

BBR keeps buffers empty -> no queuing delay

RTT = RTprop = 50 ms

h) Would the RTT be different if we would use TCP CUBIC? If yes, calculate the RTT with CUBIC. If no, explain why not.

0
1
2
3

Assuming full buffers: queuing delay has to be added

- Buffer size = 5 kByte = 40 kbit
- Bottleneck rate = 80 Mbit/s = 80kbit/ms
- Queuing delay = 40 kbit / 80kbit/ms = 0.5 ms
- RTT = RTprop + 2 * Queuing delay = 50 ms + 2 * 0.5 ms = 51 ms

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

Problem 3 Wired shark (15 credits)

You are a network administrator at “LittleLimeLever” and maintain a part of their network infrastructure shown in Figure 3.1:

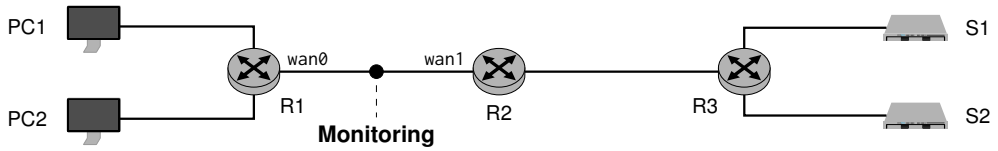

Figure 3.1: Network topology

Today, you are tasked with debugging a problem in a legacy video stream. You use Wireshark to monitor frames on the link between routers R1 and R2 and see multiple frames **from PC1 that should reach server S2**, but do not. Figure 3.2 shows such a frame.

```
                (b)                                     (a)                     (c)
0x0000  30 ee a4 4b 43 d5 e0 63   da 89 f1 ee 86 dd 60 07
                      (f)                         (d)
0x0010  45 3d 00 84 06 3f 20 01   0d b8 11 ab 00 01 32 44
                   (d)
0x0020  d7 ff fe 65 dc 73 20 01   0d b8 11 ab 00 0d 00 00
                                        (g)
0x0030  00 00 00 00 00 02 c2 65   01 bb a2 77 24 07 00 00
                (h)
0x0040  00 00 80 02 fa f0 ee b8   00 00 01 01 08 0a b8 ba
0x0050  4a 00 2b cf d4 80 a8 52   a7 c9 99 1d 82 10 37 ..
```

Figure 3.2: Ethernet frame (truncated)

First, you decide to gather some information about the involved devices and protocols.
Please note: Depending on the subproblem you must or should mark relevant parts of the hexdump—ensure that the markings are clearly associated with the subtask. Answers where the solution approach is not documented sufficiently **will not be graded**. Note down MAC and IP addresses in their respective formats and, in case of IPv6, their shortened form.

0 1 a)* What is the MAC address of R1.wan0? Only name and mark the respective header field.

Address: `e0:63:da:89:f1:ee` Field: Source MAC

0 1 b)* What is the MAC address of R2.wan1? Only name and mark the respective header field.

Address: `30:ee:a4:4b:43:d5` Field: Destination MAC

0 1 c)* Which network layer protocol is used?

Protocol: IPv6 Reasoning: EtherType = `0x86dd`

0 1 2 d) What is the IP address of PC1? Only name and mark the respective header field.

Address: 2001:db8:11ab:1:3244:d7ff:fe65:dc73 Field: Source IP

e) Through which mechanism was the IP address of PC1 most likely obtained? Explain shortly.

0 1

SLAAC, because of the clearly distinguishable 64bit EUI-64 Identifier with the constant value `ff:fe` in the middle

f) Which transport layer protocol is used?

0 1

Protocol: TCP Reasoning: IPv6 Next Header = `0x06`

g) Which application layer protocol is most likely used?

0 1

Protocol: HTTPS Reasoning: TCP Destination Port = 443

h) At what index (e.g. `0x0065`) in the frame does the PDU of layer 5 start?

0 1 2

Index: 0x0056 Reasoning: TCP Offset is `8` → TCP header is 32 B long

Then, you see another packet on the monitored link with an ICMPv6 message that provides additional information about the underlying problem. Figure 3.3 shows only the header and payload of the ICMP message.

(i)

0x0000	01	00	ff	ff	00	00	00	00	60	07	45	3d	00	84	06	3f
0x0010	20	01	0d	b8	11	ab	00	01	32	44	d7	ff	fe	65	dc	73
0x0020	20	01	0d	b8	11	ab	00	0d	00	00	00	00	00	00	00	02
0x0030	c2	65	01	bb	a2	77	24	07	00	00	00	00	80	02	fa	··

Figure 3.3: ICMPv6 message (truncated)

i)* What is the underlying cause of the ICMP message? Remember to name and mark respective fields.

0 1

Message cause: Destination Unreachable, no route to host

Reasoning: ICMP type `0x01` and code `0x00`

j) Argue which device in the topology created this ICMP message.

0 1

R2: The TTL in the original datagram's IPv6 header contained in the ICMP message is still `0x3f`, therefore the original packet has not been forwarded by R2.

k) Explain what you might need to change in your network such that the packets are successfully received.

0 1

Update the routing table of R2 to successfully forward packets toward the servers.

Problem 4 P4 Switching (15 credits)

This problem investigates a Software-Defined Network (SDN) powered by P4. Parts of the source code for the P4 switch program are given in Listing 1. This program processes frames with and without VLAN header.

```
header eth_t        { bit<48> dstAddr;
                      bit<48> srcAddr;
                      bit<16> etherType; }
header veth_ext_t   { ______ pcp; // see Subproblem a)
                      ______ dei; // see Subproblem a)
                      ______ vid; // see Subproblem a)
                      bit<16> etherType; }
struct std_meta     { bit<16> ingress_port; }
struct meta         { //unused
                      }
struct headers      { eth_t eth;
                      veth_ext_t veth_ext; }

parser ParserImpl(packet_in packet, out headers hdr, inout meta meta, inout standard_metadata_t
    std_meta) {
  state parse_eth {
    packet.extract(hdr.eth);
    transition select(hdr.ethernet.etherType) {
      16w0x8100: parse_veth_ext;
      default: accept;
    }
  }
  state parse_veth_ext {
    packet.extract(hdr.veth_ext);
    transition accept;
  }
  state start {
    transition parse_eth;
  }
}

control Pipeline(inout headers hdr, inout metadata meta, inout standard_metadata_t std_meta) {
  action drop() {
    mark_to_drop();
  }
  action decap(bit<16> egress) {
    std_meta.egress_port = egress;
    hdr.eth.etherType = hdr.veth_ext.etherType;
    hdr.veth_ext.setInvalid();
  }
  table forward {
    actions = {
      decap;
      drop;
      NoAction();
    }
    key = {
      std_meta.ingress_port: exact;
      hdr.eth.srcAddr: exact;
      hdr.veth_ext.vid: exact;
    }
    size = 4;
    default_action = NoAction();
  }
  apply {
    if (hdr.veth_ext.isValid()) {
      forward.apply();
    }
  }
}

control DeparserImpl(packet_out packet, in headers hdr) {
  // see Subproblem b)
}

// ...
```

Listing 1: VLAN P4 program

a)* Complete the `veth_eth_t` header of Listing 1 (Lines 4–6).

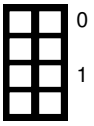


```
bit<3> pcp; // see Subproblem a)
bit<1> dei; // see Subproblem a)
bit<12> vid; // see Subproblem a)
```

b)* Create a valid deparser for the P4 program in Listing 1.

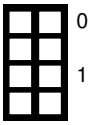


```
1    control DeparserImpl(packet_out packet, in headers hdr) {
       apply { packet.emit(hdr.eth);
3              packet.emit(hdr.veth_ext); }
     }
5
```

c)* For this problem, consider a packet that contains a **VXLAN** header. What happens to the packet according to the `ParserImpl` of Listing 1? Briefly explain which states are traversed and why.

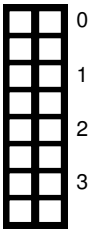

- VXLAN encapsulates Layer 2 traffic on top of UDP (Layer 4), i.e., VXLAN does not impact the actual processing inside the parser.
- If the frame has a VLAN tag, `start`, `parse_eth` and `parse_veth_ext` are traversed before the packet is accepted.
- If the frame has no VLAN tag, only `start` and `parse_eth` are traversed before the packet is accepted.

For the following problems, consider the network given in Figure 4.1. Switches 1 and 2 are **VLAN** capable switches. Servers A and C share a common VLAN (ID 16), Servers B and D also share a common VLAN (ID 32). All servers use regular Ethernet frames without any VLAN information. Ethernet frames with VLANs are only exchanged between the two Switches 1 and 2.

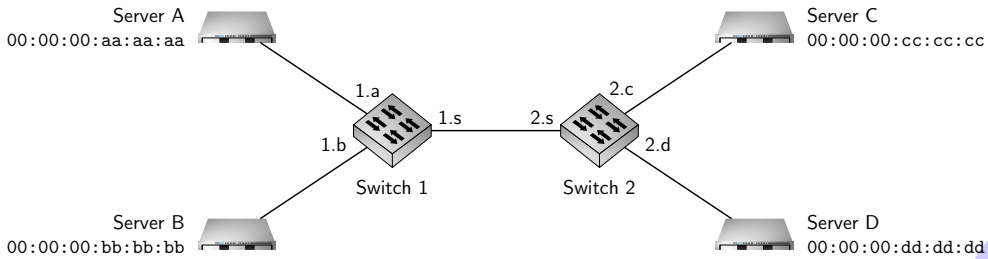

Figure 4.1: Network topology

0 1 2

d)* From the perspective of VLAN there are two different kinds of ports. **Name** the type of port for each interface of Switch 1 (1.a, 1.b, and 1.s) in Figure 4.1. Also **explain** what type of traffic is expected on the respective ports regarding the received frames' VLAN information.

- 1.a, 1.b: Access portstypically receive traffic without VLAN header information (untagged) from attached devices
- 1.s: Trunk portstypically receive traffic with VLAN header information (tagged) from attached devices

0 1 2 3

e)* Fill out the header information for a frame that was received on port 1.s in hexadecimal format. The frame was originally sent by Server D and contains an IPv6 packet.

```
hdr.eth.etherType = 0x8100
hdr.veth_ext.vid = 0x020
hdr.veth_ext.etherType = 0x86dd
```

0 1 2 3

f)* Fill out the truth table for different P4 match expressions.

Match type	Match value	1111_2	1101_2	1011_2	1010_2
`exact`	`0xB`	false	false	true	false
`ternary`	`*10*`	false	true	false	false
`lpm`	`0xB/4`	false	false	true	false

Problem 5 DNS (9 credits)

While debugging the video stream problem, you encounter a DNS message you do not expect on this link. Therefore, you decide to investigate this message in detail.

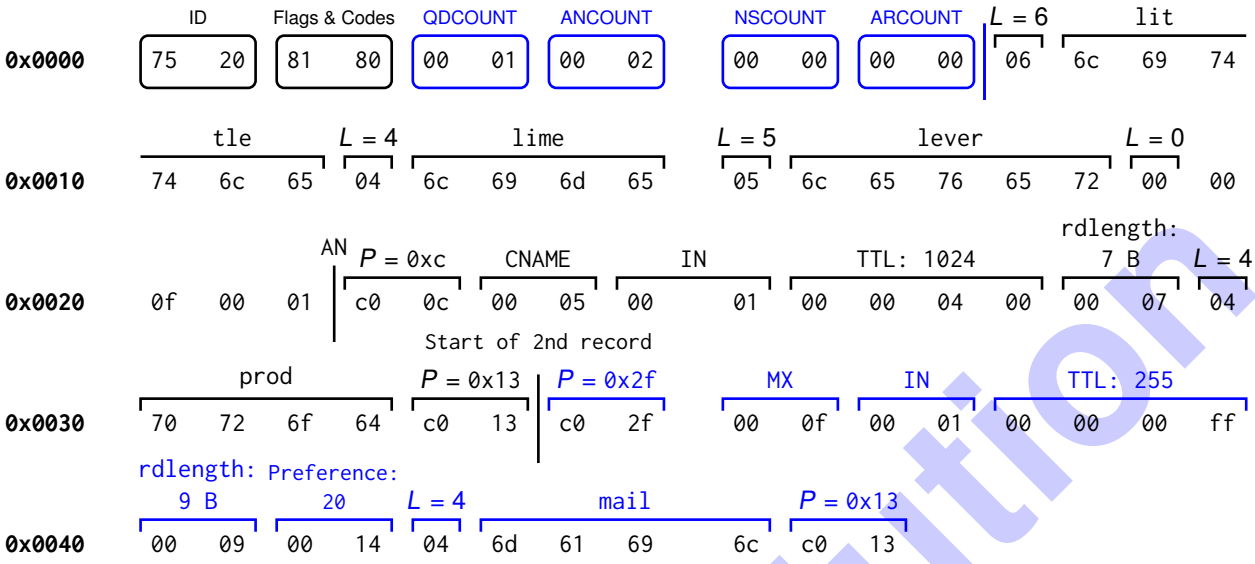

Figure 5.1: Hexdump of a DNS message

a)* Parse the header of the DNS message in Figure 5.1.

Field	Value	Name/Meaning
ID	34176	–
Flags	–	☒ QR ☐ AA ☐ TC ☒ RD ☒ RA
Opcode (Status)	0	QUERY
Rcode	0	NOERROR
QDCOUNT	1	Count of Queries RRs
NSCOUNT	0	Count of Authority RRs
ANCOUNT	2	Count of Answer RRs
ARCOUNT	0	Count of Additional RRs

The answer section begins with byte offset `0x0023` and is marked in Figure 5.1 by the label `AN`.

b)* What is the content of the answer section using the template below? Mark and label all used labels, pointers, and fields like the example shown in Figure 5.1.

NAME	TTL	CLASS	TYPE	RDATA
little.lime.lever.	1024	IN	CNAME	prod.lime.lever.
prod.lime.lever.	255	IN	MX	20 mail.lime.lever.

c) What type of DNS entity did likely send this message? Explain.

The message is a response (`qr` flag set), sent from a DNS entity supporting recursive resolution (`ra` flag). Therefore, this likely has been sent from a recursive resolver.

Problem 6 Routing (12 credits)

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 routes are equally long, use the next hop with the lower AS number.

0 1 2

a)* Draw the AS path for packets going from **AS 77** to **AS 40**. Take care to mark the path clearly.

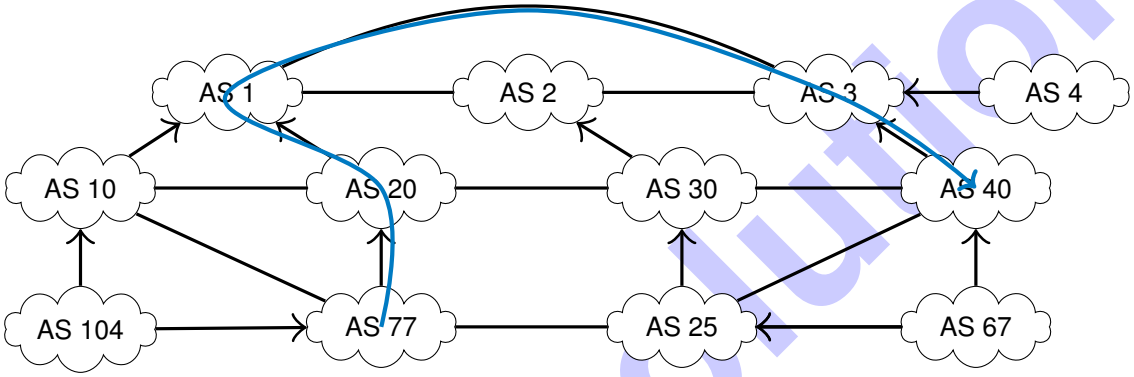

Figure 6.1: AS topology

AS 77 $\rightarrow$ AS20 $\rightarrow$ AS 1 $-$ AS 3 $\leftarrow$ AS 40

0 1

b)* List which ASes in the topology in Figure 6.1 are *Tier-1 providers*, *stub networks* and *multi-homed ASes*.

Tier-1 provider(s) AS 1, AS 2, AS 3

Stub AS(es) AS 4

Multi-homes AS(es) AS 67, AS 104

0 1 2

c)* Perform the k-core algorithm for the AS topology below. To visualize this, perform the following subtasks.

1. List the k-value for every node in the graph, i.e. the current k value when the node is removed.
2. Mark the core of the network after applying the k-core algorithm.

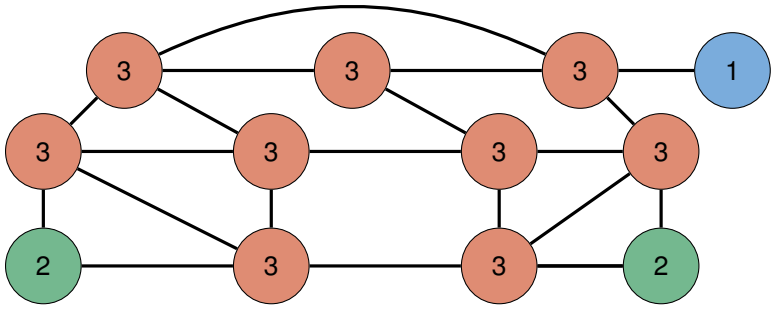

The core are the red nodes.

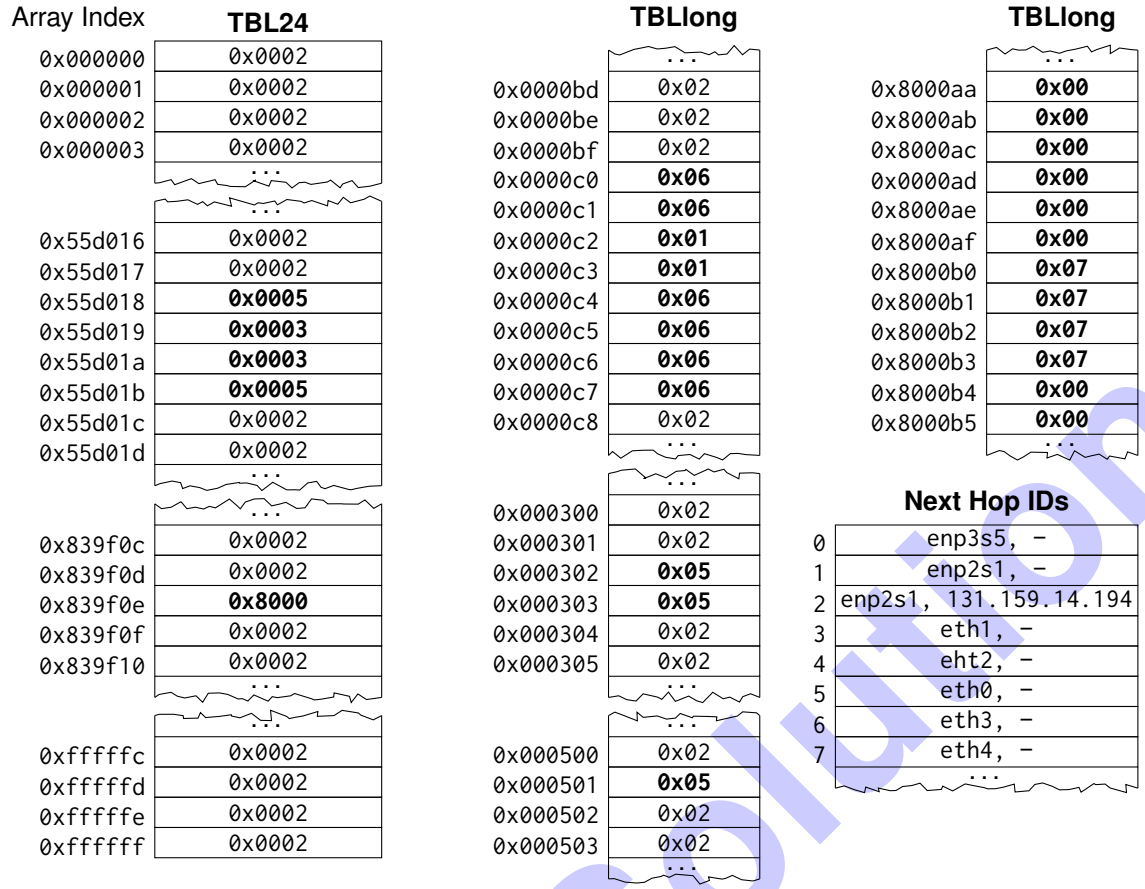

Figure 6.2: DIR-24-8 representation of a routing table

d)* Given the DIR-24-8 data structures in Figure 6.2, extract all visible routes into the routing table below. The routing table entries should be as aggregated as possible.
Hint: All entries abbreviated and not displayed in the figure have the value `0x0002/0x02`. There are more lines than necessary

0 1 2 3 4 5 6

Prefix	Interface	Next Hop	Next-Hop ID
131.159.14.194/31	enp2s1	—	1
131.159.14.192/29	eth3	—	6
85.208.25.0/24	eth1	—	3
85.208.26.0/24	eth1	—	3
85.208.24.0/22	eth0	—	5
0.0.0.0/0	enp2s1	131.159.14.194	2

e)* Name another advanced routing table data structure except DIR-24-8 (and LPM).

0 ½

Basic trie, path-compressed trie, level-compressed trie, path- and level-compressed trie or DXR.

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