Chair of Network Architectures and Services
TUM 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
Examiner: Prof. Dr.-Ing. Georg Carle
Date: Wednesday 12th February, 2025
Time: 16:00 – 17:15

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.
- If you used a second template in the back of the exam, mark this in the original template using the provided check box.
- 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.
- The multiple choice questions need to be filled out as follows:

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 ×■

Left room from ________ to ________ / Early submission at ________

Problem 1 Quiz (14 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

a)* Which of the following terms is **no** IPv4 to IPv6 transition mechanism?

☐ Dual Stack ☒ 6transition ☐ DS Lite ☐ 6to4

b)* Which of the following options can be used to **reasonably** perform large-scale IPv6 Internet scans?

☐ Hitmans ☒ Hitlists ☐ Broadcasts
☐ Full address enumeration ☐ Multiplicative group

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

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

d)* In which of the following DNS message sections can the `OPT` record be found?

☐ Answer ☐ Question ☐ Authority ☒ Additional

e)* How long is the RDATA of an `AAAA` record in bit?

☐ 136 ☒ 128 ☐ 40 ☐ 32

f)* What kind of load balancing is **definitely** used for `dell.com` given the following `dig` output?

```
endterm@acn.net.in.tum.de:~$ dig +short dell.com
143.166.136.12
143.166.30.172
```

☐ QUIC-based load balancing ☒ DNS-based load balancing
☐ Anycast-based load balancing ☐ HTTPS-based load balancing
☐ No load balancing ☐ TCP-based load balancing

g)* How does QUIC prevent head-of-line blocking?

☒ By making the transport layer stream-aware
☐ By retransmitting only lost stream frames instead of whole packets
☐ By using multiple QUIC connections for each stream
☐ By separating stream data into multiple packets

h)* Features of which layers of the ISO-OSI model are integrated in the QUIC protocol?

☐ Layers 2, 3, and 4 ☒ Layers 4, 5, and 7 ☐ Layers 3, 4, and 5 ☐ Layers 1, 2, and 3

Mixed Problems

The PC wants to communicate with the server using 6to4.

i)* Draw the path of a single request and response between the PC and the server. Use dashed lines for IPv4 traffic and solid lines for IPv6 traffic.

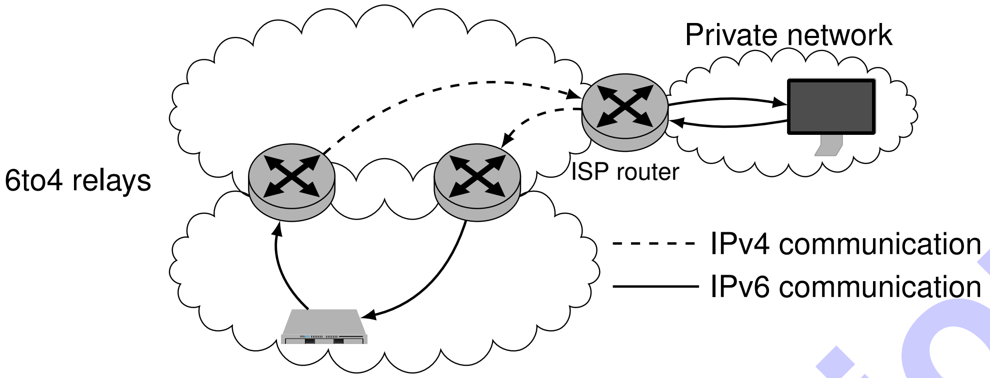

☐ Grade the second template in the back

j)* The protocols telnet and SSH should be blocked using a P4 switch via their well-known port numbers. How can both protocols be blocked using a single match?

Both protocols use neighboring ports (22 and 23)They can be blocked using a ternary match `1011*` (or via a range match from ports 22 to 23).

k)* Name the decisive difference used to recognize if a DNS message is a query or a response.

Both messages use the same format and only differ by the QR (response) bit.It is set for response messages.

l)* Explain what QNAME minimization is and name an advantage and one disadvantage of using it.

Explanation: The full QNAME is only sent to the authoritative name server, which is queried recursively for each label.

Advantage	Not all name servers learn the full FQDN
Disadvantage	Increased query load or rate of unsuccessful queries

Problem 2 Wireshark (17 credits)

Consider the network topology in Figure 2.2. The hexdump of the frame visualized in Figure 2.2, which arrived at router R2, is printed in Figure 2.1. No host in the network uses layer 4 options.
The frame in Figure 2.1 contains an IPv6 fragmentation extension header, and the `M` flag is set to 0. Thus, the frame contains the last fragment of a fragmented IPv6 packet.

☐ Grade the second template in the back

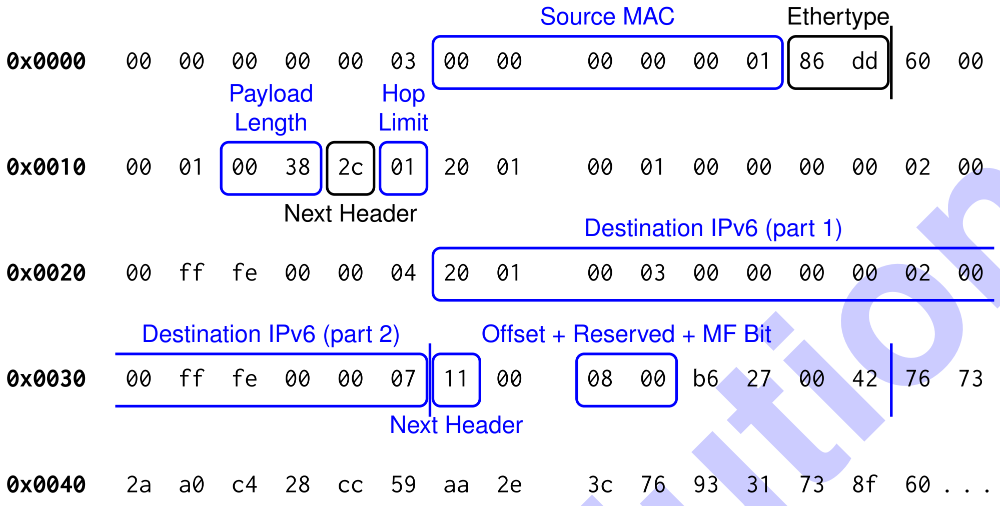

```
0x0000  00 00 00 00 00 03 00 00   00 00 00 01 86 dd 60 00
0x0010  00 01 00 38 2c 01 20 01   00 01 00 00 00 00 02 00
0x0020  00 ff fe 00 00 04 20 01   00 03 00 00 00 00 02 00
0x0030  00 ff fe 00 00 07 11 00   08 00 b6 27 00 42 76 73
0x0040  2a a0 c4 28 cc 59 aa 2e   3c 76 93 31 73 8f 60 . . .
```

Figure 2.1: Partial hexdump of a frame arriving at router R2 without FCS

a)* Separate all visible headers in Figure 2.1 with a vertical line like the given example.

b)* What is the source MAC address of the frame? Mark it in Figure 2.1.

`00:00:00:00:00:01`

c)* Mark the destination IPv6 address in Figure 2.1 and give it in shorthand, default notation.

`2001:3::200:ff:fe00:7`

d) How long is the **Layer 4 SDU** of the **reassembled** IP packet? Explain your steps in detail using the following table. Mark all header fields you used in Figure 2.1.

Hint: The IPv6 payload length includes extension headers.

Notes: Offset: `0b0000 1000 0000 0...`, ... is reserved and M flag ⟶ 256 · 8 B = 2048 B

(+/–) Size	Explanation
+ 2048 B	The offset in the fragmentation extension header
– 8 B	UDP header length as the packet does not have any options
+ 56 B	IPv6 payload length of this fragment (`0x38`)
– 8 B	IPv6 fragmentation header length
2088 B	$\sum$

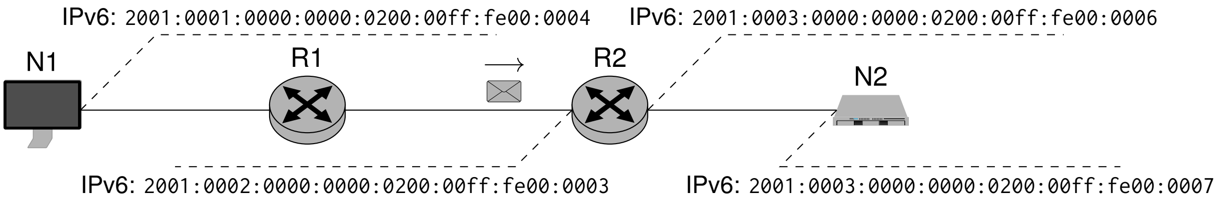

Figure 2.2: Network topology and position of the frame in Figure 2.1

We now consider the **response frame** Ψ being sent out as response to the arrival of the frame in Figure 2.1. Before the frame in Figure 2.1 arrived at router R2, all caches are filled.

e)* Frame Ψ contains an ICMPv6 time exceeded header. Explain why this is the case and what happens. Mark header fields needed for your explanation.

0
1
2

The hop limit is 1.Therefore, the original frame will be dropped by R2.Subsequently, an ICPMv6 Time Exceeded packet will be returned to the original sender.

f) Fill out the following header templates starting with the Ethernet header for frame Ψ.

0
1
2
3
4
5

1. Write IP and MAC addresses in their standard format, numbers either in hexadecimal (e.g., `0xf2`) or decimal (e.g., 242) notation.
2. Choose appropriate values for fields which are not fixed.

First header: Ethernet

0 B	Destination Address: `00:00:00:00:00:01`	
4 B		Source Address: `00:00:00:00:00:03`
8 B		
12 B	Ethertype: `0x86dd`	▨

IPv6 header

Second header: IPv6

0 B	Version: $6_{(10)}$	Traffic Class: $0_{(10)}$	Flow Label: `0x00001`	
4 B	Payload Length: $104_{(10)}$		Next Header: `0x3a`	Hop Limit: $64_{(10)}$
8 B 12 B 16 B 20 B	Source Address: `2001:2::200:ff:fe00:3`			
24 B 28 B 32 B 36 B	Destination Address: `2001:1::200:ff:fe00:4`			

ICMPv6 Time Exceeded Header

Third header: ICMPv6 Time Exceeded

0 B	Type: $3_{(10)}$	Code: $0_{(10)}$	Checksum
4 B	unused: $0_{(10)}$		

IPv6 Header of the original frame + the **full layer 3 payload**

Problem 3 Routing Table Data Structures: Trie-based Lookup (15 credits)

Consider the routing table and Next-Hop ID table given in Table 3.1.

Table 3.1: Routing table and corresponding next-hop ID table

(a) Routing table

№	Prefix	Prefix$_2$	Next-Hop	Iface
①	32.0.0.0/3	`0010...`	172.18.1.1	eno7
②	160.0.0.0/3	`1010...`	33.0.2.1	eno1
③	64.0.0.0/2	`0100...`	192.168.3.1	eno7
④	192.0.0.0/2	`1100...`	10.0.0.1	eno1
⑤	128.0.0.0/1	`1000...`	33.0.2.1	eno1
⑥	0.0.0.0/0	`0000...`	192.168.2.1	eno7

(b) Next-hop ID table for the tries

NH-ID	Next-Hop	Iface
A	192.168.2.1	eno7
B	33.0.2.1	eno1
C	192.168.3.1	eno7
D	172.18.1.1	eno7
E	10.0.0.1	eno1
X	empty node	

a)* Build a basic trie representing the routing table in Table 3.1a. X marks empty nodes.

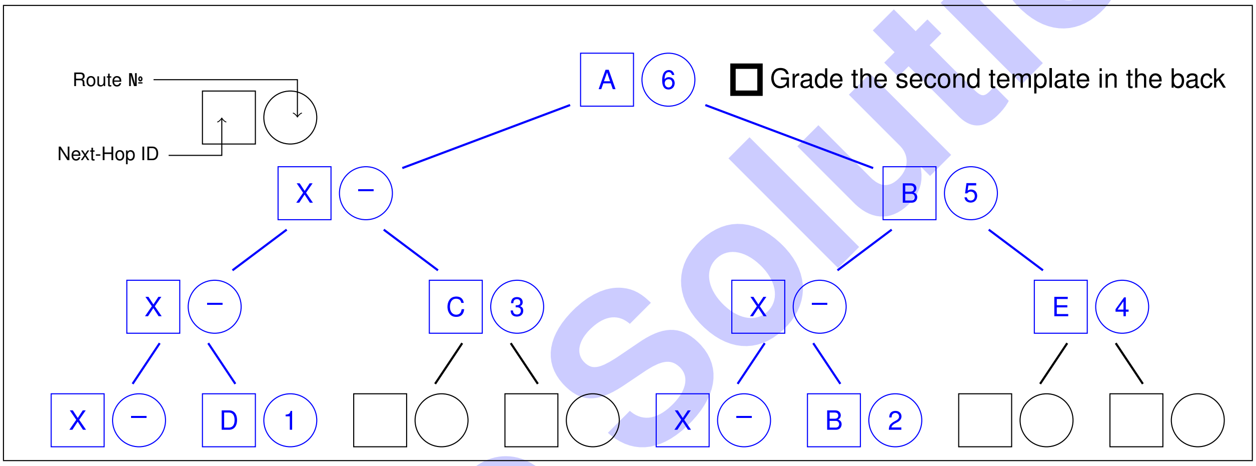

☐ Grade the second template in the back

b)* Perform a look up for 188.273.281.97 in the following trie. Circle the final leaf in the trie and explain your decision. *Remark:* The letters represent Next-Hop IDs. X marks empty nodes.

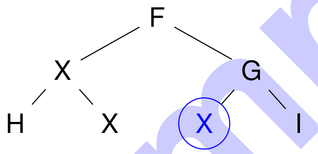

Selected Next-Hop ID: G

$188_{(10)} = 10111100_{(2)}$.Thus, the path 10 is taken. Since we take the final non-empty node, G is the next-hop ID.

c)* Briefly explain a problem of basic tries used as a routing table data structure.

Wasteful, as most long prefixes create lots of empty nodes.

d) Name an alternative trie-based data structure that solves this problem.

Level-compressed (L) trie, path-compressed (P) trie, or LP-compressed trie

e) Explain briefly what differentiates your alternative data structure from a basic trie.

☐ 0
☐ 1

Path-compressed trie Aggregate chains of empty nodes into one (or zero) nodes.

Level-compressed trie 2^n children in a node, better cache utilization.

The size of a data structure and the number of accesses to the data structure impact the performance of lookups. Routing table data structures often optimize for one of both parameters.

f)* How many accesses need to be performed on trie-based lookup structures for IPv4 routing tables in the worst case?

☐ 0
☐ 1
☐ 2

To represent 32-bit IP addresses, the trie uses 32 levels.This is also the maximum number of memory accesses.

g) For which parameter (data structure size or memory accesses) are the basic trie lookups and DIR-24-8 optimized? Explain your reasoning in the box below.

☐ 0
☐ 1
☐ 2

DIR-24-8 optimizes for ☒ memory accesses ☐ data structure size.
Basic trie lookups optimize for ☐ memory accesses ☒ data structure size.

DIR-24-8 optimizes for memory accesses (worst-case 2 memory accesses), the basic trie lookups have a worst-case of 32 accesses.The DIR-24-8 uses a fixed size for TBL24 with 2^{24} entries, the trie can be significantly smaller especially if the trie is not fully populated.

Problem 4 Network Analysis (9.5 credits)

Consider the network topology depicted in Figure 4.1, where two hosts H1 and H2 want to communicate with each other. Host H1 accesses the network of H2 via an SDN-enabled switch and two routers R1 and R2. Host H2 is directly attached to the routers R1 and R2.

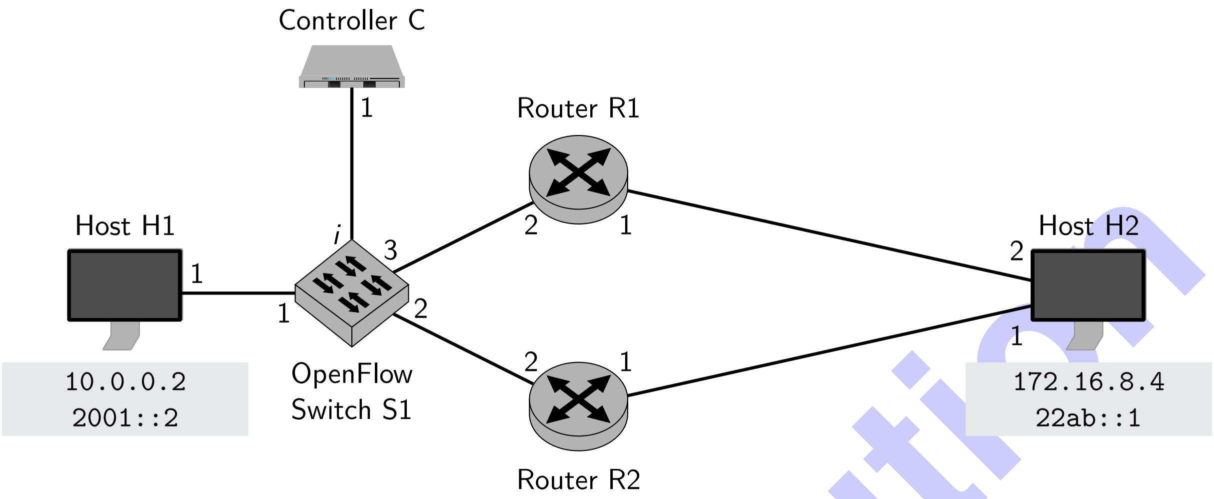

Figure 4.1: Network topology

The OpenFlow switch has a configuration interface named *i*, reachable at `192.168.0.2`. Furthermore, the switch has three output ports named 1 to 3. Listing 1 gives the commands, which configured this switch.

Listing 1: Open vSwitch table entries

```
1 ovs-ofctl add-flow tcp:192.168.0.2 dl_type=0x86dd,nw_dst=22ab::1,priority=10000,actions=output:3
2 ovs-ofctl add-flow tcp:192.168.0.2 dl_type=0x0800,nw_dst=172.16.8.4,priority=10000,actions=output:2
3 ovs-ofctl add-flow tcp:192.168.0.2 priority=0,actions=controller
```

a)* What is the general effect of specifying a `dl_type` in an OpenFlow rule? Where is the `dl_type` specified in Ethernet?

Rules are only applied for the specified protocol type, `dl_type` explicitly states the type of the link layer payload.The Ethertype field of the Ethernet header specifies the `dl_type`.

b) Look at Lines 1 and 2 of the commands shown in Listing 1. What is the effect of the different values for `dl_type`.

- Rule in Line 1 is only valid for `dl_type = 0x86dd` (IPv6).
- Rule in Line 2 is only valid for `dl_type = 0x0800` (IPv4).

c) Look at Line 1 of the commands shown in Listing 1. Explain (1) what this command does and describe (2) what the arguments `tcp:192.168.0.2`, `dl_type`, `nw_dst` and `actions` do in this example.

- This rule specifies to send packets destined to a certain address to be sent out on a specified egress interface of the switch
- `tcp:192.168.0.2` specifies the receiver of the OpenFlow rule, i.e., the interface of the controller
- `dl_type` specifies IPv6 to be used as payload of the link layer
- `nw_dst` specifies that the destination IP address of the IPv6 packet has to be `22ab::1`
- `actions` specify output port 2 as the egress interface of the switch.

d) The IP address of Interface 2 of Router R2 is not given in Figure 4.1. Give a sensible example for an IP address for that interface.

Any address from `10.0.0.0/8` except the address of H1.

For security reasons, Controller C is configured to drop ICMP packets.

e) All participating hosts and routers already know the MAC addresses of each other. Host H1 pings `172.16.8.4`. Despite ping packets arriving at H2, the response packets are not received by H1. Describe the way of the ping packets from Host H1 to H2 and back. Base your explanation on the interfaces (e.g., H1.1) given in Figure 4.1 and the rules specified in Listing 1.

H1.1 → S1.1 → S1.2 → R2.2 → R2.1 → H2.1
H2.1 → R2.1 → R2.2 → S1.2 → S1.i (matches rule 3) → C1.1 (drop)

f) What rule(s) has/have to be installed on Switch S1 to receive the replies at Host H1.

```
ovs-ofctl add-flow tcp:192.168.0.2 dl_type=0x0800,nw_dst=10.0.0.2,priority=10000,actions=output:1
```

Problem 5 Transport Layer Mechanisms (11.5 credits)

Transport protocols like QUIC or TCP offer various features, e.g., reliability.

a)* How do these protocols detect lost packets? Name and briefly explain two mechanisms that are usually used to detect loss and trigger retransmissions.

① **Retransmission Timeout (RTO)**: Wait for RTT+X, if data is not acknowledged: retransmit.

② **Duplicate ACKs**: If multiple (usually 3) ACKs of the same sequence number are received: fast retransmit.

An accurate RTT estimation is crucial for one of the previously mentioned mechanisms.

b) Shortly describe what happens if the RTT is overestimated or underestimated.

RTT overestimated: RTO is too long, retransmissions are delayed, slow reaction to loss, worse overall performance.

RTT underestimated: RTO is too short, unnecessary (=spurious) retransmissions, more traffic, worse overall performance.

c) Which mechanism is used in TCP to smooth the RTT estimation? What influences the current estimation?

Mechanism: Exponential weighted moving average (EWMA)

What influences the current estimation? The previous estimation

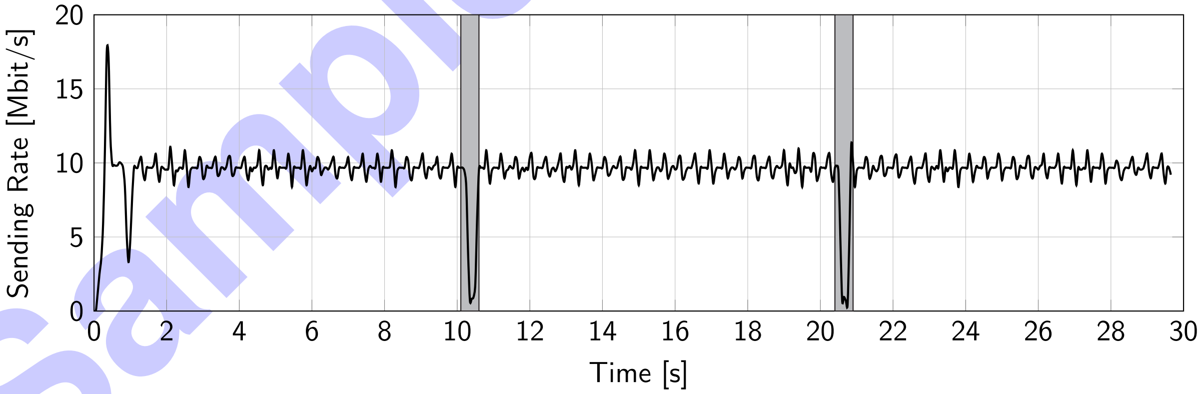

Figure 5.1: Sending rate over time for a TCP connection using the BBR algorithm.

Congestion control algorithms like Reno, CUBIC, and BBR try to prevent overloading the network with too many packets. The BBR algorithm defines four phases.

d)* Name the phase marked in Figure 5.1 and briefly explain why the sending rate spikes downwards.

Name: ProbeRTT

Explanation: To measure the RTT, we need empty buffers to not include the buffering delay. Therefore, the sending rate is reduced to empty the buffers and improve the RTT sample's quality.

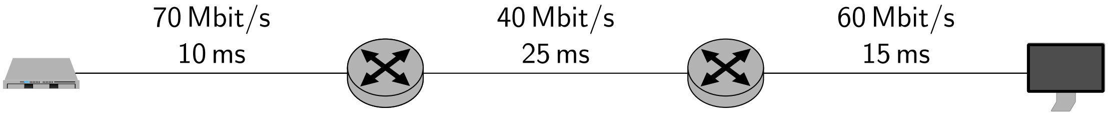

Figure 5.2: Sample network topology with link bandwidths and propagation delays.

e)* The BBR algorithm always tries to keep one so-called BDP of data inflight: (1) What is the BDP, and (2) how is it calculated?

☐ 0
☐ 1

$\mathrm{BDP} = \mathrm{RTprop} \cdot \mathrm{BtlBw}$

BDP Bandwidth-Delay Product

RTprop Round-trip propagation delay

BtlBw Bottleneck bandwidth

f)* Calculate the BDP for the link between client and server shown in Figure 5.2. Only consider one direction.

☐ 0
☐ 1

$$\mathrm{RTprop} = 10\,\mathrm{ms} + 25\,\mathrm{ms} + 15\,\mathrm{ms} = 50\,\mathrm{ms}$$
$$\mathrm{BtlBw} = \min\left(40\,\mathrm{Mbit/s}, 60\,\mathrm{Mbit/s}, 70\,\mathrm{Mbit/s}\right) = 40\,\mathrm{Mbit/s}$$
$$\mathrm{BDP} = 50\,\mathrm{ms} \cdot 40\,\mathrm{Mbit/s} = 2000\,\mathrm{kbit}$$

g) Name one disadvantage each for having less/more data inflight than one BDP?

☐ 0
☐ 1
☐ 2

Less than one BDP: Lower delivery rate, link not fully utilized.

More than one BDP: Delivery rate does not increase but latency does. Packets are buffered at the bottleneck link.

Problem 6 Network Calculus (8 credits)

This problem investigates delay bounds in networks using Network Calculus.

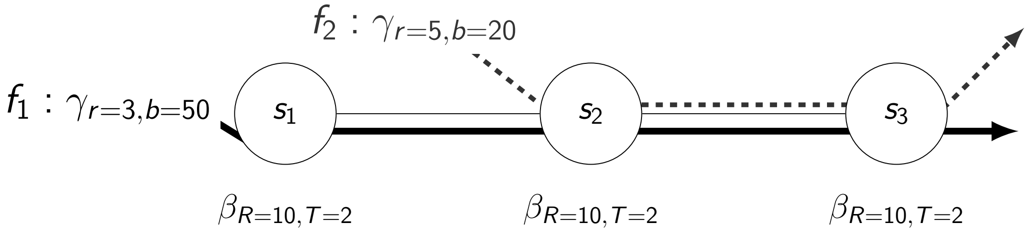

Figure 6.1: Topology with server- and flow specifications

Consider the topology in Figure 6.1. 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: $\beta^{l.o.1}_{s_1} = \beta_{10,2}$
- Left-over service curve at s_2: $\beta^{l.o.1}_{s_2} = \left[\beta_{10,2} - \gamma_{5,20}\right]^+ = \beta_{10-5,\frac{20+10\cdot 2}{10-5}} = \beta_{5,8}$
- Output arrival curve of f_2 after s_2: $\alpha^{\star} = \gamma_{5,20+5\cdot 2} = \gamma_{5,30}$
- Left-over service curve at s_3: $\beta^{l.o.1}_{s_3} = \left[\beta_{10,2} - \alpha^{\star}\right]^+ = \beta_{10-5,\frac{30+10\cdot 2}{10-5}} = \beta_{5,10}$

0 ☐ 1 ☐

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

$$\beta^{l.o.}_{e2e} = \beta^{l.o.}_{s_1} \otimes \beta^{l.o.}_{s_2} \otimes \beta^{l.o.}_{s_3} = \beta_{\min(10,5,5),2+8+10} = \beta_{5,20}$$

0 ☐ 1 ☐

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

$$d_{e2e} = T_{e2e} + \frac{b}{R_{e2e}} = 20 + \frac{50}{5} = 30$$

d) Assume the following changes to the scenario in Figure 6.1:

- The rate of Flow f_1 is set to $r = 7$
- The burst of Flow f_2 is set to $b = 10$

What is the end-to-end delay bound of f_1?

Infinite delay bound due to $r_1 + r_2 > R_2$

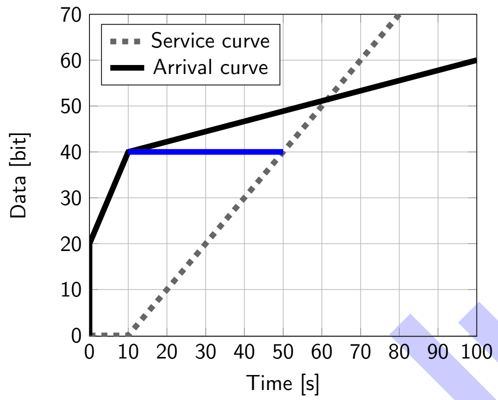

Figure 6.2: Arrival- and service curve

e)* Consider the two curves in Figure 6.2. What is the delay bound of a flow with this **non-**token-bucket arrival curve traversing a server with this rate-latency service curve?

Largest horizontal deviation is $50s - 10s = 40s$

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

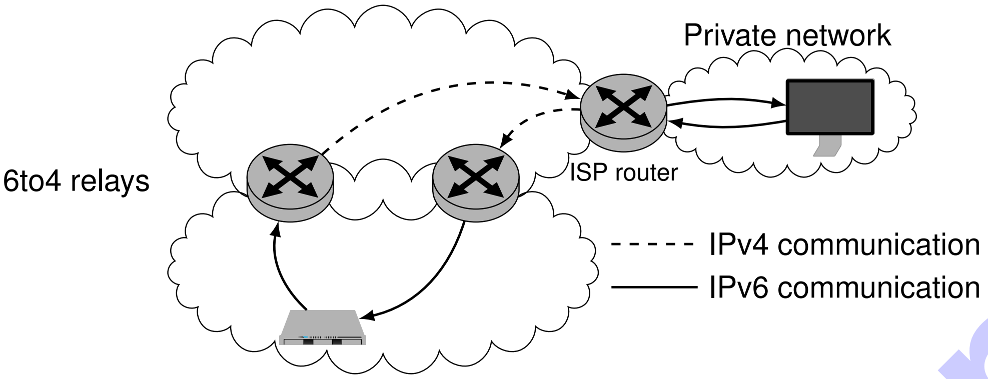

☐ Grade the second template in the back

☐ Grade the second template in the back

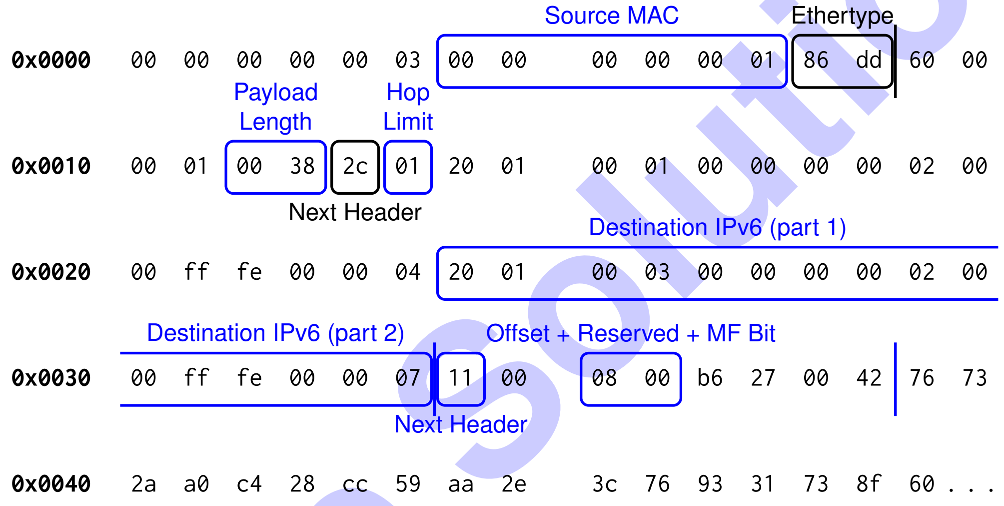

```
0x0000   00 00 00 00 00 03 00 00  00 00 00 01 86 dd 60 00
0x0010   00 01 00 38 2c 01 20 01  00 01 00 00 00 00 02 00
0x0020   00 ff fe 00 00 04 20 01  00 03 00 00 00 00 02 00
0x0030   00 ff fe 00 00 07 11 00  08 00 b6 27 00 42 76 73
0x0040   2a a0 c4 28 cc 59 aa 2e  3c 76 93 31 73 8f 60 . . .
```

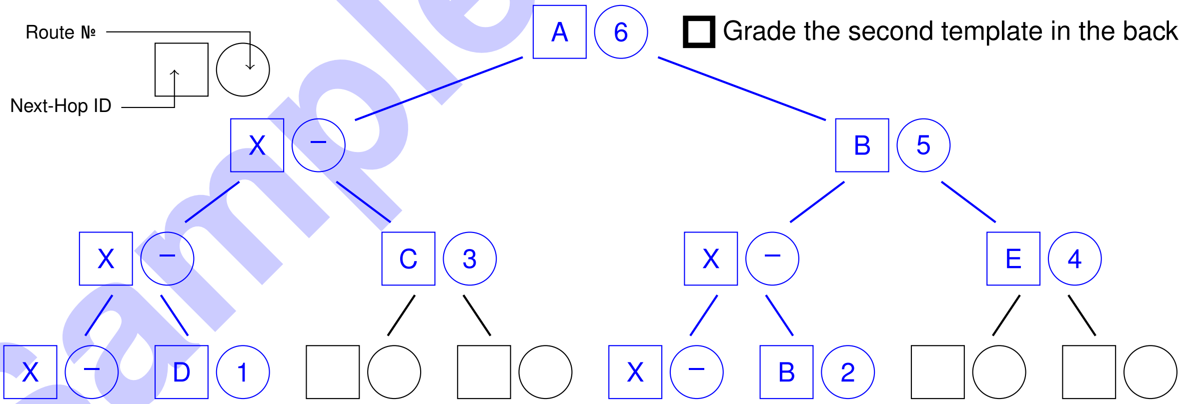

☐ Grade the second template in the back

First header: Ethernet

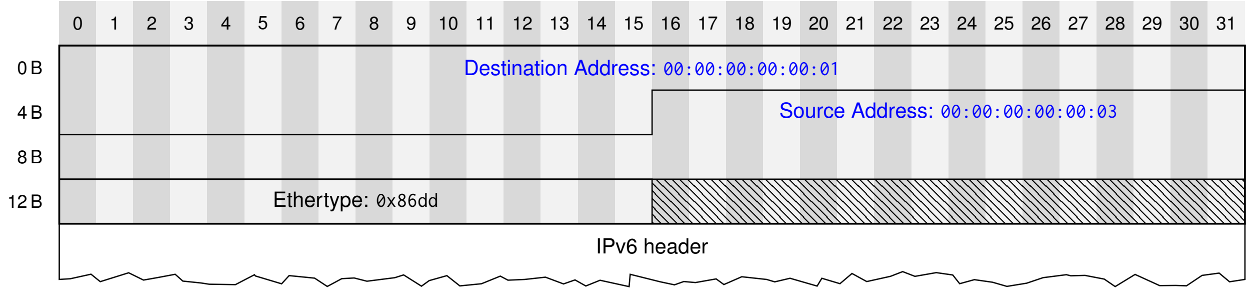

Second header: IPv6

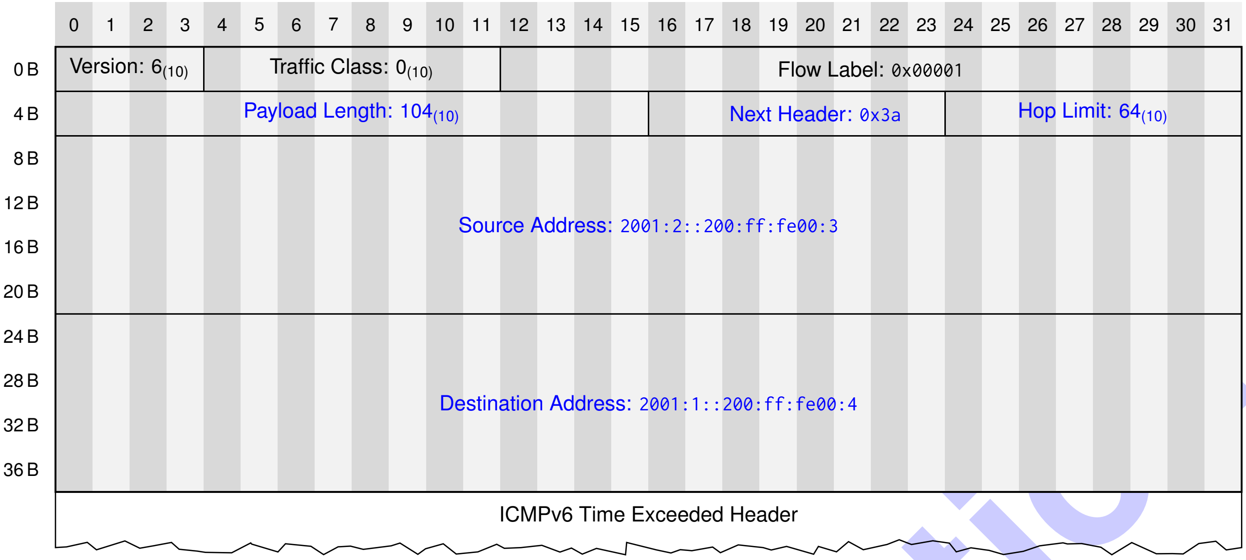

Version: $6_{(10)}$

Traffic Class: $0_{(10)}$

Flow Label: `0x00001`

Payload Length: $104_{(10)}$

Next Header: `0x3a`

Hop Limit: $64_{(10)}$

Source Address: `2001:2::200:ff:fe00:3`

Destination Address: `2001:1::200:ff:fe00:4`

ICMPv6 Time Exceeded Header

Third header: ICMPv6 Time Exceeded

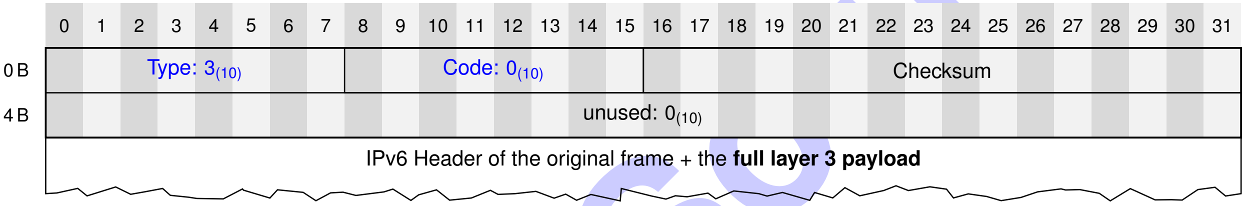

Type: $3_{(10)}$

Code: $0_{(10)}$

Checksum

unused: $0_{(10)}$

IPv6 Header of the original frame + the **full layer 3 payload**