

**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:** Monday 19<sup>th</sup> February, 2024

**Examiner:** Prof. Dr.-Ing. Georg Carle

**Time:** 13:30 – 14:45

### Working instructions

- This exam consists of **12 pages** with a total of **5 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:
  - 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 (18 credits)

The following questions cover multiple topics and can be solved independently of each other. 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



a)\* Which of the following is a correct IPv6 address?

fc21:2001:11:3223:3ff:fe74:150a

fc21::3223:3ff:fe74:150a

fc21::11:3223:3fg:fe74:150a

fc21::2001:11:3223::fe74:150a

b)\* Which of the following parameters requires DHCPv6 and is not available with ICMPv6-based configuration? **Invalid question and not graded:** An option to provide DNS information was added in a later RFC, but not properly mentioned during the lecture.

Static addresses

Netboot

Prefix information

DNS information

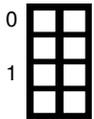
c)\* According to the Association for Computing Machinery (ACM), which of the statements below is correct.

Repeatability: **same** team using **different** experimental equipment.

Reproducibility: **different** team using **same** experimental equipment.

Recreatability: **different** team using **different** experimental equipment.

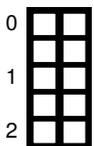
Replicability: **same** team using **same** experimental equipment.



d)\* Name the IPv6 alternative for ARP and briefly explain its mechanism.

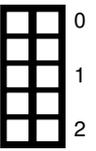


e)\* What is used to identify a TCP connection? What is different with QUIC to achieve IP mobility?

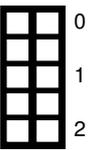


f)\* Shortly explain why traceroute might reveal non-existing paths and a solution from the lecture tackling this problem.

g)\* Perform the  $k$ -core algorithm for the topology shown in the solution box. For each  $k$ , list all removed nodes.



h)\* Table 1.1 contains all records a resolver receives as part of the authority response section while resolving `example.com`. The authoritative answer flag is not set. Shortly explain what this message means and name the next DNS message (including its destination) the resolver sends. Assume the resolver has no cache.



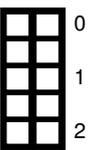
1	example.com.	300	IN	NS	dns1.lrz.de.
2	example.com.	7200	IN	NS	dns2.lrz.bayern.

Table 1.1: DNS Records received by the resolver. All records are part of the authority section.

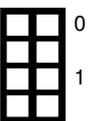
i)\* As an administrator of an authoritative name server: Would you prefer 300 or 7200 as the NS record's TTL? Argue according to lecture information and Briefly explain.



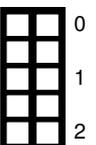
j)\* A company wants to create a Layer 2 connection between two sites in different countries via Internet. The Internet service provider (ISP) of the company only offers Layer 3 connectivity. The administrator proposes to create a VLAN tunnel between the two sites. Argue if the proposed solution works.



k)\* How do TCP Cubic and TCP BBRv1 react to packet-loss in the network?



l)\* **Name** and briefly **explain two** ways a TCP sender detects packet-loss and starts a retransmission.



## Problem 2 BGP and Routing (17.5 credits)

This problem investigates the autonomous system (AS) relationships in a given network and their impact on routing and traffic.

0			
1			
2			

a)\* Shortly explain the Valley-Free-Routing principle.

0			
1			

b)\* How can a Tier-2 provider be defined.

0			
1			

c)\* Based on the lecture, name the most important consideration for policy routing after longest prefix matching and a potential negative effect.

0			
1			
2			

d)\* What is the goal of longest prefix matching, and how can it naively be implemented?

0			
1			
2			

e)\* Which is the smallest prefix fully covering  $10.0.24.0/23$  and  $10.0.28.0/23$ .

For the following subproblems the topology in Figure 2.1 will be used.

- $\rightarrow$  represents a customer to ( $\rightarrow$ ) provider relationship.
- $\leftrightarrow$  represents a peering relationship.
- AS E owns and announces  $172.0.0.0/23$  while AS F owns and announces  $172.0.0.0/22$ .
- 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.

0			
1			
2			

f)\* A client in AS G wants to connect to  $172.0.1.27$ . To which AS is traffic routed and via which path? Explain your answer.

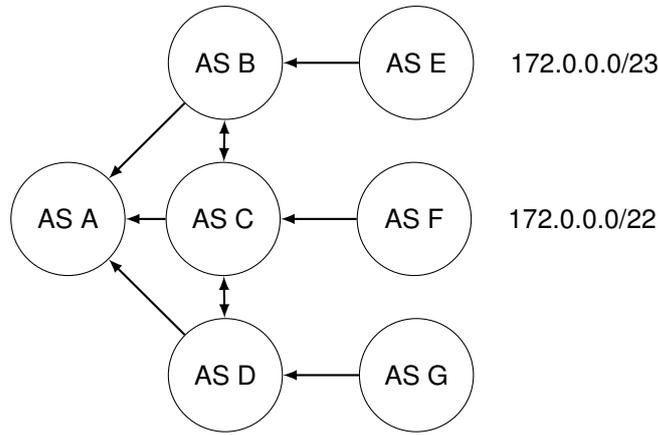


Figure 2.1: AS Network

g)\* A client in AS G wants to connect to 172.0.2.27. To which AS is traffic routed and via which path? Explain your answer.

	0
	1
	2

For the following subproblems, only consider the two prefixes 172.0.0.0/23 and 172.0.0.0/22.

h)\* How many different addresses can AS G reach within AS F and how many addresses can AS G reach in AS E?

	0
	1

i) Assume AS F additionally announces 172.0.1.0/24. How is the number of addresses within ASes E and F impacted that can be reached from AS G?

	0
	1

j)\* How can AS A intercept all traffic as man-in-the middle, without exposing itself as the origin of a route announcement? Explain your solution.

	0
	1
	2

### Problem 3 Hexdump (12.5 credits)

This problem investigates a captured Ethernet frame.

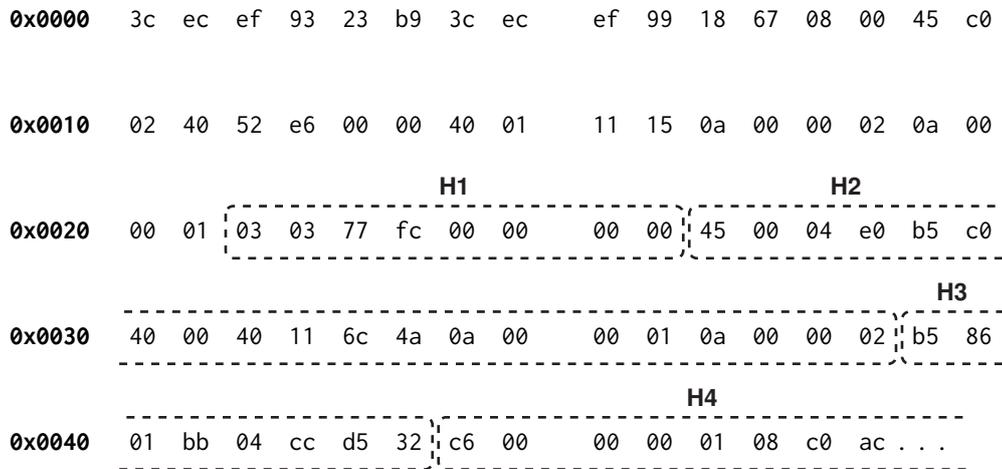


Figure 3.1: Hexdump of an Ethernet frame starting with the Ethernet header.

In this problem you **always** have to substantiate your answers using the bytes of the hexdump in Figure 3.1. Always make clear which bytes are relevant for each answer. You can **either** mark the corresponding bytes directly in the figure **or** list the locations of the corresponding bytes using [...].

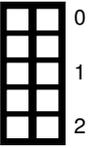
Example: the **three** bytes from position 0 to 2 can be written as [0, 2] = 0x3cec ef. Note, counting starts at 0 and start and end are included. Hexadecimal notation is also allowed: e.g., [0x10, 0x12] = 0x024052.

0  a) An Ethernet Frame consists of the header, the payload, and the *Frame Check Sequence* (FCS). What is the purpose of the FCS and how is it computed?

0  b)\* Mark the **Destination MAC Address** and note it in its common notation.

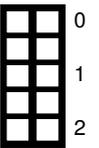
0  c)\* Identify the used **Layer 3** protocol. (Do not forget to mark and name relevant fields.)

d) Mark the **Destination IP Address** and note it in its common notation.



In Figure 3.1 different protocol headers **H1**, **H2**, **H3**, and **H4** are marked with dashed lines. In the following you will identify which protocols they belong to.

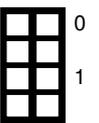
e) Identify the protocol of **H1**. Parse the header and explain the purpose of this message.



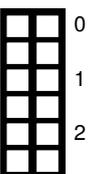
f) Based on your result from e), identify which type of protocol **H2** belongs to.



g) Identify the type of protocol of the header **H3**.



h) Argue which **protocol and application** is transported in **H4**.



### Problem 4 Network Calculus (14 credits)

This problem investigates performance bounds in networks using Network Calculus.

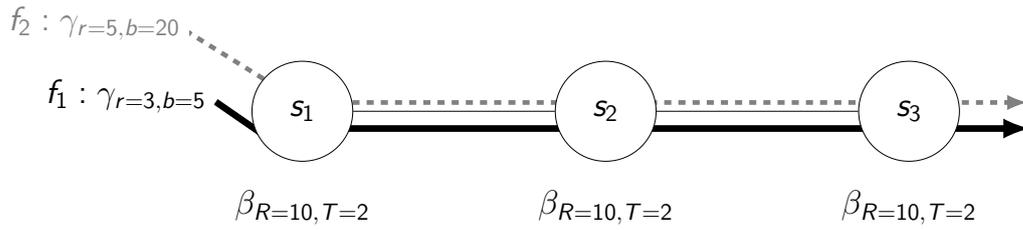


Figure 4.1: Topology with server- and flow specifications

Consider the topology in Figure 4.1. Assume each server employs strict priority queuing. Flow  $f_1$  has the lowest priority while Flow  $f_2$  has the highest priority.

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  $\alpha^* = \gamma_r, b+r \cdot T$

0	
1	
2	
3	
4	
5	

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

0	
1	

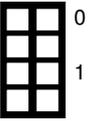
b) Perform the second step of the Separate Flow Analysis.

0	
1	

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

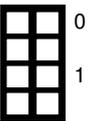
d) Assume the following changes to the scenario in Figure 4.1:

- Flow  $f_3$  joins the network with the same path and priority as Flow  $f_2$
- Flow  $f_3$  has an arrival curve with  $r_3 = 3$  and  $b_3 = 1$

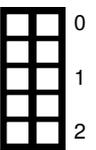


Argue how the delay bound of Flow  $f_1$  changes. Be specific.

e)\* A flow with arrival curve  $\gamma_{r=5,b=10}$  is traversing a server with service curve  $\beta_{R=20,T=2}$ . Calculate the **backlog bound** of the flow.



f)\* A flow with arrival curve  $\gamma_{r=5,b=10}$  is traversing 1,000 servers connected in series, each with the same service curve  $\beta_{R=20,T=2}$ . Calculate the **delay bound** of the flow. The method you choose should calculate a tight delay bound.



g)\* What is calculated by the following formula?

- $(\alpha \otimes \beta)(0)$

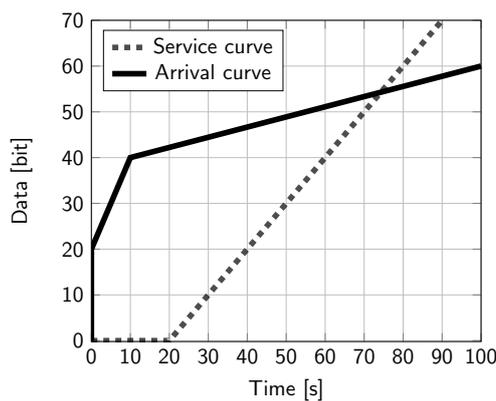


Figure 4.2: Arrival- and service curve

h)\* Consider the two curves in Figure 4.2. What is the delay bound of a flow with this **non**-token-bucket arrival curve traversing a server with this rate-latency service curve?



## Problem 5 Software-Defined Networking (13 credits)

This problem investigates a Software-Defined Network (SDN) powered by P4. For the following problems, consider the network given in Figure 5.1. Server A is configured to create and accept untagged frames, Server B only accepts VLAN-tagged frames (VLAN ID 15). PCP and DEI are always set to 0. Switch 1 is a P4 switch handling the VLAN functionality. In the following subproblems, the P4 program running on Switch 1 is investigated. Listing 1 shows parts of the used P4 program.

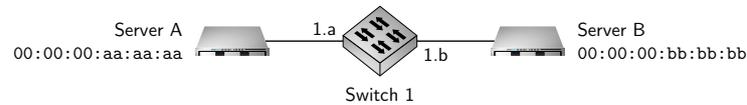
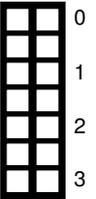


Figure 5.1: Network topology

```
1 header eth_t { bit<48> dstAddr;
3               bit<48> srcAddr;
4               bit<16> etherType; }
5 header veth_ext_t { bit<3> pcp;
6                   bit<1> dei;
7                   bit<12> vid;
8                   bit<16> etherType; }
9 struct standard_metadata_t { bit<16> ingress_spec;
10                             bit<16> egress_spec;
11                             // ...
12                            }
13 struct meta { /* unused */ }
14 struct headers { eth_t eth;
15                 veth_ext_t veth_ext; }
16
17 parser ParserImpl(packet_in packet, out headers hdr, inout meta meta, inout standard_metadata_t
18   std_meta) {
19   // to be defined in Subproblem a)
20 }
21
22 control Pipeline(inout headers hdr, inout meta meta, inout standard_metadata_t std_meta) {
23   action drop() {
24     mark_to_drop();
25   }
26   action decap(bit<16> egress) {
27     std_meta.egress_spec = egress;
28     hdr.eth.etherType = // to be defined in Subproblem b)
29     hdr.veth_ext.setInvalid();
30   }
31   action encap(bit<16> egress, bit<3> pcp, bit<1> dei, bit<12> vid) {
32     std_meta.egress_spec = egress;
33     hdr.veth_ext.etherType = // to be defined in Subproblem c)
34     hdr.eth.etherType = // to be defined in Subproblem c)
35     hdr.veth_ext.setValid();
36     hdr.veth_ext.pcp = pcp;
37     hdr.veth_ext.dei = dei;
38     hdr.veth_ext.vid = vid;
39   }
40   table forward {
41     actions = {
42       encap;
43       decap;
44       drop;
45     }
46     key = {
47       std_meta.ingress_spec: exact;
48       hdr.eth.src: exact;
49     }
50     size = 4;
51     default_action = drop;
52   }
53   apply {
54     if (hdr.eth.isValid()) {
55       forward.apply();
56     }
57 }
58
59 \\ ...
```

Listing 1: VLAN P4 program

a)\* Visualize the parse graph of Listing 1 as state machine. The parser graph starts at a start state and must be able to accept tagged and untagged Ethernet frames. Annotate the nodes with the according names and the non-trivial edges with the matches performed for this state transition.

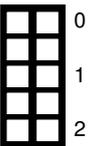


b)\* Complete the decap() action of Listing 1 (Line 26).



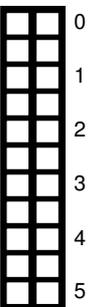
```
hdr.eth.etherType =
```

c)\* Complete the encap() action of Listing 1 (Line 31).

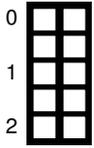


```
hdr.veth_ext.etherType =
hdr.eth.etherType =
```

d)\* The P4 program cannot work correctly without table data containing correct forwarding rules. Give the rules for Switch 1, to correctly encapsulate and forward frames of Clients A and B. Use the information given in Figure 5.1.



Match field(s)	Key	Action	Action data



e)\* The P4 program in Listing 1 is incomplete, an important control block is missing. Name the control block and briefly explain its task.

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