Chair of Network Architectures and Services
Department of Informatics
Technical University of Munich

Eexam
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 12^{th} February, 2020
Examiner:	Prof. Dr.-Ing. Georg Carle	**Time:**	10:30 – 11:45

	P 1	P 2	P 3	P 4	P 5	P 6	P 7
I							
II							

Working instructions

- This exam consists of **16 pages** with a total of **7 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 (8 credits)

0
1
a)* A **router** achieves an average throughput of 1 million packets per second for traffic with a packet size of 64 B. The same router achieves almost the same average throughput of 1 million packets per second for traffic with a packet size of 256 B. Explain why the packet size has almost no influence on the throughput.

0
1
b)* A **VPN gateway** achieves an average throughput of 1 million packets per second for traffic with a packet size of 64 B. The same gateway achieves only a throughput of 0.25 million packets per second for traffic with a packet size of 256 B. Explain why the packet size has this significant impact on throughput.

0
1
c)* What is the shortest possible representation of the given IPv6 address?

`2a01:00b0:0000:0000:0000:0a02:0000:2a0f`

0
1
d)* Name and shortly explain the main advantage of QNAME Minimization in DNS.

0
1
e)* Name and shortly explain two disadvantages of QNAME Minimization in DNS.

f)* Assume you want to register a domain. Explain the effect a parent zone has on your domain's security.

0
1

g)* An operator for data centers wants to connect two of its data centers using a VXLAN tunnel using an Internet connection from its Internet Service Provider. What is the problem if you want to transmit sensible customer data over the VXLAN tunnel?

0
1

h)* Briefly explain the purpose of TCP BBR's Drain phase.

0
1

Problem 2 Programmable Packet Processing (10.5 credits)

Computer networks currently shift from single-purpose, fixed-function network devices towards flexible and programmable network devices. This problem investigates the concepts behind these modern network devices.

The researcher Prof. Cleanslate wants to introduce a new **application layer protocol** FancyP. He wants to use programmable network devices to increment a counter contained in the header of FancyP on certain network devices.

0 1

a)* Prof. Cleanslate has 10 hosts and wants to build an OpenFlow-enabled network. Which components does he need to add to build such a network.

0 1

b) Explain the tasks of the entities listed in Problem a).

0 1 2

c)* Explain if FancyP **can or cannot** be realized using OpenFlow.

Prof. Cleanslate recently heard about a new concept called Network Function Virtualization (NFV). Now he is interested in realizing FancyP with NFV.

0 1 2

d)* Explain if FancyP **can or cannot** be realized using NFV.

Prof. Cleanslate knows from a colleague that FancyP can be realized using P4.

0 1 2

e)* Name four different P4 targets.

f) There are three potential implementations of FancyP (OpenFlow, NFV & P4). **Rank** the different implementations regarding their forwarding latency and **explain** your ranking.

0
1
2

Problem 3 P4 Forwarding (13 credits)

This problem investigates a Software-Defined Network (SDN) powered by P4. The source code of a P4 switch program is given in Listing 1.

```
header eth_t        { bit<48> dstAddr;
                      bit<48> srcAddr;
                      bit<16> etherType; }
header ipv4_t       { bit<4>  version;
                      bit<4>  ihl;
                      bit<8>  tos;
                      bit<16> totalLen;
                      bit<16> identification;
                      bit<3>  flags;
                      bit<13> fragOffset;
                      bit<8>  ttl;
                      bit<8>  protocol;
                      bit<16> hdrChecksum;
                      bit<32> srcAddr;
                      bit<32> dstAddr; }
struct meta         { /* unused */ }
struct headers      { eth_t   eth;
                      ipv4_t  ipv4; }

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) {
      ____________: parse_ipv4; // ************ see Problem a)
      default: accept;
    }
  }
  state parse_ipv4 {
    packet.extract(hdr.ipv4);
    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 ipv4_fwd(bit<16> egress) {
    std_meta.egress_port = egress;
  }
  table forward {
    actions = {
      ipv4_fwd;
      drop;
    }
    key = {
      std_meta.ingress_port: exact;
      hdr.ipv4.srcAddr: exact;
      hdr.ipv4.dstAddr: exact;
    }
    size = 2;
    default_action = drop();
  }
  apply {
    if (hdr.ipv4.isValid()) {
      forward.apply();
    }
  }
}

control DeparserImpl(packet_out packet, in headers hdr) {
  // **************************************************** see Problem f)
}
```

Listing 1: Simple P4 program

For the following problems use the network topology given in Figure 3.1. Switch 1 is a P4 switch running the P4 program of Listing 1.

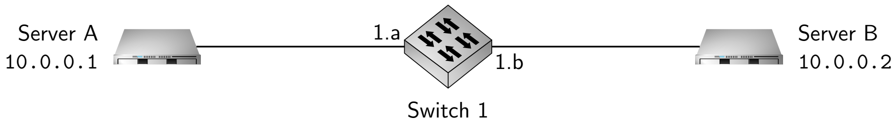

Figure 3.1: Network topology

a)* The parser of Listing 1 (**Line 24**) is incomplete. Fill in the missing value in the underlined area to get a correctly working parser.

Line 24: `_____: parse_ipv4;`

b)* The P4 program cannot work correctly without table data containing correct forwarding rules. Give the rules for Switch 1 so Servers A and B can communicate with each other via IPv4. Frames not originating from Servers A or B should be dropped. Use the information given in Figure 3.1. You can assume that both servers know the MAC address of their communication partner respectively.

Match field(s)	Key	Action	Action data

c)* The table definition of the P4 program in Listing 1 (**Line 55**) contains a `default_action = drop()`. What conditions must a packet fulfill to be dropped by this default action?

An ARP request arrives at port 1.a of Switch 1.

0 1 2 3

d)* Explain what happens to the ARP request in the parser and both control blocks.

0 1

e)* The P4 program in Listing 1 uses the `exact` match type. Name two other match types supported by P4.

0 1 2

f)* Create a valid deparser for the P4 program in Listing 1 (**Line 65**).

Problem 4 **Network Calculus (4.5 credits)**

This problem investigates the derivation of performance guarantees for a given network using Network Calculus.

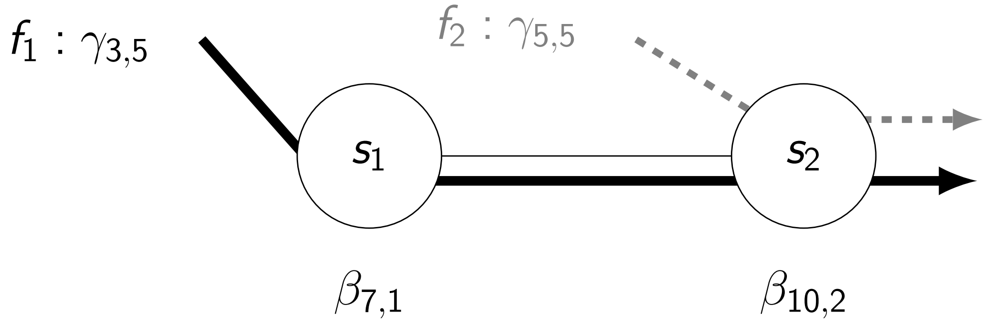

Figure 4.1: Network topology and flow description

Consider the network topology and flow definitions given in Figure 4.1. Flow f_1 traverses the Servers s_1 and s_2. Flow f_2 traverses the Server s_2.
The flows are defined as token-bucket arrival curves $\gamma_{r,b}$ with Rate r and Burst b.
The servers are defined as rate-latency service curves $\beta_{R,T}$ with Rate R and Latency T.
Assume preemptive static priority scheduling at both servers. Furthermore, assume Flow f_1 has a low priority and Flow f_2 has a high priority.

a)* Calculate the residual service curve for Flow f_1 at Server s_1. Specify your end result in the form $\beta_{R,T}$.

0
1

b)* Calculate the residual service curve for Flow f_1 at Server s_2. Specify your end result in the form $\beta_{R,T}$.

0
1

c) Use the concatenation theorem to combine the two residual service curves into a single end-to-end service curve for f_1. Specify your end result in the form $\beta_{R,T}$.

0
1

d) Calculate the delay bound for Flow f_1 traversing the Servers s_1 and s_2. Make use of the concatenation theorem and consider the scheduling strategy as well as the influence of other flows.
Hint: Re-use results from previous sub-problems.

0
1

Problem 5 Transport Layer (12.5 credits)

This problem is about transport layer protocols. The Transmission Control Protocol (TCP) is widely used for its reliability property. Lost segments can be detected and are retransmitted by the sender. Beside the reliable data transfer TCP also offers *Flow Control* and *Congestion Control*.

0
1
a)* How can a TCP **receiver** detect lost segments in the byte stream?

0
1
b)* Briefly explain the *Fast Retransmit* algorithm for TCP.

0
1
c)* Explain one situation when *Fast Retransmit* is **not** triggered.

0
1
d)* What is the goal of *Flow Control*?

0
1
e)* How is the concept of *Flow Control* implemented in TCP?

0
1
f)* Name **three** different classes of congestion control algorithms.

0
1
g)* Name one advantage and one disadvantage of TCP Vegas.

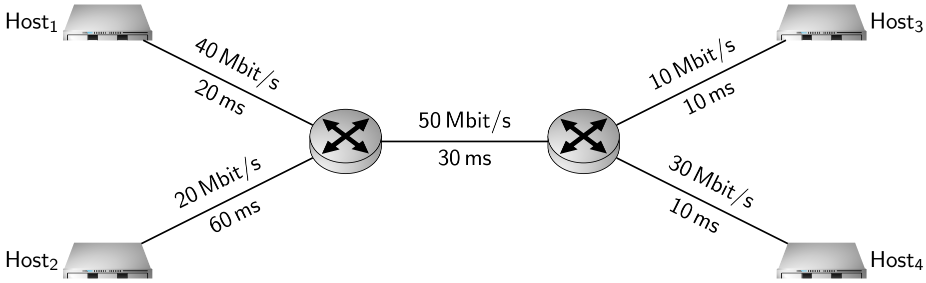

Figure 5.1: The maximum transmission rate and propagation delay is noted for each link.

In the following the network shown in Figure 5.1 is used. There, two TCP flows are started:

- $Flow_1$ sends data from $Host_1$ to $Host_3$
- $Flow_2$ sends data from $Host_2$ to $Host_4$.

$Flow_1$ ends before $Flow_2$ starts so you can assume that they do not influence each other.

h)* Compute the *Bandwidth Delay Product* (BDP) for $Flow_1$ and $Flow_2$ in kbit.

0 1 2

i)* How is the BDP related to congestion control?

0 1

In HTTP/2 all files are sent multiplexed over one TCP connection. This can result in Head-of-Line (HoL) Blocking.

j)* Briefly explain how TCP causes HoL Blocking on the transport layer in this case.

0 1

k)* How does the QUIC protocol solve this issue?

0 1

Problem 6 Wireshark (10.5 credits)

According to the OSI model network protocols are distributed to seven different layers each containing several protocols. In this problem a packet is analyzed referring to the involved protocols.

```
0x0000  ac 1f 6b 7c 81 68 00 25   90 57 22 4a 08 00 45 00

0x0010  00 3c fa 3a 00 00 79 06   a5 91 08 08 08 08 83 9f

0x0020  0e 41 03 55 cb 7c 11 45   ab d7 bb 0f f3 be a0 12

0x0030  eb 20 1d 5c 00 00 02 04   05 64 04 02 08 0a ba 8c

0x0040  d2 b1 20 bd b5 19 01 03   03 08 01 02 03 04
```

Figure 6.1: Hexdump of a complete Ethernet frame including FCS

a)* Mark and name all parts of the protocol specific information for layer 2 in Figure 6.1 **Note:** Put your solution directly in Figure 6.1

In the next four subproblems you are asked to identify which protocols were used for each layer. For each question do the following:

- mark the corresponding bytes in the hexdump and
- write the corresponding bytes in the solutionbox.

b)* Name the L3 protocol.

c) Name the L4 protocol.

d) List the flags which are set in the L4 protocol.

e) Identify the application layer protocol.

0
1
2

f) Determine the length of the L4 header in byte.

0
1

g) Determine the size of the L4 payload in byte.

0
1

Problem 7 Internet Measurements (16 credits)

Internet-wide measurements are a major research area in the field of networking. This problem investigates properties and important considerations regarding active network scans.

0
1

a)* Name two ethical considerations relevant for Internet wide measurements.

0
1

b)* Explain why non BGP announced Prefixes should not be scanned.

0
1
2
3

c)* Given we want to scan the routable addresses in 138.0.0.0/8, how many targets are we going to scan with the prefixes in following table having an entry in the routing table?

Routed Prefixes
138. 64. 0.0/11
8. 0. 0.0/12
138. 0. 0.0/12
130.120. 0.0/16
138.132. 0.0/16
138. 30. 0.0/20
138.138. 0.0/20
138. 0. 2.0/22
138.138.12.0/22

Hint: The table might contain more specific announcements.

Hint: The answer does not need to compute the actual number of targets, the computation including power of two terms is also valid. Your solution approach should be comprehensible.

0
1

d)* Assume that we now want to scan a /**8** subnet in IPv6. Explain why that is not feasible.

0
1

e)* What Resource Record (RR) Type has to be queried to resolve domain lists to IPv4 and IPv6 addresses respectively.

f)* Name and briefly explain two effects discussed during the lecture observed in domain top lists.

0 1 2

g)* Name two further input sources for active Internet wide scans besides domain lists

0 1

Now consider an ICMP echo request scanner that sends minimal echo requests to each target.

h)* Calculate the total **frame** size of frames send by the ICMP scanner for IPv4 and IPv6.

0 1 2

Now, assume the scanner sends IPv4 ICMP packets with a frame size of 100 B and IPv6 ICMP packets with a frame size of 250 B. The scanner is configured with two rate limits. The maximum number of sent packets per second is 10 000 pps. The scanner sends a maximum L2 data rate of 10 Mbit/s.

i)* How many IPv4 targets can be scanned in 1 s and how much data will be sent in 1 s with the given configuration?

0 1 2

j)* How many IPv6 targets can be scanned in 1 s and how much data will be sent in 1 s with the given configuration?

0 1 2

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