

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: Friday 1st March, 2019

Examiner: Prof. Dr.-Ing. Georg Carle

Time: 10:30 – 11:30

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

Working instructions

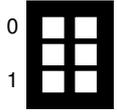
- This exam consists of
 - **16 pages** with a total of **5 problems** and
 - a two-sided printed **cheat sheet**.

Please make sure now that you received a complete copy of the exam.

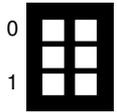
- Detaching pages from the exam is prohibited.
- Subproblems marked by * can be solved without results of previous subproblems.
- **Answers are only accepted if the solution approach is documented.** Give a reason for each answer unless explicitly stated otherwise in the respective subproblem.
- Do not write with red or green colors nor use pencils.
- The total amount of achievable credits in this exam is 60 credits.
- Allowed resources:
 - one **analog dictionary** English ↔ native language
- 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)



a)* Consider HTTP load balancers. Explain the interaction between load balancer and client.



b)* Figure 1.1 shows the arrival curve α and the service curve β at a network node. Clearly mark the delay bound and the backlog bound directly in Figure 1.1.

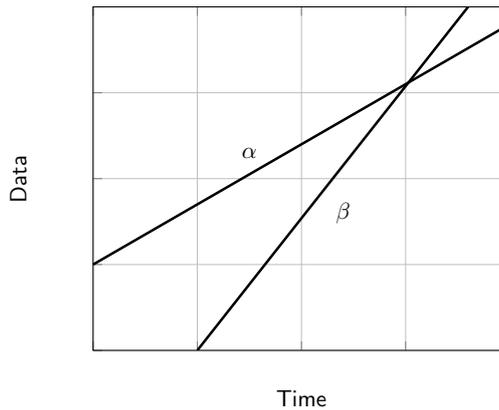
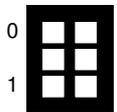
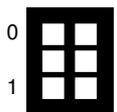


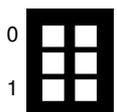
Figure 1.1: Arrival and service curve at a network node.



c)* Name the protocol used by Internet Protocol Security (IPsec) to establish Security Associations (SA) during the handshake.

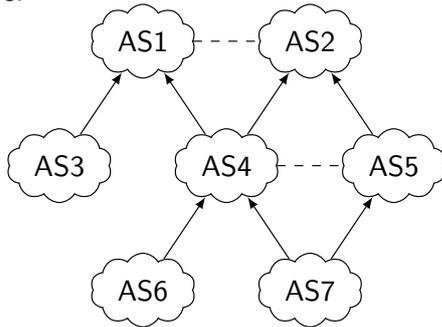
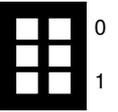


d)* Explain the concept of path prepending in BGP.



e)* Explain the difference between the transmission (or serialization) delay and the propagation delay of a frame.

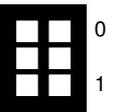
f)* Figure 1.2 shows an AS network. The arrows represent Customer→Provider relations, the dashed lines represent Peering agreements. Which ASes are Stub- and/or Multi-Homed AS? Fill the given table with the correct AS numbers.



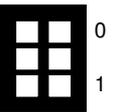
Tier-1 provider:	AS1, AS2
Stub-AS:	
Multi-Homed AS:	

Figure 1.2: AS Topology

g)* You try to build up a QUIC connection to google.com using your favorite browser, the Internet Explorer. Does this work? Give a brief explanation.



h) Table 1.1 shows two IPFIX flows. Assume the IP 5-tuple is used as flow identifier. Briefly explain whether or not the two flows can be merged using bidirectional flow matching with an unlimited active timeout and an inactive timeout of 30 seconds.



Start	End	Src IP	Dst IP	Protocol	Src Port	Dst Port	Bytes	Packets
t+0s	t+20s	10.10.10.10	8.8.8.8	TCP	2222	443	750	3
t+5s	t+32s	8.8.8.8	10.10.10.10	TCP	2222	443	4500	3

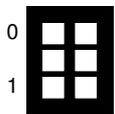
Table 1.1: Two IPFIX flows

Problem 2 Wireshark (11 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  00 25 90 57 22 4a 58 ef    68 e7 48 66 86 dd 60 04
0x0010  80 0f 00 24 11 40 20 01    4c a0 20 01 00 00 00 00
0x0020  00 00 10 01 20 02 2a 00    14 50 40 01 00 00 00 00
0x0030  00 00 00 00 20 0a e2 89    01 bb 00 24 38 8e 0c 47
0x0040  f5 ec e4 5b 3a c7 49 05    a5 40 1e 1e 5d e3 d6 22
0x0050  37 e6 12 55 dd 0c 3d 19    a3 21 de ad be ef
```

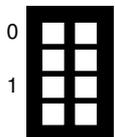
Figure 2.1: Hexdump of a complete Ethernet frame including FCS



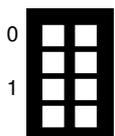
a)* Mark the FCS in the hexdump. What is the purpose of the FCS?

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

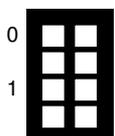
- mark the corresponding bytes in the hexdump
- write the corresponding bytes in the solutionbox
- name the used protocol.



b)* Identify the L3 protocol.



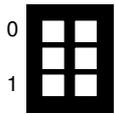
c) Identify the L4 protocol.



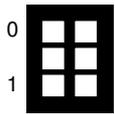
d) Name the L4 destination port in decimal notation. Derive the type of the L4 payload from it.

Problem 3 TCP (14 credits)

The Transmission Control Protocol is used by many different network applications. In this problem properties of the protocol are analyzed.

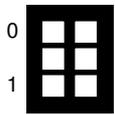


a)* Name two ways how TCP can detect lost segments?



b)* The following sequence numbers arrive at an endpoint of a TCP connection. For each segment the receiver replies with an acknowledgment. Complete the missing ACK numbers. You can assume equal sized segments, a sufficiently large receiver buffer, and no influences due to flow control.

SEQ	5	6	8	9	7	10
ACK	6	7				



c)* Figure 3.1 shows how the delivery rate of a TCP connection changes with increasing data inflight in an ideal environment. Update Figure 3.2 accordingly.

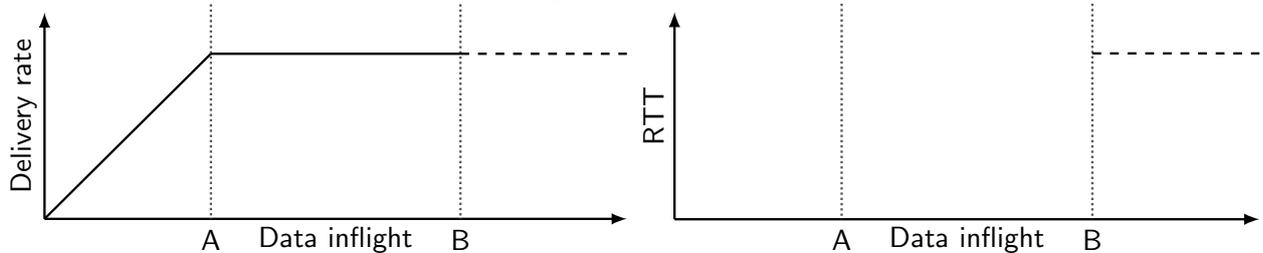
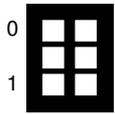


Figure 3.1

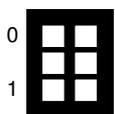
Figure 3.2



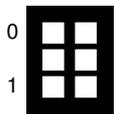
d)* How large is the amount of inflight data at the labels A and B in Figure 3.1?

A:

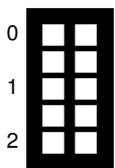
B:



e)* Mark and label the operation point of loss-based and delay-based congestion control algorithms in Figure 3.1.



f)* Briefly explain which amount of inflight data is optimal for TCP.



g)* Briefly explain what happens during TCP BBR's Probe BW and Probe RTT phases.

Problem 4 Traceroute (13.5 credits)

This problem investigates the traceroute tool.

0
1
2
3

a)* Figure 4.1 shows a network topology. *C* executes a traceroute with destination *S*. The different stages of one of the packets and its reply traversing the network are shown as tables. Fill in the missing fields directly in Figure 4.1. The Protocol field should be the highest layer protocol of the packet. You can abbreviate Src MAC, Dst MAC, Src IP, and Dst IP as *device.interface*, for example *R₀.eth0*. Make meaningful choices for the values of the TTL and Protocol fields.

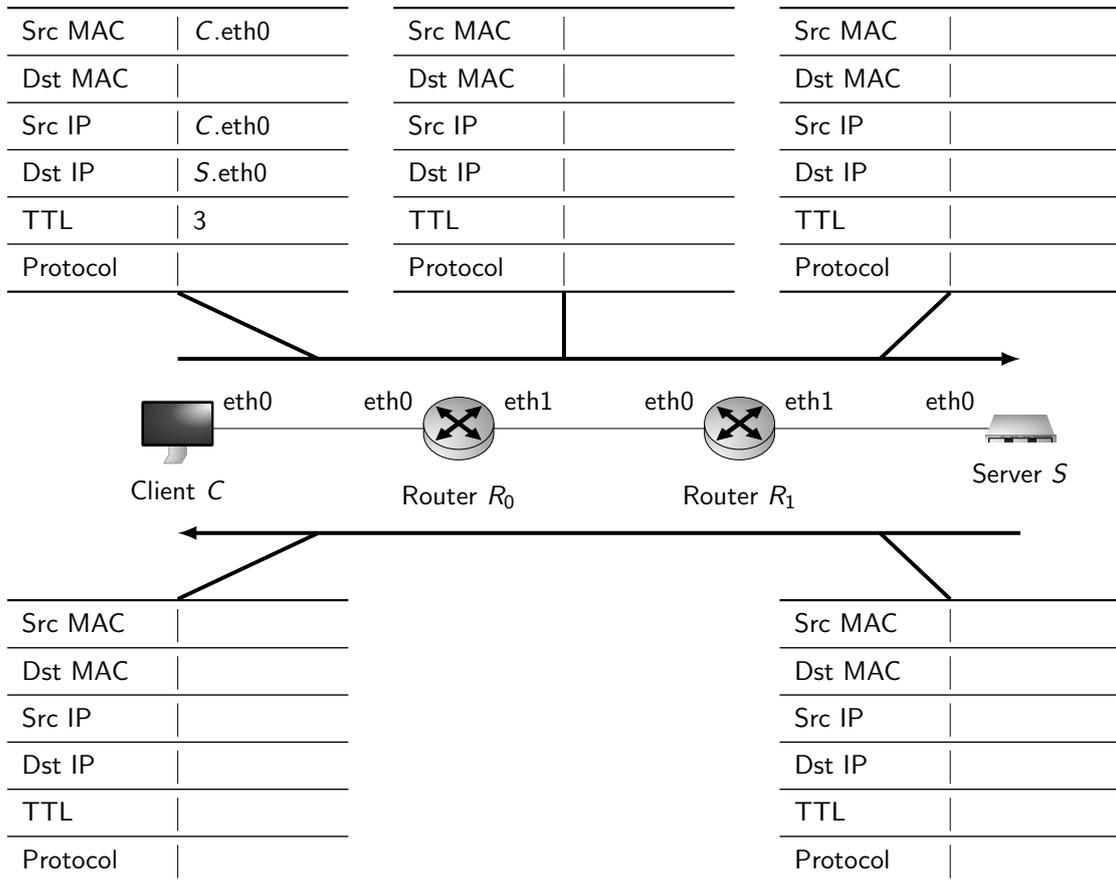


Figure 4.1: Network topology with packets between hops

Figure 4.2 shows a network topology. *C* executes a UDP-based traceroute with destination *S*. Figure 4.3 shows a hexdump of an ICMPv4 packet (starting with the first byte of the ICMPv4 header), which is one of the packets received by *C* as a result of the executed traceroute.

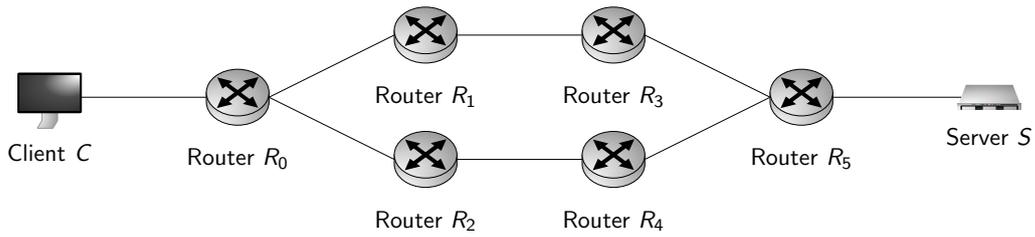


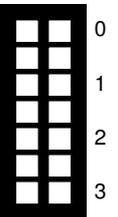
Figure 4.2: Network topology

```

0x0000  03 03 32 68 00 00 00 00    45 80 00 3c b2 f3 00 00
0x0010  01 11 c3 11 83 9f 14 2f    d8 3a d0 23 84 a8 82 bc
0x0020  00 28 c3 07
  
```

Figure 4.3: Hexdump of an ICMPv4 packet received at *C*

b)* Use the hexdump in Figure 4.3 to explain in detail which node from Figure 4.2 most likely originated the ICMPv4 packet and why.



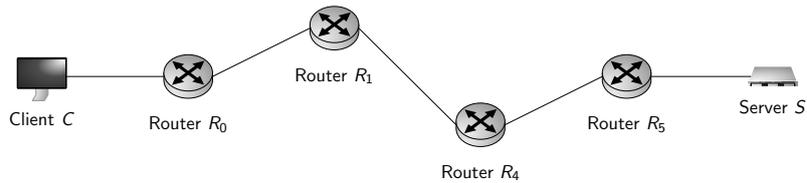
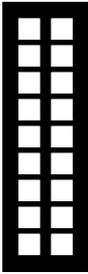
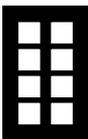
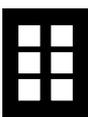


Figure 4.4: Erroneously detected network topology.

0  c) Explain a scenario where a traceroute on the network topology in Figure 4.2 erroneously detects it to look like in Figure 4.4.

0  d)* Briefly explain how Paris-Traceroute works and what advantage it has over traditional traceroute.

For the following subproblems, assume R_5 in Figure 4.2 blocks all types of ICMPv4 traffic and does not generate any type of ICMPv4 traffic itself. Assume traceroute sends exactly one packet for each TTL, and each hop introduces a **one way** latency of 10 ms.

0  e)* C executes an ICMPv4 traceroute with target S . Table 4.1 shows the output of the traceroute. Fill in the missing Round Trip Time (RTT) values.

Hop	IP address	RTT
1	$R_0.ip$	
2	$R_1.ip$	
3	$R_3.ip$	
4	$R_5.ip$	
5	$S.ip$	

Table 4.1: Output of ICMPv4 traceroute

0  f)* C executes a UDP traceroute with target S . Table 4.2 shows the output of the traceroute. Fill in the missing Round Trip Time (RTT) values.

Hop	IP address	RTT
1	$R_0.ip$	
2	$R_1.ip$	
3	$R_3.ip$	
4	$R_5.ip$	
5	$S.ip$	

Table 4.2: Output of UDP traceroute

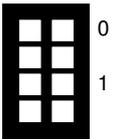
Problem 5 Software-Defined Networking (13.5 credits)

This problem investigates a Software-Defined Network (SDN) powered by P4.

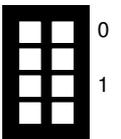
Plane	Description

Table 5.1: Planes of a Software-Defined Network

a)* List the three planes used in SDN in Table 5.1.



b) Briefly explain one of the main tasks of each plane in Table 5.1.



For the following problems consider the network given in Figure 5.1. 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. PCP and DEI are always set to 0.

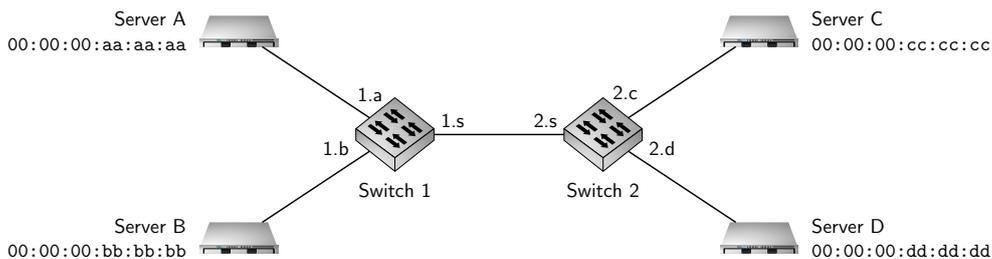
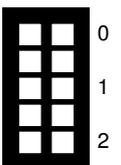


Figure 5.1: Network topology

c)* A VLAN capable switch has two different types of ports. Name the two different types and classify all ports of the network in Figure 5.1 (e.g. 1.a) in those two classes.



Port type	List of ports

Switches 1 and 2 are P4 switches handling the VLAN functionality. In the following the P4 program running on both switches is investigated. Listing 1 shows a small part of the used P4 program.

```

header eth_t      { bit<48> dstAddr;
                  bit<48> srcAddr;
                  bit<16> etherType; }
header veth_ext_t { bit<3> pcp;
                  bit<1> dei;
                  bit<12> vid;
                  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) {
  // to be programmed
}

control Pipeline(inout headers hdr, inout metadata meta, inout standard_metadata_t std_meta) {
  action drop() {
    mark_to_drop();
  }
  action encap(bit<16> egress, bit<3> pcp, bit<1> dei, bit<12> vid) {
    std_meta.egress_port = egress;
    hdr.veth_ext.setValid();
    hdr.veth_ext.pcp = pcp;
    hdr.veth_ext.dei = dei;
    hdr.veth_ext.vid = vid;
  }
  table forward {
    actions = {
      encap;
      drop;
    }
    key = {
      std_meta.ingress_port: exact;
      hdr.eth.src: exact;
    }
    size = 4;
  }
  apply {
    if (hdr.veth_ext.isInvalid()) {
      forward.apply();
    }
  }
  // decapsulation functionality omitted
}

```

Listing 1: VLAN P4 program

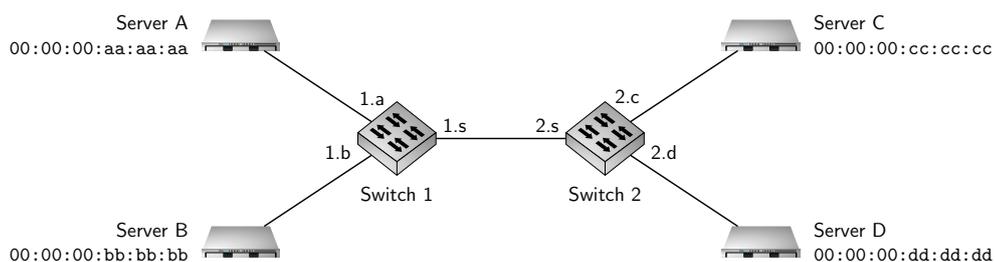
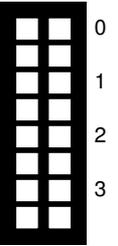
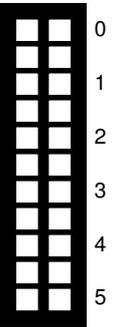


Figure 5.2: Network topology (Copy from previous page)

d)* Create a parser for the program given in Listing 1. The parser should be able to correctly parse VLAN and non-VLAN frames.



e)* 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. Packets not originating from either A or B should be dropped. Use the information given in Figure 5.2. Use * to mark the table cells which match on any value for which no more specific entry exists in this table.



Match field(s)	Key	Action	Action data

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

