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 / Retake	**Date:**	Thursday 13th April, 2023
Examiner:	Prof. Dr.-Ing. Georg Carle	**Time:**	11:30 – 12:45

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:
  - 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 (14 credits)

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

a)* Given an IPv4 IHL of 5, how long is the header?

- ☐ 5 B
- ☐ 10 B
- ☒ 20 B
- ☐ 40 B

b)* How long is an IPv6 header?

- ☐ 20 bit
- ☐ 40 bit
- ☐ 160 bit
- ☒ 320 bit

c)* Which flags are set during an Xmas TCP scan?

- ☐ FIN+PUSH
- ☒ FIN+PUSH+URG
- ☐ PUSH+URG
- ☐ FIN

d)* Finish the following sentence: A QUIC packet is

- ☐ a UDP PDU
- ☒ a UDP SDU
- ☐ an IP PDU
- ☐ an IP SDU

e)* Which HTTP status code would you expect to receive from an HTTP-based load balancer?

- ☐ 503
- ☐ 200
- ☐ 404
- ☒ 302

f)* If the DNS message contains authoritative data, which part of it contains the authoritative information?

- ☐ Additional
- ☐ Header
- ☐ Authority
- ☒ Answer

0 1 g)* What does CIDR stand for and what is it used for?

CIDR = Classless InterDomain Routing
Divides IP address space into subnets of arbitrary length

0 1 h)* Shortly explain **what** OUI stands for and **where** it can be found in the context of this course.

OUI stands for Organisation Unique Identifier which is a fixed ID assigned to a hardware manufacturer and used in the first 3 Bytes of a MAC address.

i)* Which DNS use case often breaks the boundary of 512 B UDP payload size?

0 1

- DNSSEC as its records contain signatures which can be large especially with RSA

j)* Consider an authoritative name server of a CDN. Why does an A record of a CDN have a TTL of 300 or less while for a personal website hosted on an instance in the cloud it is suggested to have at least an hour or longer TTLs?

0 1 2

- CDNs typically provide load-balancing. In order to perform a DNS based load balancing small TTLs are important to shift traffic between addresses in a timely manner
- Personal websites which hosted in the cloud are typically on a fixed IP address and therefore they rarely change and don't need timely changes. The goal is to reduce load on the name server with higher TTLs

k)* SCTP is a transport layer protocol combining the benefits of TCP and UDP. Give two reasons why SCTP is **not** widely deployed today.

0 1 2

No support on middleboxes (firewall, NAT) and hardly for end hosts (operating system)

```
1 user@foo:~$ dig +short netflix.com
2 54.73.148.110
3 54.155.246.232
4 18.200.8.190
```

Listing 1: Command and its output

l)* Name (without explanation) the reason why the command in Listing 1 returns more than one IP address.

0 ½

(DNS-based) load balancing

Problem 2 Hexdump (7.5 credits)

This problem investigates a captured Ethernet frame. The given hexdump starts with the Ethernet header in Figure 2.1.

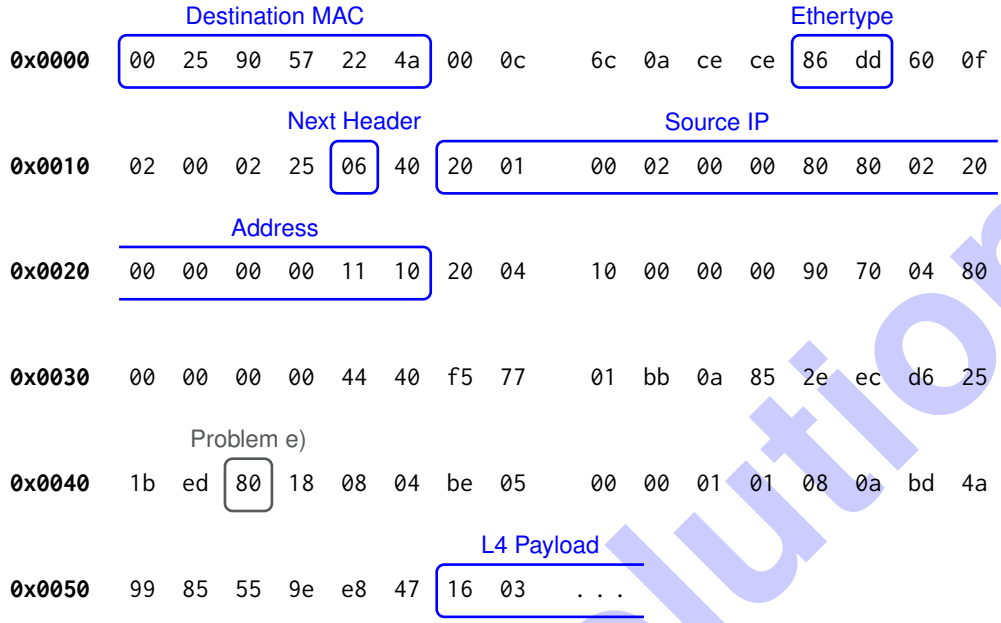

Figure 2.1: Hexdump of an Ethernet frame starting with the Ethernet header. Only the first 88 B are displayed, the rest is omitted.

In this problem, you **always** have to substantiate your answers using the bytes of the hexdump in Figure 2.1. Always make clear which bytes are relevant for each answer.

0 1 a)* Mark the **Destination MAC address** and note it in its common notation.

Destination MAC: `[0, 5]= 0x00259057`
`00:25:90:57:22:4a`

0 1 b)* Identify the used **Layer 3** protocol. (Do not forget to mark and name relevant fields.)

Ethertype: `[12, 13]= 0x86dd` ⇒ IPv6

0 1 2 c) Mark the **Source IP address** and note it in its common and shortest possible notation.

Source IP address: `[22, 37]= 0x20010002000080800220000000001110`
`2001:0002:0000:8080:0220:0000:0000:1110`
→ `2001:2:0:8080:220::1110`

0 1 d) Identify the used **Layer 4** protocol.

Next Header: `[20]= 0x06` ⇒ TCP

The field at position [66] in the hexdump (see Figure 2.1) is described as follows in RFC 793:

> The number of 32 bit words in the Header. This indicates where the data begins.

e) Using this information, determine the start of the **Layer 4 Payload** and mark it in the hexdump.

0
1
2

TCP offset: 0x8 $\Rightarrow$ 8 · 4 B = 32 B header length
Start of TCP header: [54]
Start of payload: [54 + 32] = [86]

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

Problem 3 Network Calculus (14 credits)

This problem investigates performance bounds in networks using Network Calculus.

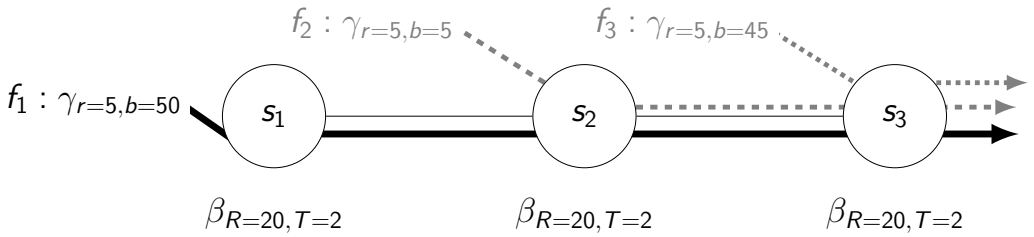

Figure 3.1: Topology with server- and flow specifications

Consider the topology in Figure 3.1. Assume each server employs strict priority queuing. Flow f_1 has the lowest priority while Flows f_2 and f_3 have the highest priority.
We are interested in calculating the end-to-end delay bound of **Flow f_1** using the Separate Flow Analysis.

0 1 2 3 4

a)* Perform the first step of the Separate Flow Analysis (calculating the left-over service curves).

- Left-over service curve at s_1:

$$\beta_{s_1}^{l.o.} = \beta_{s_1} = \beta_{20,2}$$

- Left-over service curve at s_2:

$$\beta_{s_2}^{l.o.} = \left[\beta_{s_2} - \alpha_{f_2}\right]^+ = \left[\beta_{20,2} - \gamma_{5,5}\right]^+ = \beta_{20-5,\frac{5+20\cdot 2}{20-5}} = \beta_{15,3}$$

- Left-over service curve at s_3:

$$\beta_{s_3}^{l.o.} = \left[\beta_{s_3} - (\alpha_{f_2}^\star + \alpha_{f_3})\right]^+ = \left[\beta_{s_3} - ((\alpha_{f_2} \oslash \beta_{s_2}) + \alpha_{f_3})\right]^+ = \left[\beta_{20,2} - (\gamma_{5,5+5*2} + \gamma_{5,45})\right]^+ = \left[\beta_{20,2} - \gamma_{10,60}\right]^+ = \beta_{20-10,\frac{60+20\cdot 2}{20-10}} = \beta_{10,10}$$

0 1

b) Perform the second step of the Separate Flow Analysis (applying the concatenation theorem to compute the end-to-end left-over service curve).

$$\beta_{e2e}^{l.o.} = \beta_{s_1}^{l.o.} \otimes \beta_{s_2}^{l.o.} \otimes \beta_{s_3}^{l.o.} = \beta_{min(20,15,10),2+3+10} = \beta_{10,15}$$

0 1

c) Perform the third step of the Separate Flow Analysis (calculating the end-to-end delay bound).

$$d = T + \frac{b}{R} = 15 + \frac{50}{10} = 20$$

d) Assume the following modifications to the scenario in Figure 3.1:

- The burst of Flow f_2 is increased to 50: $\gamma_{r=5,b=50}$
- The rate of Flow f_3 is increased to 15: $\gamma_{r=15,b=45}$
- The processing latency of Server s_1 is increased to 5: $\beta_{R=20,T=5}$

Argue how this combination of changes influences the end-to-end delay bound of Flow f_1 calculated in c).

0 1

The delay bound is infinity because the left-over service curve at Server s_3 has a rate of $R = 0$.
Therefore, Flow f_1 with a rate of $r = 5$ cannot be served.
Burst and processing latency increase are not relevant.

e)* Assume a single token-bucket constrained flow with arrival curve $\gamma_{r=0.5,b=20}$ traverses a single rate-latency server with a service curve $\beta_{R=1,T=60}$. Draw the following four components into the empty plot in Figure 3.2a:

- The arrival curve
- The service curve
- The delay bound
- The backlog bound

An additional preprint is available in Figure 3.2b). Clearly mark each of the four components.

0 1 2 3

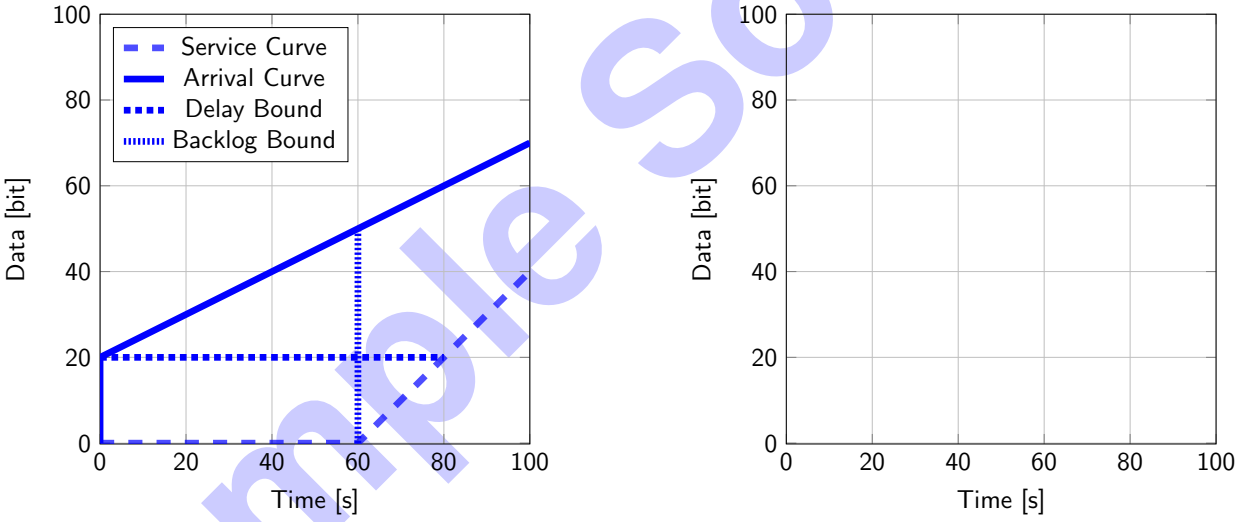

(a) Insert solution for Subproblem e) here

(b) Additional preprint

Figure 3.2: Third step of the Separate Flow Analysis

f)* Name **two** approaches that can be used to obtain hard real-time guarantees.

0 1

Any two of: Model Checking, Real-Time Calculus, Trajectory Approach, Deterministic Network Calculus, Schedulability Analysis

g)* Explain what arbitrary multiplexing is.

0 1

Arbitrary multiplexing assumes we don't have knowledge of how the mutliplexing or scheduling works.

Problem 4 AS Relations and BGP (13.5 credits)

This problem investigates the autonomous system (AS) relationships in a given network and their impact on routing and traffic. All ASes apply standard routing behavior. Furthermore, the following policies are applied:

- If routes with the same prefix exist, an AS selects the cost-efficient route.
- If routes with the same prefix exist, and the cost to route traffic is equal for all routes, the shorter route is selected.

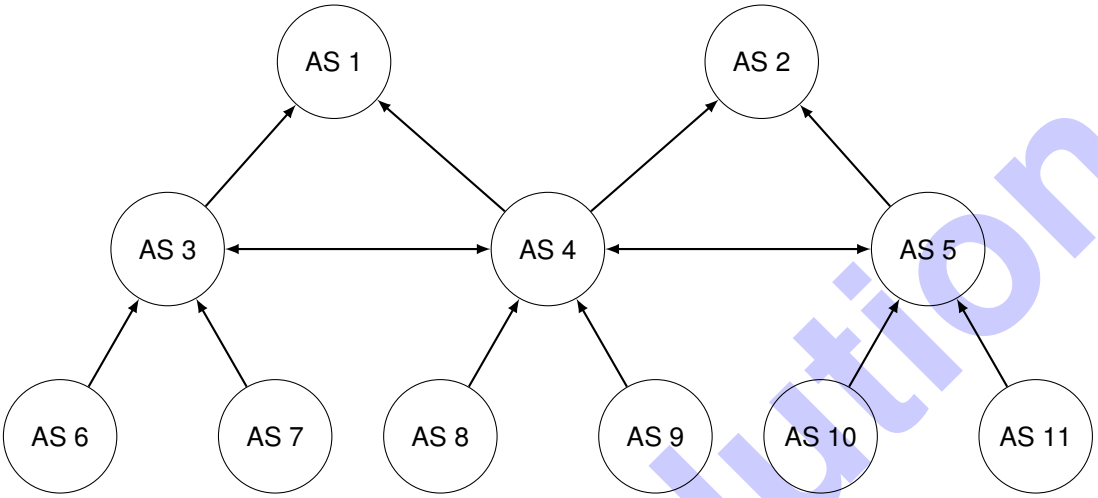

Figure 4.1: AS Network

0 1 2

a)* What does iBGP and eBGP stand for and what are the differences?

- Internal BGP: BGP exchanges information with routers in the same AS
- External BGP: BGP exchanges information with routers of neighboring ASes

0 1 2

b)* Which ASes can reach all other ASes? Explain!

- AS 4 can reach all other ASes as it is connected to both AS 1 and AS 2
- AS 8 and AS 9 are customers of AS 4 and can therefore also reach all other ASes

0 1

c)* Explain why AS 1 and AS 2 cannot be considered real Tier-1 providers. Which Tier-1 properties do they fulfill and which are missing?

- AS 1 and AS 2 only have customers, but cannot reach the complete network.

0 1

d)* Explain why AS 4 cannot be considered as a real Tier-1 provider. Which Tier-1 properties does it fulfill and which are missing?

- AS 4 can reach the complete network, but it has providers.

Assume AS 11 announces the prefix 172.0.0.0/24 for the following subproblems.

e)* How is traffic routed from AS 8 to the prefix?

0 1

- AS 8 → AS 4 → AS 5 → AS 11
- AS 5 announces the route to AS 2 and AS 4. AS 4 announces the route via AS 5 to AS 8. It gets money from AS 8

f)* How is traffic routed from AS 7 to the prefix?

0 1

- No valid route given the standard policies
- AS 4 neither announces the route via AS 2 (it has to pay) nor the route via AS 5 (only traffic, no profit) to AS 3

g)* Which neighboring ASes do offer AS 4 a path towards the announced prefix and which is the most economic neighbor for AS 4?

0 1 2

- AS 5 as AS 11 is a customer
- AS 2 as AS 5 is a customer of AS 2 and therefore announces all prefixes of its own customers to the provider AS 2. The provider offers all prefixes of its customers to all other customers
- The path via AS5 is the more economic one

Assume AS 6 additionally announces the prefix 172.0.0.0/24 for the following subproblems.

h)* Name and shortly explain two use cases from the lecture and exercise given both announcements. Note: A single organization can own multiple ASes.

0 1 2

- Prefix Hijack, a malicious actor could announce the prefix to intercept traffic
- Anycast load balancing, announcing the same prefix from different vantage points can be used for load balancing which relies on BGP to select a *good* route.

i)* Given both announcements, how is traffic routed from AS 8 towards the prefix?

0 1

- With the given policies, no unique route can be selected
- same prefix, same path length, same cost → additional policies at AS 4 are required.

Problem 5 Transmission Control Protocol (15 credits)

The Transmission Control Protocol (TCP) is one of the most used network protocols. It provides reliable, in-order data transmission on the transport layer. In this problem, you will take a closer look at the features it provides.

0 1

a)* How does a **TCP receiver** detect lost segments?

TCP segments can be uniquely identified by the consecutive sequence number
When sequence numbers are missing, the receiver knows that segments are missing.

0 1 2

b)* **Name** and **briefly explain** two reasons a TCP sender retransmits segments.

- Timeout: if the time between sending a segment and receiving an acknowledgement for it is too long, the segment gets retransmitted.
- Fast Retransmit: if three duplicate acknowledgement arrive, a lost segment is assumed and it is retransmitted.

0 1

c)* What is the purpose of flow control?

Prevention of overwhelming the receiver of a TCP stream.

0 1

d)* How is flow control implemented in Linux TCP?

Each receiver announces the amount of free buffer space (receive window) with each segment in the TCP header.

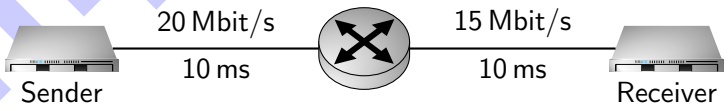

Figure 5.1: Network topology with link capacities and round-trip propagation delays.

For the following questions, consider the network topology shown in Figure 5.1. The link capacity and round-trip propagation delay are given for each link. Each network interface uses an additional 300 kbit buffer to queue packets. Consider a TCP flow transmitting data from the sender to the receiver.

0 1 2

e)* Compute the BDP for the TCP flow in **kbit**.

BDP = BtlBw · RTprop
= $min(20, 15)$ Mbit/s · (10 + 10) ms
= 15 · 20 kbit
= 300 kbit

To prevent congestion collapse, TCP implements congestion control. The amount of data inflight is, in general, limited by the congestion window.

f) In Figure 5.2a, draw how the delivery rate behaves with increasing inflight data until packets are dropped.

0 1

g) In Figure 5.2b, draw how the RTT behaves with increasing inflight data until packets are dropped.

0 1 2

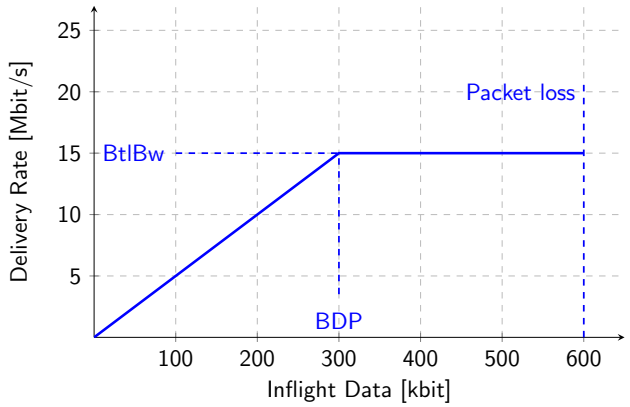

(a) Delivery rate of the TCP flow

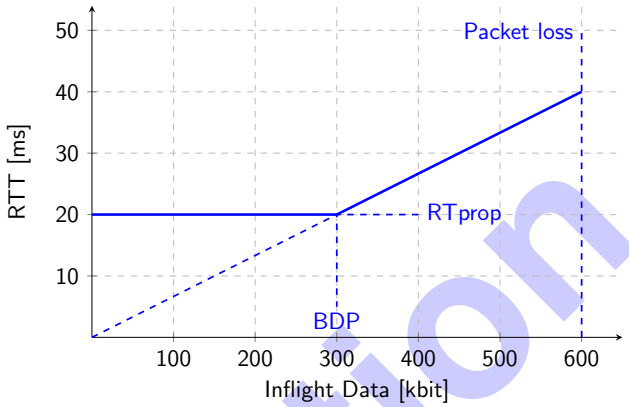

(b) RTT of the TCP flow

Figure 5.2: Model for the delivery rate and RTT of the TCP flow with increasing data inflight. (You can find additional preprints at the bottom of this page.)

h)* There exist different groups of congestion control algorithms like loss-based and delay-based. Name **two loss-based** algorithms.

0 1

Reno, Cubic

i)* Name one advantage and one disadvantage of **TCP Vegas**.

0 1

Optimal RTT and delivery rate
Starving when competing with loss-based flows

j)* TCP BBR aims to achieve the optimum for both—delivery rate and latency.
Name all different phases of BBR.

0 1 2

Startup, drain, probe bandwidth, probe rtt

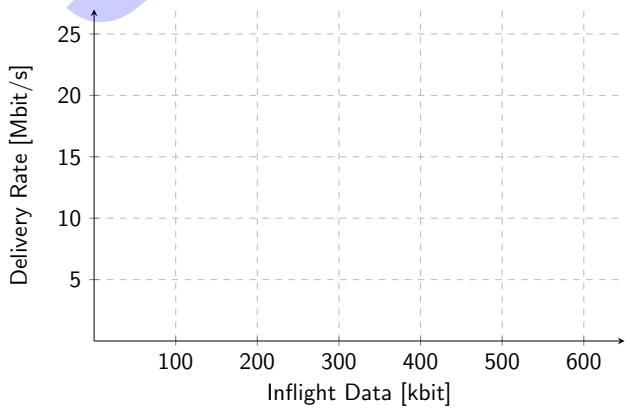

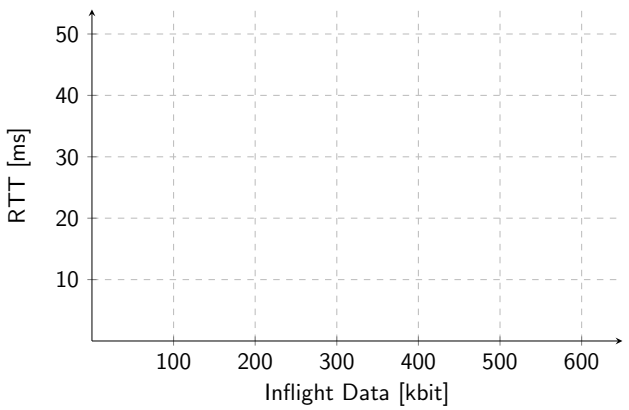

Problem 6 Software-defined Networks (11 credits)

In this problem, we investigate a software router based on P4. Listing 2 shows the P4 program of the software router.

```
1  header ethernet_t {
2      bit<48> dstMacAddr;
3      bit<48> srcMacAddr;
4      bit<16> etherType; }
5
6  header ipv4_t {
7      bit<4>  version;
8      bit<4>  ihl;
9      bit<8>  diffserv;
10     bit<16> totalLen;
11     bit<16> identification;
12     bit<3>  flags;
13     bit<13> fragOffset;
14     bit<8>  ttl;
15     bit<8>  protocol;
16     bit<16> hdrChecksum;
17     bit<32> srcAddr;
18     bit<32> dstAddr; }
19
20 struct metadata { /* unused */ }
21
22 struct headers  { eth_t  eth;
23                   ipv4_t ip4; }
24
25 parser ParserImpl(packet_in packet, out headers hdr, inout metadata meta, inout
       standard_metadata_t standard_metadata) {
26
27     state parse_ip4 { packet.extract(hdr.ip4)
28                       transition accept; }
29     state parse_eth { packet.extract(hdr.eth);
30                       transition select(hdr.eth.etherType) { 0x0800:  parse_ip4;
31                                                              default: accept; } }
32     state start     { transition parse_eth; }
33
34 }
35
36 control DeparserImpl(packet_out packet, in headers hdr) {
37     apply { packet.emit(hdr.eth);
38             packet.emit(hdr.ip4); }
39 }
40
41 control Pipeline(inout headers hdr, inout metadata meta, inout standard_metadata_t
       standard_metadata) {
42     action drop() { mark_to_drop(); }
43
44     action ipv4_forward(bit<48> destinationMac, bit<48> sourceMac, bit<9> egressPort) {
45         [AAA] = destinationMac;          // Subproblem a)
46         [BBB] = sourceMac;               // Subproblem b)
47         [CCC].egress_spec = egressPort; // Subproblem c)
48         [DDD] = [DDD] - 1; }             // Subproblem d)
49
50     table routing_table {
51         key =     { [EEE]: lpm; }        // Subproblem e)
52         actions = { ipv4_forward;
53                     drop;
54                     NoAction; }
55         default_action = NoAction(); }
56
57     apply { routing_table.apply(); }
58 }
59
60 V1Switch(ParserImpl(), Pipeline(), DeparserImpl()) main;
```

Listing 2: Simple P4 program

The action `ipv4_forward()` in Listing 2 is incomplete. You can assume that the arguments of this action are filled with correct values according to their respective name. You will complete the functionality of this action in the following subproblems. Note that the router is simplified for this problem, i.e., the verification/calculation of the header checksum is not required.

a)* Replace `[AAA]` in Line 45 to create a valid software router.

0
½

```
hdr.eth.dstMacAddr
```

b)* Replace `[BBB]` in Line 46 to create a valid software router.

0
½

```
hdr.eth.srcMacAddr
```

c)* Replace `[CCC]` in Line 47 to create a valid software router.

0
½

```
standard_metadata
```

d)* Replace both instances of `[DDD]` in Line 48 to create a valid software router.

0
1

```
hdr.ip4.ttl
```

e)* Replace `[EEE]` in Line 51 to create a valid software router.

0
1

```
hdr.ip4.dstAddr
```

P4 supports different kinds of match types. The following problem investigates how these different match types can be used to represent subnetworks.

The subnet, investigated in the following subproblems, is a `/21` subnet, that contains the address `172.16.15.232`.

f)* The LPM match type entry requires the subnet address and its subnet mask. Write down the subnet mask in dot-decimal notation.

0
1

```
255.255.248.0
```

A ternary match is specified with a mask and a value. The mask defines the relevant bits for the ternary match (the relevant bits should be set to 1). The value defines if the relevant bits must be set to 1 or 0.

g)* Encode the subnet as a ternary match. Write down the mask and the value for the ternary match as **hexadecimal** numbers.

0
1
2

Mask: `0xFF FF F8 00`

Value: `0xAC 10 08 00` (host part of the address can be ignored, network part is sufficient)

0 1 2 h)* To map an entire subnet to a range match the IP addresses with the lowest and the highest value are required. Write down both addresses.

Lowest value: `172.16.8.0`

Highest value: `172.16.15.255`

0 1 i)* It is also possible to represent the subnet as a number of exact matches in a P4 table. How many entries are required for the given subnet?

For the `/21` subnet, $2^{(32-21)} = 2^{11} = 2048$ entries are required.

0 1 j) Assume a `/32` subnet containing the IP address `172.16.55.101`. Briefly argue which of the previously investigated match types is the most efficient regarding memory consumption.

The lpm, ternary and range match types require additional information besides the address (e.g., masks or a second address). The exact match type only requires the address per entry, i.e., it requires the least amount of memory per entry.

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

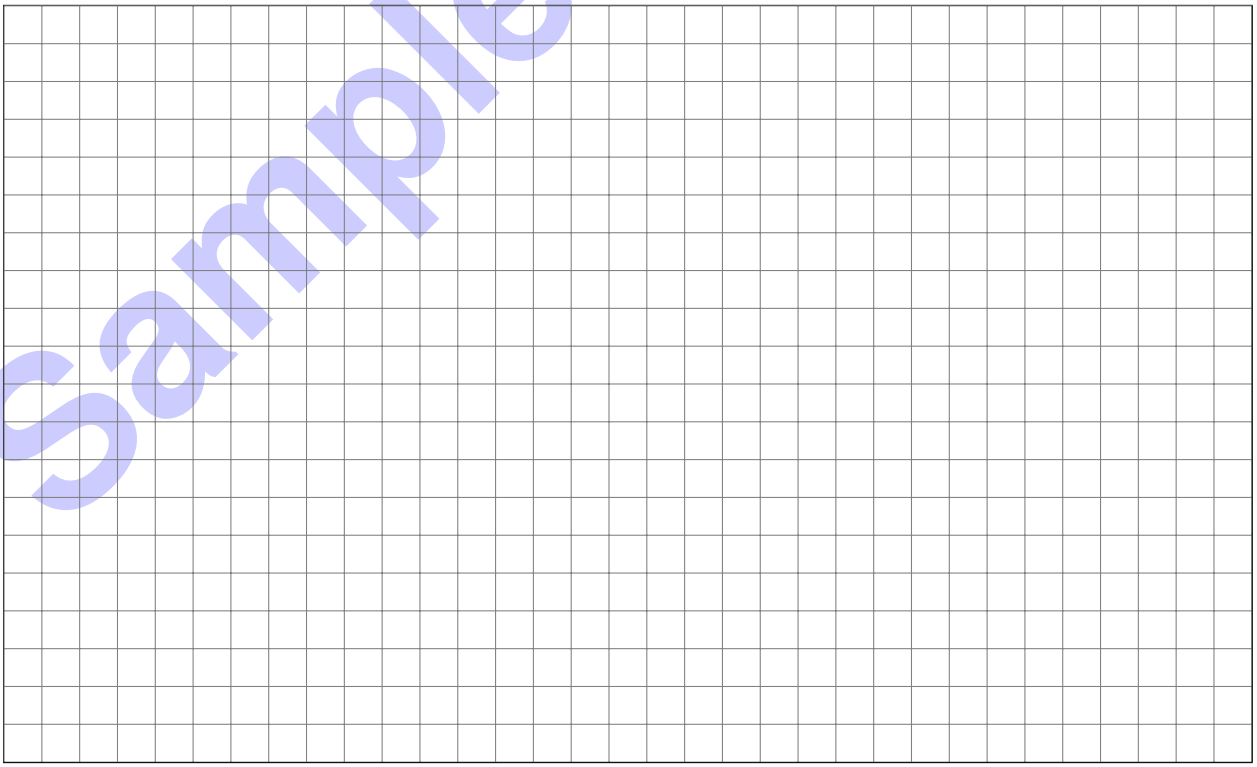

Additional space for solutions–clearly mark the (sub)problem your answers are related to and strike out invalid solutions.