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 / Retake-Online | **Date:** | Monday 11th April, 2022 |
| **Examiner:** | Prof. Dr.-Ing. Georg Carle | **Time:** | 14:15 – 15:30 |

## 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 without annotations
  - the **provided cheatsheet** without annotations (print or digital copy)
- 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.
- Code of conduct:
  - I participate without the help of others and only use the allowed resources.
  - I do not share, discuss, or exchange any information related to the exam with anybody.
  - I feel in good health and I am able to participate in the exam.
  - I understood the examination policy, agree to the video supervision, and adhere to this process.

## Problem 1 Quiz (15 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* ×■

Each question has **one single correct answer** that is worth 1 credit.

a)* What is the length of an IPv6 address?

☐ $2^{128}$ byte ☐ 128 byte ☐ $2^{128}$ bit ☐ 16 byte

b)* Which of the following IPv6 addresses is a Solicited Node Address?

☐ `ff02::1:ffb9:fd7f` ☐ `ff01::1:ffb9:fd7f` ☐ `2002::1:bbb9:fd7f` ☐ `ff02::1:bbb9:fd7f`

c)* How long does it roughly take to scan the complete IPv4 address space with ZMap and a rate of 10 000 packets per second? Assume the complete address space is probed.

☐ 50 h ☐ 200 h ☐ 24 h ☐ 120 h

d)* Which statement about a Software-Defined Network (SDN) **is wrong**?

☐ In SDNs, a single control plane may control several forwarding planes.
☐ SDNs rely on a physically centralized control plane.
☐ In SDNs, control plane and data plane are separate.
☐ An SDN is not required to realize Network Function Virtualization (NFV).

e)* Which match type is **not** part of the core P4 language according to the P4 specification?

☐ exact ☐ range ☐ ternary ☐ lpm

f)* Which congestion control algorithm is **loss-based**?

☐ TCP BBR ☐ TCP Vegas ☐ TCP Illinois ☐ TCP Reno

0 1 g)* What is the Organisation Unique Identifier (OUI) and at which position can it be found?

0 1 h)* Does the Spanning Tree Protocol (STP) prevent routing loops?

i)* What is the main difference between classful and classless IP addressing, regarding prefix length?

0 1

j)* Name the two latest DNS encryption mechanisms discussed in the lecture.

0 1

k)* Explain where DNS encryption protocols introduced in the lecture are operating and why that is useful.

0 1 2

l)* What is the difference between an active and an inactive timeout and when would they be triggered?

0 1 2

m)* Why do QUIC ZMap scans require more bandwidth than TCP ZMap scans?

0 1

## Problem 2 Longest Prefix Matching (LPM) (14 credits)

Longest Prefix Matching (LPM) is the algorithm that is performed by a router. The router performs the LPM for each packet to determine the next hop of a packet and perform an appropriate forwarding decision.

There are serveral notations to represent the network parts of IPv4 addresses. The slash-notation `/x`, the dotted-decimal notation `a.b.c.d`, or a subnetmask represented as a hex value `0xffffffff`.

0 1

a)* Convert `/17` into the dotted-decimal notation.

0 1

b)* Convert `255.255.254.0` into a subnetmask as a hex value.

0 1 2

c)* Explain briefly if `128.255.255.255` is a valid subnet mask.

A major part of the LPM is the matching operation (cf. Listing 1). The given matching function has three unsigned 32-bit integers as input arguments. The first argument contains the IPv4 address of the packet that should be forwarded (`uint32 addr`). The second argument contains the IPv4 address of a routing entry that should be matched (`uint32 entry`). The third argument contains the subnet mask of the routing entry (`uint32 mask`). `match()` returns true if the entry matches, false otherwise.

```
1 boolean match(uint32 addr, uint32 entry, uint32 mask) {
    // enter code here
3 }
```

Listing 1: Matching function

0 1 2

d)* Complete the `match` function of Listing 1.

e)* Explain why the routing entry of the default gateway is typically using a subnet mask of 0.

0 1 2

For the following suproblems, we perform the LPM using a list of routing table entries as a routing table. This list contains a number of n routing table entries. LPM iterates over the list of routing table entries sequentially to determine the longest matching prefix.

f)* LPM selects the "longest" entry. What does that mean?

0 1

g) Briefly explain, how many matches must be performed for an **unsorted** list with n unique entries, before the longest entry can be selected.

0 1

For the following subproblem, we sort the routing table using the **IPv4 address values**. The list is sorted in ascending order, i.e., the list begins with the routing entry containing the lowest address value.

h) Briefly explain, how many matches must be performed for this **address-sorted** list with n unique entries, before the longest entry can be selected.

0 1 2

For the following subproblem, we sort the routing table using the **IPv4 subnet masks**. The list is sorted in ascending order, i.e., the list begins with the routing entry containing the lowest subnet mask value.

i) Briefly explain, how many matches must be performed for this **mask-sorted** list with n unique entries, before the longest entry can be selected.

0 1 2

## Problem 3 Load Balancing (15 credits)

Load balancing is important to properly serve large amounts of clients with a good user experience. The following problem is based on the network shown in Figure 3.1.
A company operates two load blancers (LB1 and LB2) each connected to multiple content servers. Furthermore, they operate a nameserver (NS) used to resolve company domains to both load balancers.
On the opposite side, two clients (C1 and C2) frequently contact the company and access their content and a public DNS resolver (R) is available.

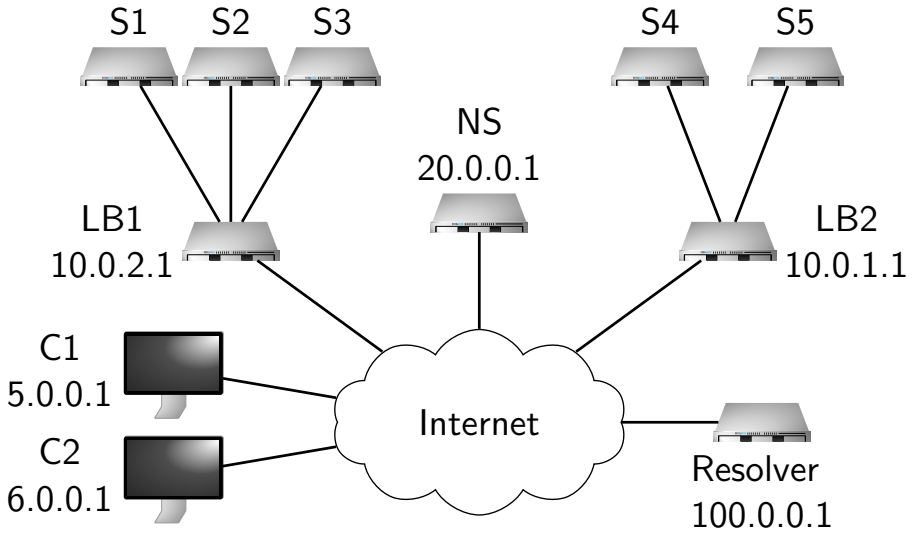


Figure 3.1: Important company with exam relevant content

0 1 a)* The company nameserver's zone file contains resource records listed in Table 3.1. Fill in the missing information, so that the nameserver can actively balance the load in between the load balancers.

Table 3.1: Company NS Resource Records

| | | | | |
|---|---|---|---|---|
| relevant.exam | 900 | IN | | |
| relevant.exam | 900 | IN | | |

0 1 b)* Explain why it is not a good idea for the company nameserver to return both records in one reply.

0 1 c)* If C1 and C2 use the public DNS resolver, why is geographical load balancing not possible with traditional DNS?

d) Name an extension of DNS that solves the previous problem and shortly explain how it works.

0 1 2

e) What is the downtime (in minutes) of the service in the worst case if one of the load balancer fails?

0 1

f)* Based on the lecture, how can LB1 use modulo hashing to load balance connections to the content servers?

0 1

g) What happens if one of the content servers fails especially using modulo hashing?

0 1

h)* How does consistent hashing help with the previous problem?

0 1

i) Assume consistent hashing is used but a planned update needs to be done on a content server. Describe what needs to be done to not disrupt services while using all available servers as long as possible.

0 1 2

j) Instead of DNS load balancing, how else could traffic be load balanced between LB1 and LB2 using routing? Does it prevent service disruptions if a load balancer fails?

0 1 2

## Problem 4 Network Calculus (12.5 credits)

This problem investigates performance bounds in networks using Network Calculus.
Consider the following network topology with flow and server descriptions. Flow $f_1$ traverses Servers $s_1$, $s_2$, and $s_3$. Flow $f_2$ traverses Servers $s_2$ and $s_3$. Assume each server handles flows according to strict priority scheduling with preemption. Assume Flow $f_1$ **has a low priority** and Flow $f_2$ **has a high priority**. In the following, we want to apply the Separate Flow Analysis to compute an end-to-end delay bound **for Flow** $f_1$.
**Hint:** Use the following formula to calculate a left-over service curve: $\beta^{l.o.} = \left[\beta_{R,T} - \gamma_{r,b}\right]^+ = \beta_{R-r,\frac{b+R\cdot T}{R-r}}$. An output arrival curve is given by $\alpha^\star = \alpha_{r,b+r\cdot T}$.

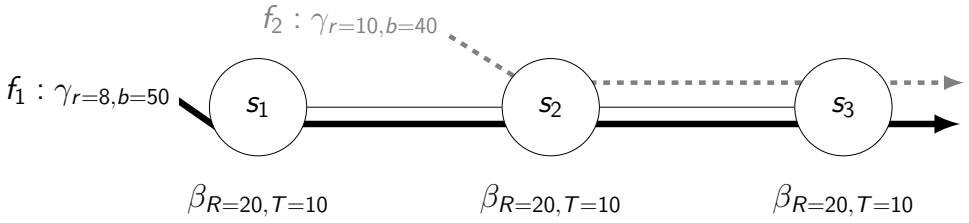


Figure 4.1: Network topology with flow and server specifications

0 1 2 3 4

a)* Perform the first step (left-over service curves) of the Separate Flow Analysis for $f_1$.

0 1 2

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

0 1

c)* Perform the third step (calculate the delay bound) of the Separate Flow Analysis. If you have no solution to b), assume an end-to-end left-over service curve of $\beta_{R=10,T=100}$.

d)* Consider the following scenario to the network in Figure 4.1:

0 1

- The rate of server $s_1$ is set to $R = 10$

Explain the influence on the delay bound of flow $f_1$ as determined by the Separate Flow Analysis.

e)* Consider a token-bucket constrained flow, traversing a single rate-latency server. The corresponding arrival- and service curves are shown in Figure 4.2. Determine the **burst value of the flow** and the **rate of the server**.

0 1 2

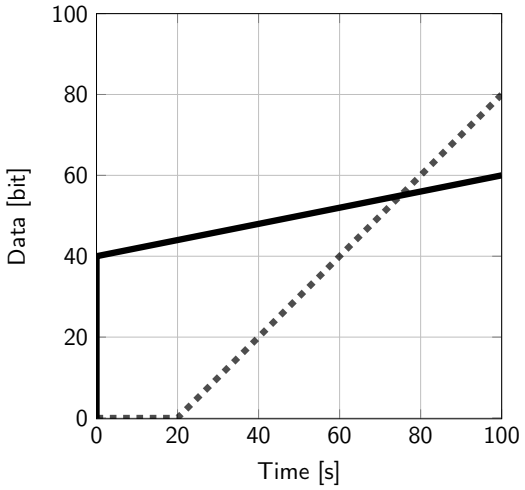


Figure 4.2: Arrival- and service curve

f)* Consider the following statement:

0 1

- A rate-latency service curve is defined as $\beta_{R,T}(s) = R \cdot s + T$ with rate parameter $R$ and latency parameter $T$.

Argue whether or not the statement is correct.

g)* Name one mathematical framework that can be used to obtain soft real-time guaratees **and cannot** be used to obtain hard real-time gurantees.

0 1

## Problem 5 Wireshark (18.5 credits)

The ISO-OSI model defines seven layers in a communication system. For each layer multiple protocols exist. In this problem, a frame is analyzed, referring to the involved protocols. You are given a hexdump of an Ethernet frame including FCS, starting with the Ethernet header. For the following problems we already marked the headers of different layers in Figure 5.1.

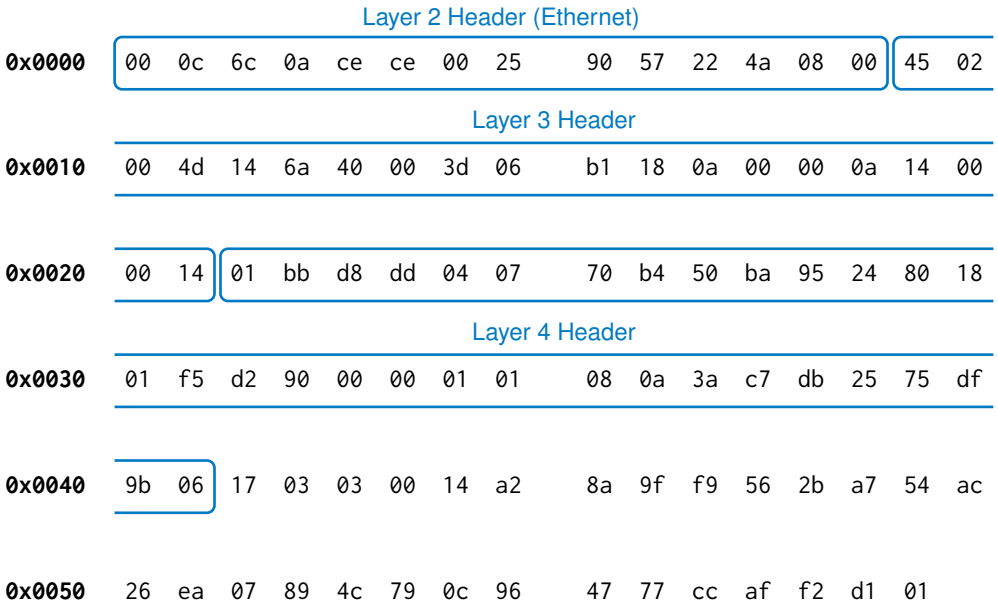


Figure 5.1: Hexdump of an Ethernet frame including FCS

In this problem you **always** have to substantiate your answers using the bytes in the hexdump in Figure 5.1. 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] = 0x000c6c`.

0 1 2

a)* What is the name of **Layer 2 and Layer 3** in the OSI model?

**Layer 2:**

**Layer 3:**

0 1

b)* In Figure 5.1, mark the FCS.

0 1 2

c)* What is the size of the **Layer 3 PDU** in bytes.

0 1 2

d)* What is the size of the **Layer 4 SDU** in bytes.

e) Compute the share of the L4 SDU in the whole hexdump.

0 1

f)* Identify the **Layer 3** protocol.

0 1

g) List all addresses contained in the **Layer 3** header in the common notation.

0 1 2

h) Identify the **Layer 4** protocol.

0 1

i) Make an educated guess which protocol can be expected **on top of Layer 4**.

0 1

j)* Can the two subnets `10.0.0.0/8` and `11.0.0.0/8` be merged? If yes, give the resulting network address and subnet.

0 1

k)* Name **and** briefly explain **two** design goals of QUIC, which target weaknesses present with TCP / TLS.

0 1 2

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