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

**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.

Sample Solution

# Advanced Computer Networking

**Module:** IN2097 **Date:** Friday 3$^{rd}$ March, 2017

**Examiner:** Prof. Dr.-Ing. Georg Carle **Exam:** Final exam

| | 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.
- Answers can be given in **English** or **German**
- 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 **printed dictionary** German ↔ native language
- Physically turn off all electronic devices, put them into your bag and close the bag.

## Problem 1 Quiz (7 credits)

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

a)* Perform the *k*-core algorithm for the topology shown in the solution box.

1. List the *k* value for every node in the graph, i. e. the current *k* value when the node is removed.
2. Mark the core of the network after applying the *k*-core algorithm.

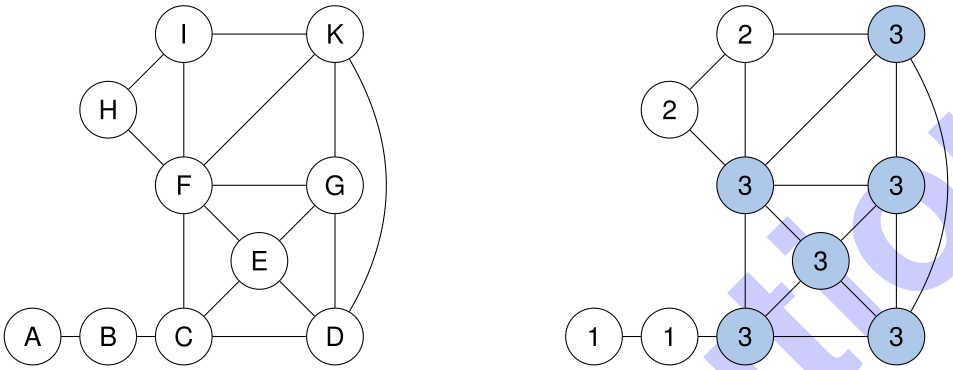

b)* Modern network interface cards (NICs) use standardized switchable transceivers, such as the SFP module. What is a benefit of such a system compared to a NIC with a fixed transceiver?

- Support for multiple systems using different physical media on the same NIC, i.e. multi mode fiber, single mode fiber, coax
- Replacing of defect lasers/LEDs for fiber based systems

c)* Give the technical term for the "cost free traffic exchange arrangement" between ASes.

Peering

d)* Name two routing protocols besides BGP.

IS-IS, RIP, OSPF

Sample Solution

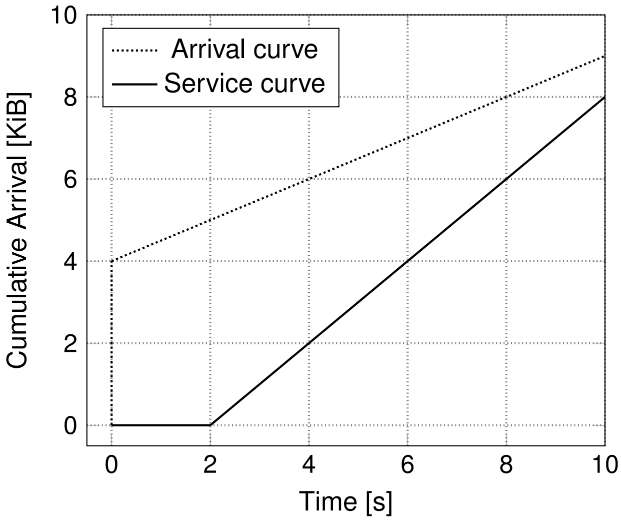

Figure 1.1: Arrival and service behavior of a packet processing system

e)* Determine the rate *R* and the latency *L* for the service curve from the graph in Figure 1.1.

□ 0 □ 1

$R = 1KiB/s$

$L = 2s$

f)* Mark the delay bound in Figure 1.1 and give the value in the answer box.

□ 0 □ 1

Maximum horizontal distance between arrival curve and service curve: 6s.

Sample Solution

## Problem 2 Routing - small scale (17.5 credits)

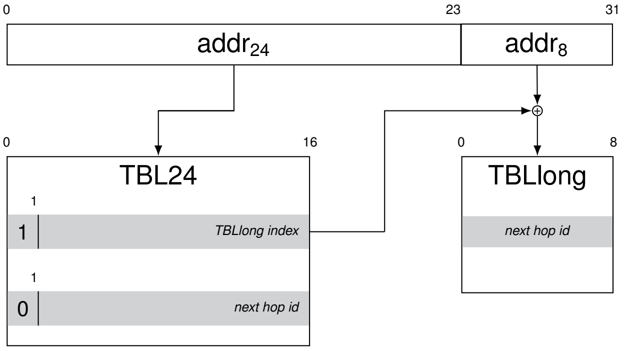

Figure 2.1: Table data structures used by DIR-24-8

Sample Solution

The routing table is one of the key data structures of routers. This data structure is necessary for answering the forwarding decision based on the routes added. The following exercise uses the DIR-24-8-BASIC routing table, which utilizes a two-stage lookup process:

- **First stage:** This stage is used for routes with a prefix of /24 or shorter. The IP address to be looked up is split in two parts, the most significant 24 bit ($\text{addr}_{24}$, see Figure 2.1) and the least significant 8 bit ($\text{addr}_8$, see Figure 2.1). For this stage, only $\text{addr}_{24}$ is used. $\text{addr}_{24}$ acts as an index to table TBL24. TBL24 contains $2^{24}$ 16 bit entries, i. e., exactly one entry for each of the $2^{24}$ values $\text{addr}_{24}$ can have. Such an entry in TBL24 has the most significant bit set to 0 and directly contains the next hop id.

- **Second stage:** This stage is used for routes with a prefix /25 or higher. In this case the $2^{24}$ entries in TBL24 are not enough and therefore a second table called TBLlong is used. Such a more specific entry has the most significant bit set to 1. Its least significant 15 bit together with the $\text{addr}_8$ part of the IP address is used as an index to TBLlong to get the corresponding next hop id.

0
1

a)* Calculate the size of the TBL24 data structure in MiB (**Note:** $2^{20}$ B = 1 MiB).

$$2^{24} \cdot 16 = 2^{24} \cdot 2^1 \cdot 8 = 2^{20} \cdot 2^5 \cdot 8 = 2^5 MiB = 32 MiB$$

0
1

b)* Assume that 256 /25 subnets or smaller are supported. Calculate the maximum size of the TBLlong data structure in KiB (**Note:** $2^{10}$ B = 1 KiB).

$$256 \cdot 256 \cdot 8 = 2^8 \cdot 2^8 \cdot 8 = 2^{10} \cdot 2^6 \cdot 8 = 2^6 KiB = 64 KiB$$

0
½

c)* Calculate the maximum number of next hop ids which can be handled by the described DIR-24-8 implementation.

$$2^8 = 256$$

Sample Solution

Consider a DIR-24-8 that contains the following routes:

Table 2.1: Example routes

| label | destination subnet | next hop id |
|---|---|---|
| A | `10.4.0.0/23` | 1 |
| B | `10.8.16.0/24` | 2 |
| C | `10.16.0.16/32` | 3 |

d)* Give the number of accesses to TBL24 and TBLlong for the routing table containing the addresses given in Table 2.1.

□ 0 □ 1 □ 2 □ 3

| Destination address | Accesses to TBL24 | Accesses to TBLlong |
|---|---|---|
| `10.8.16.232` | 1 | 0 |
| `10.16.0.16` | 1 | 1 |
| `10.4.1.16` | 1 | 0 |

e)* Enter the entries for the route with label **A** given in Table 2.1 into a DIR-24-8 routing table. Assume that TBLlong is empty. Include the indices of each entry in each array. **Note:** Use hexadecimal values.

□ 0 □ 1 □ 2

| Index | Entry<br>TBLlong index / next hop id |
|---|---|
| 0xA0400 | 1 |
| 0xA0401 | 1 |

| Index | Entry<br>Next hop id |
|---|---|
| - | - |
| - | - |

f)* Enter the entries for the route with label **C** given in Table 2.1 into a DIR-24-8 routing table. Assume that TBLlong is empty. Include the indices of each entry in each array. **Note:** Use hexadecimal values.

□ 0 □ 1 □ 2

| Index | Entry<br>TBLlong index / next hop id |
|---|---|
| 0xA1000 | 0x8000 (or 0x8001) |

| Index | Entry<br>Next hop id |
|---|---|
| 0x10 | 0x3 |

0 □□ □□ 1 □□ □□ 2 □□

g) According to the creators, DIR-24-8 is an algorithm which only can be used for IPv4. Give a reason why this algorithm is not used for IPv6 and its larger 128 bit addresses.

- The memory consumption for data structure would be unreasonably high (>TiB range)
- Every additional bit for either of the two tables would double the memory needed for the respecive data structure

0 □□ □□ 1 □□ □□ 2 □□ □□ 3 □□

h) Look at your solution for Problem d). What are the differences for memory accesses and what could be the reason for handling the routes differently?

- Fewer accesses for larger subnets 1 access for `/24`, 2 for smaller ones
- There will be more larger subnets than smaller subnets in a usual routing table , therefore handling larger subnets faster than smaller subnets makes sense

0 □□ □□ 1 □□ □□ 2 □□

i) Look at your previous results concerning the memory consumption and the number of data structure accesses of DIR-24-8. Argue whether the authors optimized their algorithm either for the memory consumption or data structure accesses.

- Few accesses (1 access for `/24` or larger, 2 for smaller ones) was one of the optimization goals
- Rather simple algorithm (easy to implement in hardware)
- Memory consumption rather high (always >32 MiB), however memory is cheap

Sample Solution

## Problem 3 Internet-wide Measurements (13.5 credits)

a)* Name three differences between the measurement tools nmap and ZMap.

□ 0
□ 1

stateless architecture/approach, no timeout detection, randomization based on multiplicative group, no sophisticated protocol exchanges, no TCP payload, keep track of sent packets, zmap faster, Internet-wide scans possible, zmap can be distributed/sharding

You are now conducting an IPv4-wide measurement using ZMap to evaluate the security of home routers over the course of two weeks. You are probing one specific UDP port per target IP address.

b)* Describe two time-related phenomena that could bias your measurements.

□ 0
□ 1
□ 2

diurnal and weekly patterns , IP address churn (could lead to the same target being scanned multiple times)

Before starting the scan, you decide which target IP addresses will not be scanned. Table 3.1 shows your subnet blacklist compiled from complaints from previously conducted scans.

Table 3.1: Blacklist

| Subnet |
|---:|
| 123.45.6.2/26 |
| 8.8.8.0/24 |
| 47.23.1.128/28 |
| 8.8.4.4/25 |
| 123.45.6.5/32 |
| 12.33.254.17/28 |
| 8.8.4.196/27 |

c)* How many IP addresses are omitted in the scans due to the blacklist shown in Table 3.1?

□ 0
□ 1

The /32 subnet is included in /26 subnet.

$$/28 + /28 + /27 + /26 + /25 + /24 \rightarrow /23 \rightarrow 2^9 addresses = 512 addresses \qquad (1)$$

Sample Solution

0 ½

d)* Which additional IP addresses should you blacklist in your scanning effort if you have access to your upstream's BGP export?

unannounced prefixes

The number of target IP addresses shall be denoted by $n$. We send a fixed UDP payload of 18 B. The scan starts on Feb 2, at 16:00:00 and ends on Feb 16, at 15:59:59.

0 1

e)* Calculate the size of the layer 2 Ethernet frame which contains the given UDP datagram. Assume minimal sized headers for all protocols. **Hint:** There is a cheat sheet.

$$p_{size} = 14_{eth} + 20_{ipv4} + 8_{udp} + 18_{payload} + 4_{fcs} = 64 Byte$$

0 1

f)* The measurement tool generates Ethernet frames with a rate of 1024 kbit/s. How many packets per second are sent?

$$r_{pps} = \frac{r}{64 * 8} = \frac{1,024,000}{64 \cdot 8} = 2000pps = 2kpps$$

0 1

g) Determine the number of targets $n$ over the whole measurement period. **Note:** Document your approach, and fill in the values. You do not need to calculate the result.

$$n = t_{measurement} * r_{pps} = 14 \cdot 24 \cdot 3600s \cdot 2000pps = 2,419,200,000packets \approx 2.4 \text{ billion packets}$$

Now let us look at the responses. The UDP response rate is $p$, the ICMP destination unreachable (type 3) rate is $q$. The UDP response payload is double the request payload. The ICMP response contains the original IP header and the first 8 B of the original IP packet's payload. Assume minimally-sized headers without additional extensions or options.

0 ½

h)* Calculate the Ethernet frame size $s_{UDP}$ for the UDP response?

$$14 + 20 + 8 + 36 + 4 = 82$$

Sample Solution

i)* Calculate the Ethernet frame size $s_{ICMP}$ for the ICMP response?

□ 0
□ ½

14 + 20 + 8 + 20 + 8 + 4 = 74

j) How much layer 2 payload $P$ do we receive from the scans? **Note:** Document your approach. You do not need to calculate the result.

□ 0
□ ½

$P = nps_{UDP} + nqs_{ICMP}$

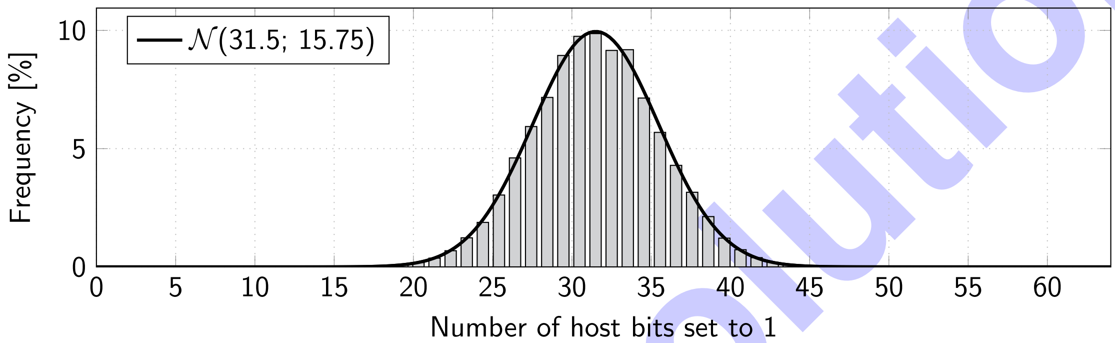

Figure 3.1: Host bit distribution of observed devices

In addition to the active IPv4 measurements you are conducting passive IPv6 measurements at an IXP observation point. Figure 3.1 shows the hostbit distribution, i. e., the number of bits set to '1' in the interface identifier of the observed IPv6 addresses.

k)* What type of devices do we mostly see at the observation point? Explain your reasoning!

□ 0
□ 1
□ 2

We see mostly client/end user devices with a random IID (IPv6 privacy extensions)

l) Why does the normal distribution $\mathcal{N}$ approximate the hostbit distribution?

□ 0
□ 1

Law of large numbers, central limit theorem, sum of single distribution approximates normal distribution

m) Why do we use $\mu = 31.5$ instead of $\mu = 32$?

□ 0
□ ½

global/local (universal/local) bit → mean is decreased by 0.5

Sample Solution

## Problem 4 Wireshark (18 credits)

We consider the IP packet depicted in Figure 4.1 as hexdump in network byte order.

```
                    Receiver MAC                    Transmitter MAC       Ethertype  IHL  TOS
0x0000   34  e6  d7  a6  c1  d0  0c  c4     7a  90  f2  32  08  00   45   10

      Total Length Identification Flags/FragOffset TTL Prot.  HdrChecksum  Source Address  Destination
0x0010   00  3c  14  b1  40  00  40  06     f4  e0  0A  10  20  40   0A   10

         Address   SrcPort   DstPort      Sequence Number     Acknowledgement Number  Offset/Flags
0x0020   20  41  00  16  9a  9e  74  19     b7  ea  50  4a  be  c5   80   18

         Window    Checksum  Urgent Pointer               Options
0x0030   01  5c  31  49  00  00  01  01     08  0a  5f  1a  f8  4b   00   50

                                  Data                         FCS
0x0040   43  72  00  00  00  b0  d2  aa     0f  c9  3b  b6  35  2a
```

Figure 4.1: Complete hexdump of an Ethernet frame

Sample Solution

**Note:**

- To solve this problem, use the cheat sheet that is handed out separately.
- Give a reason for all answers, e. g. answering subproblem c) with "UDP" without further comment gives no credit even if "UDP" was correct.
- Marking headers / fields directly in Figure 4.1 is probably a good idea.

0 □ 1 □ 2 □

a)* Mark and name all header / trailer fields on layer 2 of the Ethernet frame in Figure 4.1.

b)* What is the purpose of the FCS?

0 □ 1 □

Detect transmission errors.

0 □ 1 □

c) What is the type of the L3-PDU?

Ethertype `0x0800` ⇒ IPv4

0 □ 1 □

d) Mark the end of the layer 3 header in Figure 4.1.

e) Determine the source and destination address of layer 3. Write the addresses in their usual, human readable form.

□ 0 □ 1 □ 2

Source: 10.16.32.64
Destination: 10.16.32.65

f) What is the type of the L4-PDU?

□ 0 □ 1

IP protocol number `0x06` ⇒ TCP

g) What is most likely the application layer protocol?

□ 0 □ 1

TCP source port number `0x0016` = 22 ⇒ SSH

h) Determine the beginning of the L4-SDU.
**Hint:** TCP headers are aligned at 4 B boundaries.

□ 0 □ 1 □ 2

TCP Offset `0x8` ⇒ 32 B header length

Sample Solution

Table 4.1: TCP options

| Type | Description |
|---|---|
| 0 | End of option list |
| 1 | No-Operation (NOP) |
| 2 | Maximum segment size |
| 3 | Window scale |
| 4 | SACK permitted |
| 5 | SACK |
| 8 | Timestamps |

The hexdump shown in Figure 4.2 represents the second TCP segment exchanged while establishing a connection. The header includes multiple TCP options, among others the window scaling option. The first byte of each option specifies its type, the second byte its total length in multiples of 1 B. An exception is the NOP option used for padding, which has a fixed length of 1 B. Table 4.1 lists some common TCP options and the type values.

```
0x0000  01 bb e3 cc 3a 11 77 ad   a4 f6 55 2c a0 12 71 20
0x0010  a1 fd 00 00 02 04 05 b4   04 02 08 0a 6a 12 1f 2e
0x0020  62 bf 17 47 01 03 03 09
```

Figure 4.2: Hexdump of a TCP segment during handshake (including TCP header)

0 □□ 1 □□ 2 □□ 3 □□

i)* List the options that are contained in the segment of Figure 4.2 in the order of their appearance.

MSS, SACK permitted, Timestamp, NOP, window scaling

0 □□ 1 □□ 2 □□

j)* Explain how window scaling works.

Increase TCP window sizes by a factor of $2^x$ where $x$ is the value of the window scaling option. It is used to increase window sizes beyond what can normally be specified.

0 □□ 1 □□

k)* Which particular problem is addressed by the window scaling option?

Low transmit rates in networks with high bandwidth-delay-product.

0 □□ 1 □□

l) Determine the window scaling factor as power of 2.

$x = 9 \Rightarrow 2^9$

Sample Solution

# Problem 5 SSH (4 credits)

Consider the network topology displayed in Figure 5.1. Client C is connected to Server S via an ssh connection. The ssh connection uses the default router on Server S via the Internet.

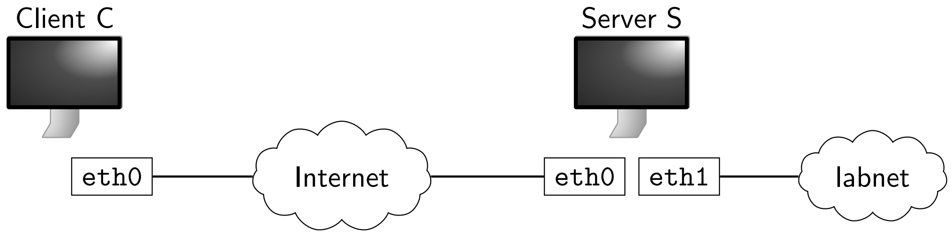

Figure 5.1: Network topology

a) Which packets take the default route in general?

□ 0
□ 1

- The default route matches against all packets (if not a more specific route matches) and forwards the traffic accordingly.

Server S has a second interface attached to an internal lab network. Now, the default route on Server S is changed via the ssh connection towards the lab network. After that you do not get any answer back from the Server S. You type *reboot* and press enter in your still open ssh console on Client C.

b) Explain and reason in detail what happens next.

□ 0
□ 1
□ 2
□ 3

- The server will not reboot because:
  - Packets toward the ssh client will use the new default route
  - These packets cannot reach the client any more and connection breaks down:
  - The connection can break down for several reasons: a timeout can happen if the reboot is not not entered immediately, also the TCP window could be full, i.e. no more packets can be sent without an ACK.
- The server will reboot because:
  - Packets toward the ssh client use the new default route, therefore no ACKs can reach the server
  - However, the TCP connection works as long as no timeout is detected and the TCP window still allows segments to be sent
  - Therefore, the reboot can reach the server and is executed despite the lacking ACKs

Sample Solution

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

Sample Solution

Sample Solution

Sample Solution