Name

Vorname

Studiengang (Hauptfach)

Fachrichtung (Nebenfach)

Matrikelnummer

Unterschrift der Kandidatin/des Kandidaten

....
Note

TECHNISCHE UNIVERSITÄT MÜNCHEN

Fakultät für Informatik

- ☐ Midterm-Klausur
- ☒ Final-Klausur

- ☐ Semestralklausur
- ☐ Diplom-Vorprüfung
- ☐ Bachelor-Prüfung
- ☐ ....

- ☐ Einwilligung zur Notenbekanntgabe per E-Mail / Internet

Prüfungsfach: Master Course: Computer Networks

Prüfer: Prof. Dr.-Ing. Georg Carle **Datum: February 16, 2013**

Hörsaal: .... **Reihe:** .... **Platz:** ....

	I	II
1		
2		
3		
4		
5		
6		
7		
8		
9		
10		
Σ		

Nur von der Aufsicht auszufüllen:

Hörsaal verlassen von ...... : ...... bis ...... : ......

Vorzeitig abgegeben um .... : ......

Besondere Bemerkungen:

Endterm

Master Course: Computer Networks

Prof. Dr.-Ing. Georg Carle
Chair for Network Architectures and Services
Department of Computer Science
Technische Universität München

Saturday, February 16, 2013
9:00 a.m. – 10:00 a.m.

- This exam consists of **15 pages** and a total of **4 problems** as well as an **additional handout** which contains a reference of protocol headers. Please make sure that you got a complete copy of all documents.
- Write your name and matriculation number in the header of **every** page.
- Do neither write with red / green colors nor use pencils.
- The total amount of credits is 50.
- This exam is **closed book**, i. e., lecture notes, homework, cheat sheets, pocket calculators etc. are **not** allowed.
- Turn off your mobile phones and put them into your bag.
- Problems marked by * can be solved without knowledge of previous results.
- **Results are only rated if your approach is reproducible**. If not instructed otherwise, state a reason for all your answers.

Problem 1 Protocol dissemination (8 credits)

8

Consider the hexdump (network byte order) given in Table 1. It shows an IEEE 802.3 FastEthernet frame (preamble and checksum are stripped). In the following we will disseminate this frame step by step. The additional handout accompanying the exam might be helpful.

Byte	0	1	2	3	4	5	6	7	8	9	a	b	c	d	e	f
0000	00	25	90	57	1f	dc	28	37	37	02	32	41	86	dd	60	00
0010	00	00	00	23	11	40	20	01	4c	a0	20	01	00	11	5d	92
0020	47	55	86	2e	34	65	20	01	4c	a0	20	01	00	17	00	00
0030	00	00	00	00	01	97	d8	ad	00	35	00	23	fa	1b	08	7f
0040	01	00	00	01	00	00	00	00	00	00	06	67	6f	6f	67	6c
0050	65	02	64	65	00	00	01	00	01							

Table 1: Hexdump (network byte order) of an IEEE 802.11 FastEthernet frame, preamble and checksum are stripped.

a)* Explain the difference between host and network byte order.

1

Network byte order is big endian ✓, host byte order is the native byte order of some processor architecture ✓, e.g. x86 is little endian, PowerPC is big endian.

b) Which network layer protocol is being used (give a reason)?

1

Ethertype 0x86dd (big endian) = IPv6 ✓

c) Which transport layer protocol is being used (give a reason)?

1

Next header 0x11 = UDP ✓

d) What kind of payload does the frame carry (give a reason)?

1

UDP destination port 53 = DNS ✓

1

e) What is the protocol identified in (d) being used for?

Translation of FQDN into IPv6 address ✓

Assume that the frame given in Table 1 is going to be transmitted through an MPLS network. The MPLS header is depicted in Figure 1.1.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Label																				**Exp.**			**S**	**TTL**							

Figure 1.1: MPLS header (offset in bit)

1

f)* At which point of the frame is the MPLS header inserted?

After the Ethernet header, i. e., beginning at offset 0x0e. ✓

1

g)* Explain what the label field is used for.

Identifies the path a frame should take through the MPLS network. ✓

1

h)* What is the advantage of using MPLS?

Speed up IP forwarding. ✓
(also accepted: network virtualization)

Problem 2 Traceroute and IP alias resolution (15 credits)

15

This problem covers the well-known tool `traceroute` and techniques to resolve IP aliases.

a)* Briefly describe how `traceroute` works.

2

- Host sends IP packets with incrementing TTL starting at TTL=1. ✓
- TTL is decremented by routers. ✓
- When the TTL reaches 0, a router responds with ICMP Time Exceeded / TTL Exceeded in transmit ✓ that is returned to the sender of the original packet ✓.

b)* Optionally `traceroute` may use TCP instead of UDP or ICMP. State two advantages / disadvantages.

1

- If the port number is chosen appropriately (e.g. TCP 80 when tracing a webserver), the probe will probably traverse firewalls near the target. ✓
- TCP probes may be filtered by intermediate nodes or when directed to some port the target is not listening on. ✓

c)* In practice, paths between two nodes on the Internet are often asymmetric, i. e., the path from some node A to B is different from the reverse path. Give two possible reasons for this phenomenon.

2

- Asymmetric routing between ASs due to provider policies. ✓
- Asymmetric links, i. e., uplink and downlink characteristics may differ and influence routing decisions. ✓
- (Hot potato routing)

Table 2 lists a the output of `traceroute` issued from from two nodes A and B to each other.

Hop	$A \rightarrow B$	$A \leftarrow B$
1	85.214.1.25	46.4.173.33
2	85.214.0.168	213.239.224.1, 213.239.224.65
3	85.214.0.71	213.239.240.150
4	80.81.192.164	80.81.192.110
5	213.239.240.240	85.214.0.70
6	213.239.224.8, 213.239.224.104	85.214.0.169
7	213.239.199.116	85.214.236.18

Table 2: Output from `traceroute` executed on A and B, respectively.

3

d)* Based on the output listed in Table 2, derive a likely network topology by completing the Figure below.
Note: You do not have to write down every single IP address, but make sure we are able to map your graph to Table 2.

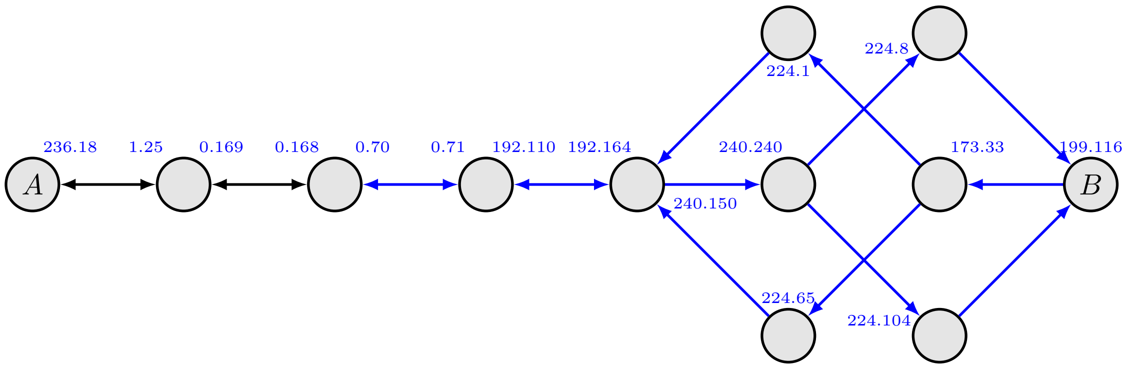

- Correct continuation of the chain ✓
- Load balancer ✓
- Asymmetric reverse path ✓

(slightly other topologies may also be possible)

Based on a more or less reasonable assumption, you probably assigned two or more different IP addresses to the same node in (d).

e) Select one of those nodes from (d) and give a reason, why you think this node has more than one IP address. 1

A sees 0.168 for its second hop. Its gateway may have an IP address close to 0.168 on the interface directed towards B, but A cannot know about it. However, the traceroute from B shows 0.169 for the last hop before reaching A. This is an indication that 1.25 and 0.169 are aliases. ✓
(similar argument may hold for 0.168 and 0.70)

Given the initial assumption, we want to verify that two IP addresses are indeed aliases for the same node. As we have seen in the homework, there are quite a few different approaches.

f)* Briefly describe how absolute values of IP identifiers can be used to detect IP aliases. 2

- ICMP TTL exceeded messages are new IP packets that should carry individual IP identifiers. ✓
- When a node is probed from both sides, the identifiers should be close to each other. ✓

g) Describe two problems when using IP identifiers. 2

- There is no guarantee that routers increment the identifiers linearly. ✓
- Other traffic may increment the identifier causing jumps when recording a sequence of identifiers. ✓

h) Why may it be helpful to use the difference between consecutive IP identifiers? 2

If cross traffic disturbs the results, ✓ the slope of identifiers (their differences) should be very similar when the router is probed from both sides. ✓

14

Problem 3 IP addressing, NAT, and SCTP (14 credits)

In this problem we consider IP addressing, NAT, and SCTP in the network depicted in Figure 3.1. The private network on the left hand side is connected via a NAT-enabled router to a public network. Router SP1 acts as SOCKS4a proxy for HTTP connections. SP1 and SP2 use a proprietary HTTP-over-SCTP implementation (similar to what we did in the project), i. e., incoming HTTP messages on SP1 are sent as part of an SCTP association to SP2 which decapsulates the HTTP messages and forwards them to the webserver and vice versa. You may assume that all links shown in Figure 3.1 are FastEthernet segments.

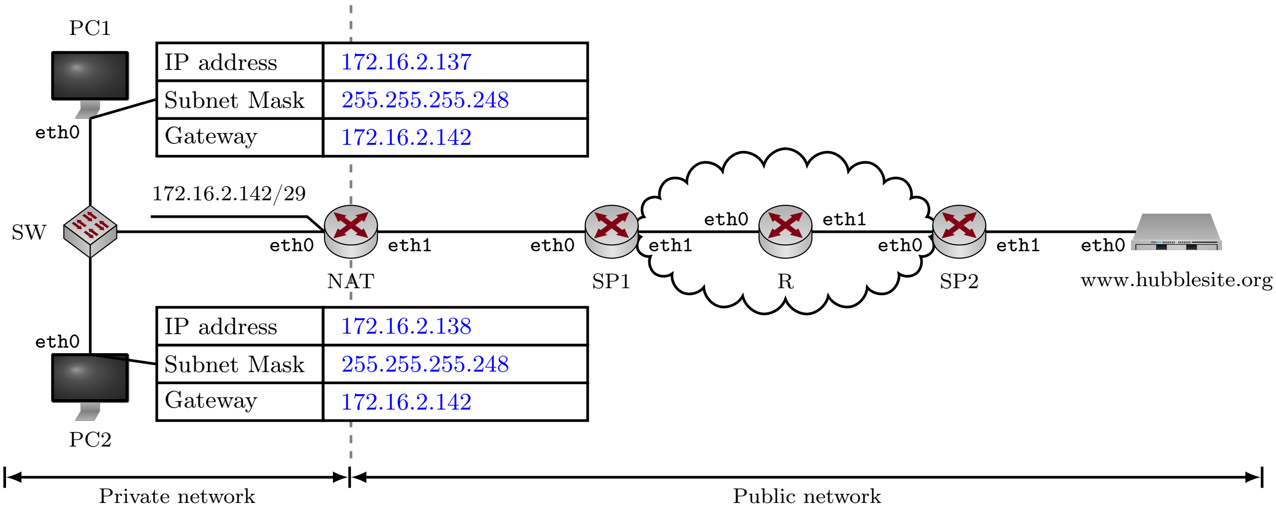

Figure 3.1: Network topology

In the following, we will first assign addresses to the private network and consider some fundamental problems. Afterwards, you are asked to state the contents of specific protocol headers at different points in the network.

1

a)* Based on the local address of the NAT router, derive the network and broadcast address of the private network.

Network address is 172.16.2.136 ✓, broadcast address is 172.16.2.143 ✓

1

b) Assign meaningful IP addresses, subnet masks, and default gateways to both PC1 and PC2. Write the configuration directly into Figure 3.1.

1

c)* Why do we need a NAT-enabled router in this setup?

Private IP addresses are not routed in the Internet. ✓

(also accepted: private IPs are not unique) d)* Briefly describe at least two advantages of SCTP over TCP.

1

Multihoming ✓, flexibility (configurable reliability) ✓.
(others possible)

e)* What may be a problem if the local PCs would try to directly establish an SCTP connection to SP2?

1

The NAT would have to support SCTP. ✓

f)* Assuming that the restrictions considered in (e) do not apply, what is the advantage of exchanging the positions of SP1 and the NAT?

1

Connections from all local PCs can be served by a single SCTP association and thus only one mapping in the NAT table. ✓

Figure 3.2 is a larger copy of our network. We assume that PC1 tries to establish an HTTP session to the webserver www.hubblesite.org via its SOCKS4a proxy SP1. The message shown in Figure 3.2 may be the HTTP request sent by the client.

The next problem asks you to state the contents of specific header fields at three different points in the network. You may abbreviate MAC and IP addresses of individual devices using the naming convention `<device>.<interface>.<layer>`, e.g. `NAT.eth1.MAC` means the MAC address on interface eth1 of the NAT router while `NAT.eth1.IP` denotes the IP address on that interface.

g) Fill out the header fields in the three boxes shown in Figure 3.2. If the content of a field is not uniquely defined, make a **meaningful** choice. Abbreviate MAC and IP addresses as stated above!

8

Comments:

- 0.5 credits off for each wrong or missing entry
- at most 1 credit off per header field

Src MAC	PC1.eth0.MAC
Dst MAC	NAT.eth0.MAC
Src IP	PC1.eth0.IP
Dst IP	SP1.eth0.IP
TTL	64 (common defaults)
L4 protocol	TCP
Src Port	30712 (almost arbitrary)
Dst Port	9050 (SOCKS)
L7 protocol	SOCKS4a + HTTP

Src MAC	SP1.eth0.MAC
Dst MAC	R.eth1.MAC
Src IP	SP1.eth1.IP
Dst IP	SP2.eth0.IP
TTL	128 (pobably new packet)
L4 protocol	SCTP
Src Port	20000 (almost arbitrary)
Dst Port	80 (may differ)
L7 protocol	proprietary

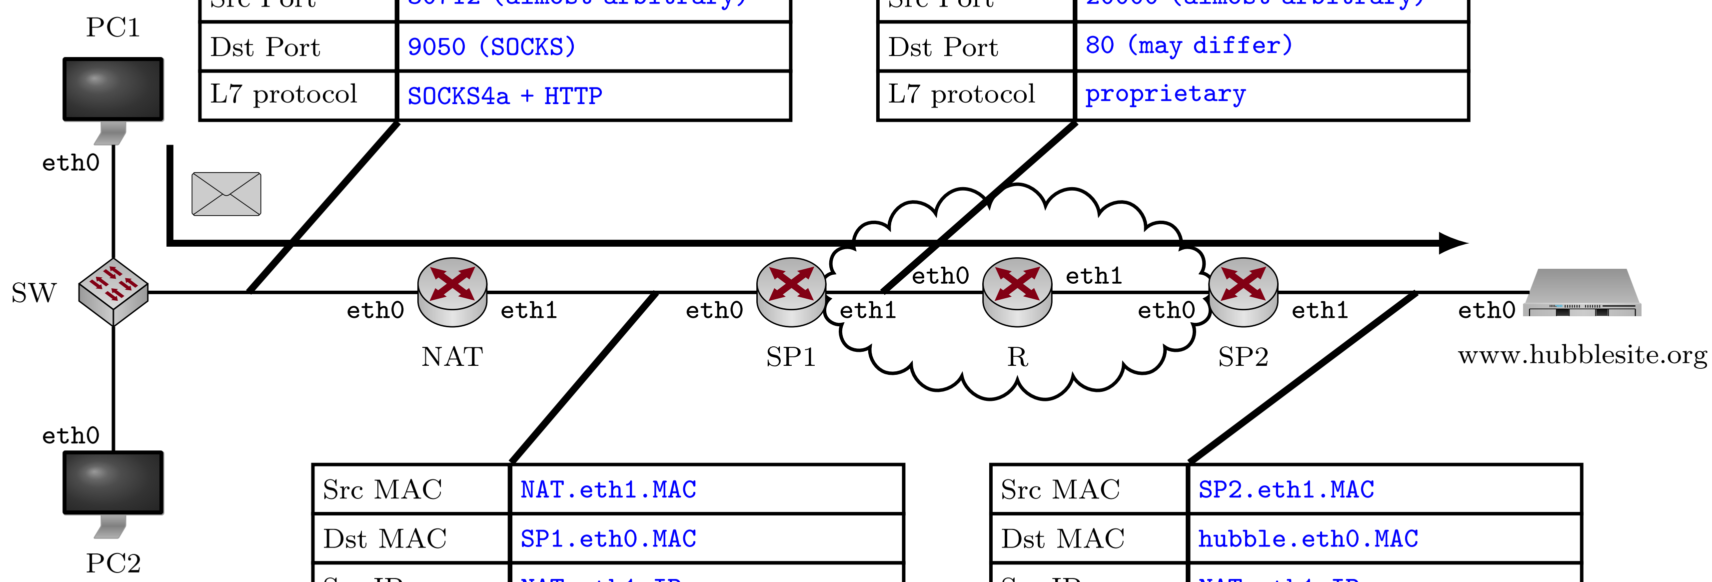

Src MAC	NAT.eth1.MAC
Dst MAC	SP1.eth0.MAC
Src IP	NAT.eth1.IP
Dst IP	SP1.eth0.IP
TTL	63 (one less)
L4 protocol	TCP
Src Port	31217 (almost arbitrary)
Dst Port	9050 (SOCKS)
L7 protocol	SOCKS4a + HTTP

Src MAC	SP2.eth1.MAC
Dst MAC	hubble.eth0.MAC
Src IP	NAT.eth1.IP
Dst IP	hubble.eth0.IP
TTL	128 (probably new packet)
L4 protocol	TCP
Src Port	31217
Dst Port	80
L7 protocol	HTTP

Figure 3.2: Preprint for subprblem (g)

Problem 4 TCP congestion control (13 credits)

13

In this problem we consider the congestion control mechanism of TCP. We denote the size of TCP's sender window by $w[n]$ given in multiples of the MSS and depending on discrete time steps n given in multiples of the RTT. We assume that $w[n]$ depends on the current value of the congestion window only, i. e., the receiver window is larger than the maximum value W of the sender window.

a)* What is the purpose of congestion control?

1

Avoids overload situations in the network. ✓

b)* What is the receiver window used for?

1

It allows the receiver to limit the sender's window, i. e., flow control. ✓

Assume that a new TCP connection has just been established at time index $n = 0$. We consider the ideal case where no segment loss occurs until the sender window reaches its maximum value, i. e., $w[n] = W$. For simplicity, we assume that W is a power of two. When the maximum value is reached, we assume that **a single segment** is lost and timely retransmitted by the sender. This leads to the time-discrete development of $w[n]$ depicted in Figure 4.1.

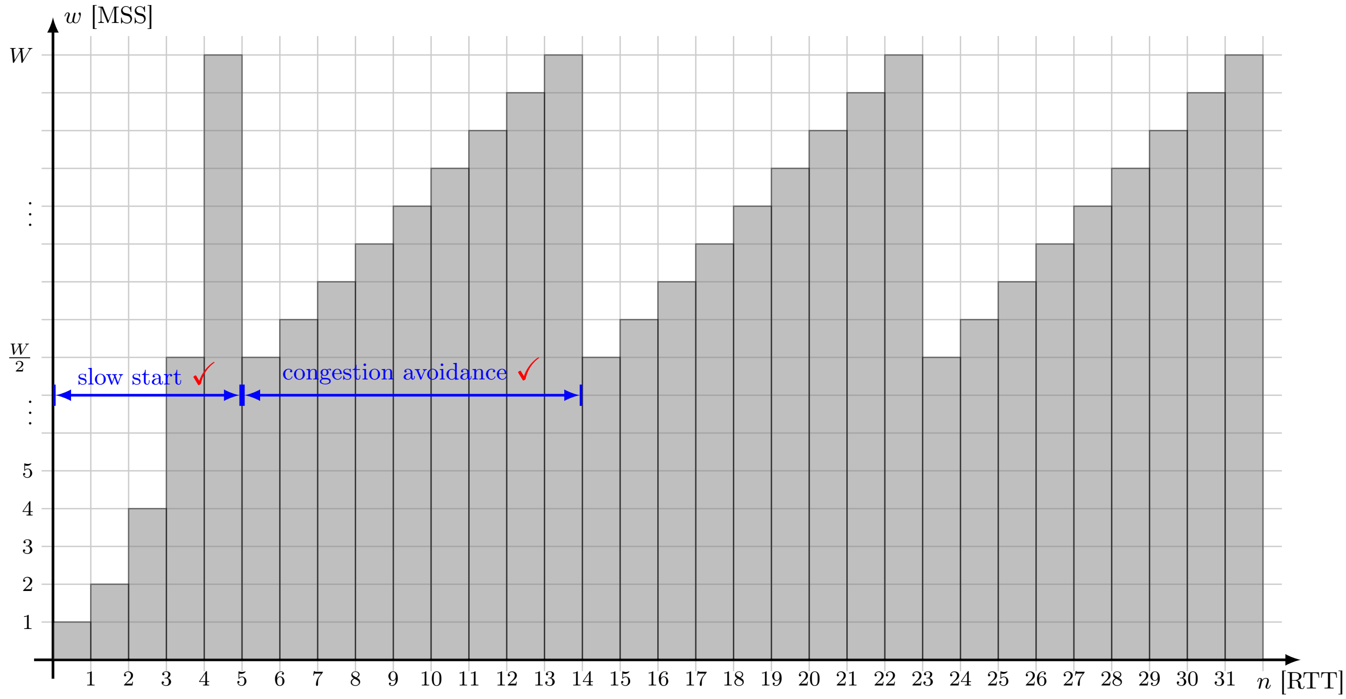

Figure 4.1: Development of the TCP sender window over time.

1

c)* Mark and name the different phases of TCP's congestion control in Figure 4.1.

The average number of segments during the first phase of congestion control is given by

$$N_\alpha = 2W - 1. \tag{1}$$

2

d)* Prove Equation (1).

Total number of segments during slow start are given by

$$N_\alpha = \sum_{i=0}^{\log_2(W)} 2^i \checkmark = 2W - 1 \checkmark$$

The average number of segments during the second phase is given by

$$N_\beta = \frac{3}{8}W^2 + \frac{3}{4}W. \tag{2}$$

3

e)* Prove Equation (2).

Total number of segments during congestion avoidance are given by

$$\begin{aligned} N_\beta &= \sum_{i=W/2}^{W} i \checkmark \\ &= \sum_{i=1}^{W} i - \sum_{i=1}^{W/2-1} i \\ &= \frac{W \cdot (W+1)}{2} - \frac{\left(\frac{W}{2} - 1\right) \cdot \frac{W}{2}}{2} \checkmark \\ &= \frac{W^2 + W}{2} - \frac{W^2}{8} + \frac{W}{4} \\ &= \frac{3}{8}W^2 + \frac{3}{4}W \checkmark \end{aligned}$$

Assume that there is chance of $\epsilon = \frac{1}{9}$ that the sender fails to retransmit a lost segment in time. The fails are assumed to be statistically independent. This causes a timeout at the receiver and thus restarts the congestion control algorithm (we assume that the congestion threshold is also reset).

f)* Derive the expected number N of segments between two restarts of the congestion control mechanism. Simplify the result. 1

$$N = N_\alpha + 8N_\beta = 3W^2 + 8W - 1 ✓$$

g) Determine the expected time T between two restarts in dependency of W and RTT for $\epsilon = \frac{1}{9}$. 2

$$T = T_\alpha + 8T_\beta = \left(\log_2(W) + 1 + 8 \cdot \left(\frac{W}{2} + 1\right)\right) \cdot \mathrm{RTT} ✓$$
$$= (\log_2(W) + 4W + 9) \cdot \mathrm{RTT} ✓$$

Assume a RTT of $\frac{1}{11}$ s, a maximum window size of $W = 16\,\mathrm{MSS}$, and a maximum segment size of 1400 Byte.

h) Determine the expected transmit rate r in MB/s. 2

$$r = \frac{N}{T} = \frac{895 \cdot 1400\,\mathrm{Byte}}{77 \cdot \frac{1}{11}\,\mathrm{s}} ✓ = 179\,\mathrm{kB/s} ✓$$

Additional space for solutions – plaese clearly indicate to which problem your notes belong and strike invalid solutions.