Solution of Exercise Sheet 4

Exercise 1 (Routers, L3-Switches, Gateways)

1. Explain the purpose of Routers in computer networks.
   (*Also explain the difference to Layer-3-Switches.*)

   They forward packets between networks with different logical address ranges and provide a WAN interface.

2. Explain the purpose of Layer-3-Switches in computer networks.
   (*Also explain the difference to Routers.*)

   They are Routers too, which means they forward packets between networks with different logical address ranges, but they do not provide a WAN interface.

3. Explain the purpose of Gateways in computer networks.

   They enable communication between networks, which base on different protocols.

4. Explain why Gateways in the Network Layer of computer networks are seldom required nowadays.

   Modern computer networks operate almost exclusively with the Internet Protocol (IP). For this reason, a protocol conversion at the Network Layer is mostly not required.

Exercise 2 (Collision and Broadcast Domain)

1. Mark the devices that divide the collision domain.

   ☐ Repeater
   ☐ Hub
   ☒ Bridge
   ☒ Layer-2-Switch
   ☒ Router
   ☒ Layer-3-Switch

2. Mark the devices that divide the broadcast domain.

   ☐ Repeater
   ☐ Hub
   ☐ Bridge
   ☐ Layer-2-Switch
   ☒ Router
   ☒ Layer-3-Switch

3. Sketch in the diagram all collision domains and all broadcast domains.

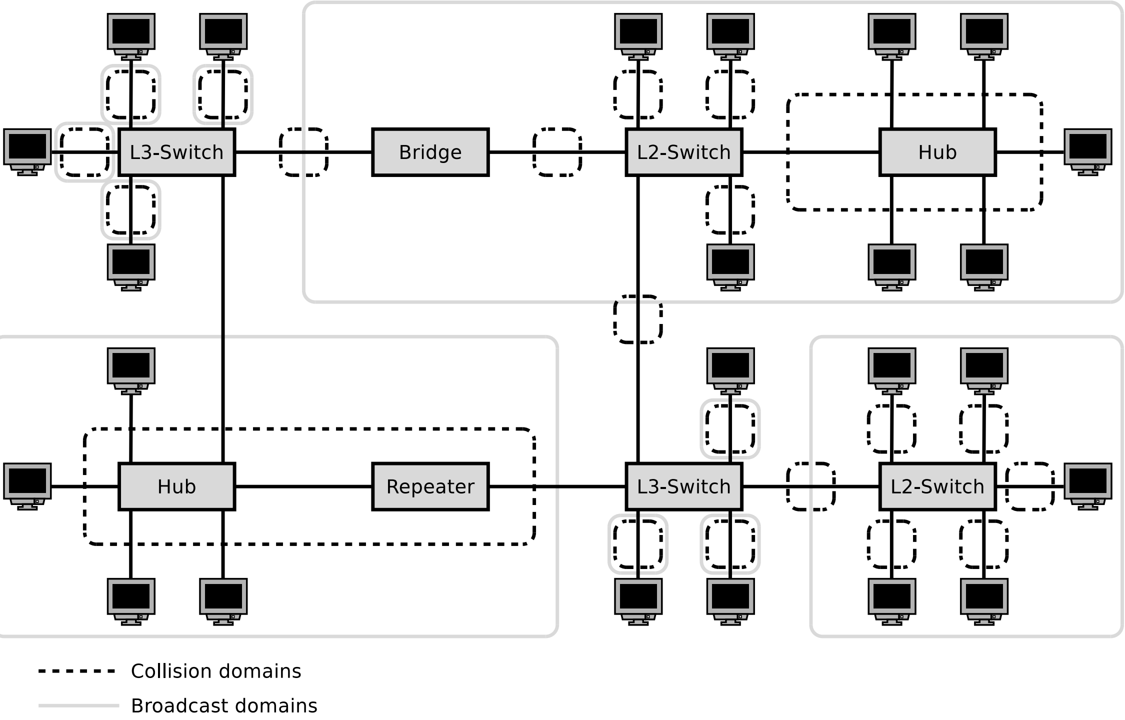

Exercise 3 (Addressing in the Network Layer)

1. Explain the meaning of Unicast in the Network Layer of computer networks.

   An IP address is assigned to a single receiver.

2. Explain the meaning of Broadcast in the Network Layer of computer networks.

   An IP address is assigned to all receivers in the subnet.

3. Explain the meaning of Anycast in the Network Layer of computer networks.

   An IP address is used to reach a single device of a group of devices.

4. Explain the meaning of Multicast in the Network Layer of computer networks.

   An IP address is assigned to a group of receivers.

5. Explain why the IPv4 address space does contain only 4,294,967,296 addresses.

   IPv4 addresses have a length of 32 bits (4 bytes). Thus, the address space contains $2^{32} = 4,294,967,296$ *possible addresses.*

6. Explain why Classless Interdomain Routing (CIDR) was introduced.

Because with address classes, many addresses are wasted and it is impossible to dynamically adjust address classes.

7. Describe in simple words the functioning of CIDR.
Focus on the way, how IP addresses are treated and subnets are created.

Since the introduction of CIDR, the address class of an IPv4 address is no longer important. All hosts in a network have a subnet mask assigned, which has a length of 32 bits (4 bytes) and is used to specify the number of subnets and hosts. The network mask splits the host ID of an IP address into subnet ID and host ID. 1-bits in the subnet mask indicate, which part of the address space is used for subnet IDs and 0-bits indicate, which part of the address space is used for host IDs.

Exercise 4 (Addressing in the Network Layer)

Calculate for each subtask of this exercise the first and last host addresses, the network address and the broadcast address of the subnet.

```
IP Address:         151.175.31.100   10010111.10101111.00011111.01100100
Subnet mask:        255.255.254.0    11111111.11111111.11111110.00000000
Part for host IDs:                                          x xxxxxxxx
Network address?    151.175.30.0     10010111.10101111.00011110.00000000
First host address? 151.175.30.1     10010111.10101111.00011110.00000001
Last host address?  151.175.31.254   10010111.10101111.00011111.11111110
Broadcast address?  151.175.31.255   10010111.10101111.00011111.11111111

IP Address:         151.175.31.100   10010111.10101111.00011111.01100100
Subnet mask:        255.255.255.240  11111111.11111111.11111111.11110000
Part for host IDs:                                                  xxxx
Network address?    151.175.31.96    10010111.10101111.00011111.01100000
First host address? 151.175.31.97    10010111.10101111.00011111.01100001
Last host address?  151.175.31.110   10010111.10101111.00011111.01101110
Broadcast address?  151.175.31.111   10010111.10101111.00011111.01101111

IP Address:         151.175.31.100   10010111.10101111.00011111.01100100
Subnet mask:        255.255.255.128  11111111.11111111.11111111.10000000
Part for host IDs:                                               xxxxxxx
Network address?    151.175.31.0     10010111.10101111.00011111.00000000
First host address? 151.175.31.1     10010111.10101111.00011111.00000001
Last host address?  151.175.31.126   10010111.10101111.00011111.01111110
Broadcast address?  151.175.31.127   10010111.10101111.00011111.01111111
```

binary representation	decimal representation	binary representation	decimal representation
10000000	128	11111000	248
11000000	192	11111100	252
11100000	224	11111110	254
11110000	240	11111111	255

Exercise 5 (Addressing in the Network Layer)

In each subtask of this exercise, a sender transmits an IP packet to a receiver. Calculate for each subtask the subnet ID of sender and receiver and specify whether the IP packet leaves the subnet during transmission or not.

```
Sender:      11001001.00010100.11011110.00001101    201.20.222.13
Subnet mask: 11111111.11111111.11111111.11110000    255.255.255.240
         AND -----------------------------------
             11001001.00010100.11011110.00000000
                                        ^^^^        => 0 = subnet ID sender

Receiver:    11001001.00010100.11011110.00010001    201.20.222.17
Subnet mask: 11111111.11111111.11111111.11110000    255.255.255.240
         AND -----------------------------------
             11001001.00010100.11011110.00010000
                                        ^^^^        => 1 = subnet ID receiver

       Does the IP packet leave the subnet [yes/no]? yes


Sender:      10000100.10011000.01010011.11111110    132.152.83.254
Subnet mask: 11111111.11111111.11111100.00000000    255.255.252.0
         AND -----------------------------------
             11000100.10011000.01010000.00000000
                               ^^^^^^               => 20 = subnet ID sender

Receiver:    10000100.10011000.01010001.00000010    132.152.81.2
Subnet mask: 11111111.11111111.11111100.00000000    255.255.252.0
         AND -----------------------------------
             11000100.10011000.01010000.00000000
                               ^^^^^^               => 20 = subnet ID receiver

       Does the IP packet leave the subnet [yes/no]? no


Sender:      00001111.11001000.01100011.00010111    15.200.99.23
Subnet mask: 11111111.11000000.00000000.00000000    255.192.0.0
         AND -----------------------------------
             00001111.11000000.00000000.00000000
                      ^^                            => 3 = subnet ID sender

Receiver:    00001111.11101111.00000001.00000001    15.239.1.1
Subnet mask: 11111111.11000000.00000000.00000000    255.192.0.0
         AND -----------------------------------
             00001111.11000000.00000000.00000000
                      ^^                            => 3 = subnet ID receiver

       Does the IP packet leave the subnet [yes/no]? no
```

```
Sender:      11010010.00000101.00010000.11000110    210.5.16.198
Subnet mask: 11111111.11111111.11111111.11111100    255.255.255.252
         AND -----------------------------------
             11010010.00000101.00010000.11000100
                                        ^^^^^^      => 49 = subnet ID sender

Receiver:    11010010.00000101.00010000.11000101    210.5.16.197
Subnet mask: 11111111.11111111.11111111.11111100    255.255.255.252
         AND -----------------------------------
             11010010.00000101.00010000.11000101
                                        ^^^^^^      => 49 = subnet ID receiver

       Does the IP packet leave the subnet [yes/no]? no


Sender:      11010010.00000101.00010000.11000110    210.5.16.198
Subnet mask: 11111111.11111111.11111111.11111100    255.255.255.252
         AND -----------------------------------
             11010010.00000101.00010000.11000100
                                        ^^^^^^      => 49 = subnet ID sender

Receiver:    11010010.00000101.00010000.11001001    210.5.16.201
Subnet mask: 11111111.11111111.11111111.11111100    255.255.255.252
         AND -----------------------------------
             11010010.00000101.00010000.11001000
                                        ^^^^^^      => 50 = subnet ID receiver

       Does the IP packet leave the subnet [yes/no]? yes


Sender:      00000101.00000101.00000101.00000101    5.5.5.5
Subnet mask: 11111111.11111110.00000000.00000000    255.254.0.0
         AND -----------------------------------
             00000101.00000100.00000000.00000000
                      ^^^^^^^                       => 2 = subnet ID sender

Receiver:    00000101.00000110.00000110.00000110    5.6.6.6
Subnet mask: 11111111.11111110.00000000.00000000    255.254.0.0
         AND -----------------------------------
             00000101.00000110.00000000.00000000
                      ^^^^^^^                       => 3 = subnet ID receiver

       Does the IP packet leave the subnet [yes/no]? yes
```

Exercise 6 (Addressing in the Network Layer)

Calculate for each subtask the subnet masks and answer the questions.

1. Split the class C network `195.1.31.0` for implementing 30 subnets.

```
Network ID:    11000011.00000001.00011111.00000000    195.1.31.0
Number of bits for subnet IDs? 30 => 32 $(= 2^5)$ => 5 bits
Subnet mask:   11111111.11111111.11111111.11111000    255.255.255.248
Number of bits for host IDs? 3
Number of host IDs per subnet? $2^3 - 2 = 6$
```

2. Split the class A network `15.0.0.0` for implementing 333 subnets.

```
Network ID:    00001111.00000000.00000000.00000000    15.0.0.0
Number of bits for subnet IDs? 333 => 512 $(= 2^9)$ => 9 bits
Subnet mask:   11111111.11111111.10000000.00000000    255.255.128.0
Number of bits for host IDs? 15
Number of host IDs per subnet? $2^{15} - 2 = 32766$
```

3. Split the class B network `189.23.0.0` for implementing 20 subnets.

```
Network ID:    10111101.00010111.00000000.00000000    189.23.0.0
Number of bits for subnet IDs? 20 => 32 $(= 2^5)$ => 5 bits
Subnet mask:   11111111.11111111.11111000.00000000    255.255.248.0
Number of bits for host IDs? 11
Number of host IDs per subnet? $2^{11} - 2 = 2046$
```

4. Split the class C network `195.3.128.0` into subnets, with 17 hosts each.

```
Network ID:    11000011.00000011.10000000.00000000    195.3.128.0
Number of bits for host IDs? 17 => 32 $(= 2^5)$ => 5 bits
Number of bits for subnet IDs? $8 - 5 = 3$ bit
Number of possible subnets? $2^3 = 8$
Subnet mask:   11111111.11111111.11111111.11100000    255.255.255.224
```

5. Split the class B network `129.15.0.0` into subnets, with 10 hosts each.

```
Network ID:    10000001.00001111.00000000.00000000    129.15.0.0
Number of bits for host IDs? 10 => 16 $(= 2^4)$ => 4 bits
Number of bits for subnet IDs? $16 - 4 = 12$ bit
Number of possible subnets? $2^{12} = 4096$
Subnet mask:   11111111.11111111.11111111.11110000    255.255.255.240
```

binary representation	decimal representation	binary representation	decimal representation
10000000	128	11111000	248
11000000	192	11111100	252
11100000	224	11111110	254
11110000	240	11111111	255

Exercise 7 (Collision and Broadcast Domain)

1. Sketch in the diagram of the network topology all collision domains and all broadcast domains.

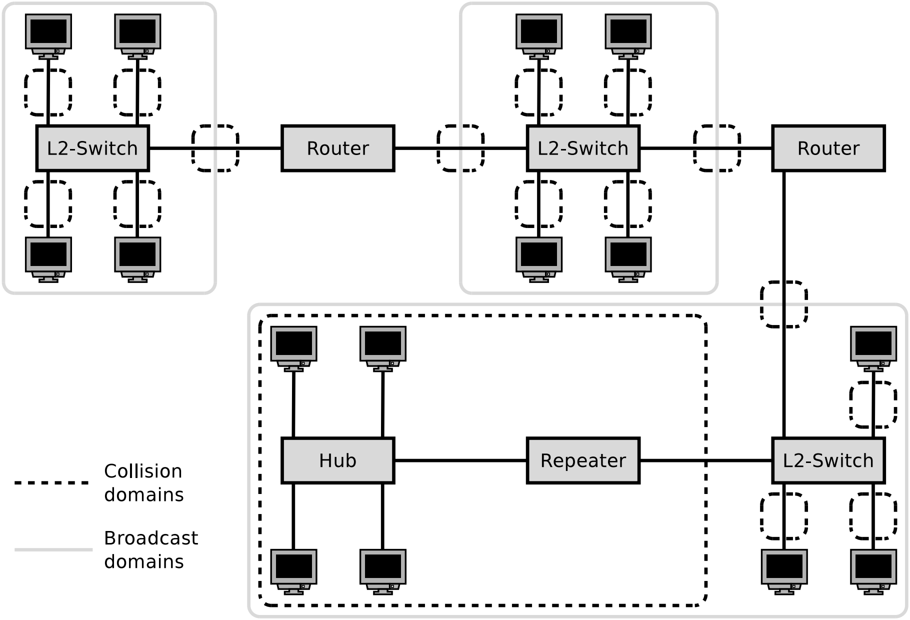

2. Sketch in the diagram of the network topology all collision domains and all broadcast domains.

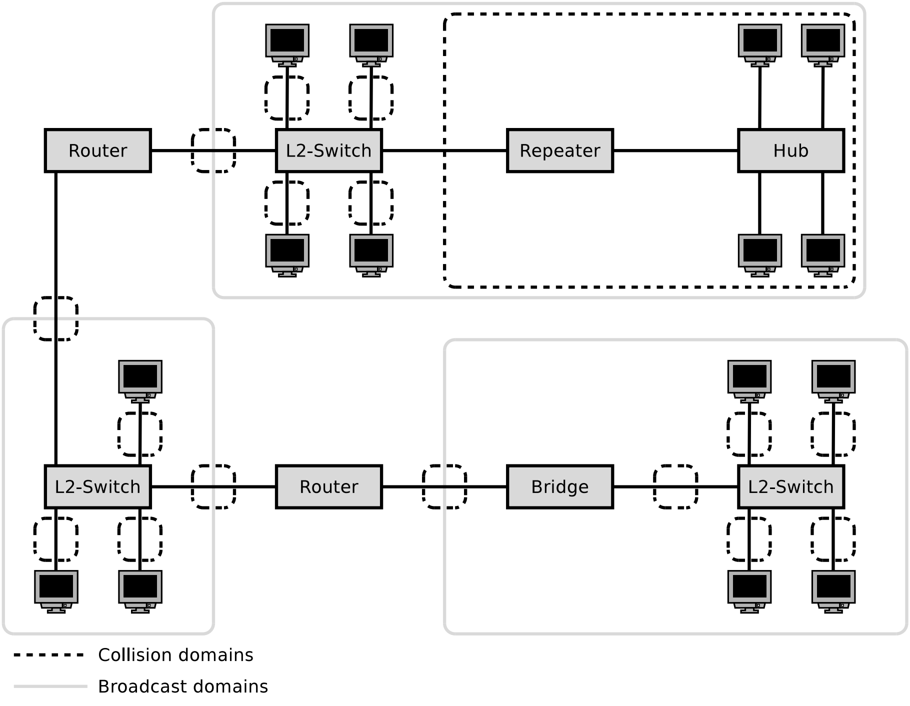

Exercise 8 (Broadcast Domain)

1. Sketch in the diagram of the network topology all broadcast domains.
2. Give the required number of subnet for this network topology.

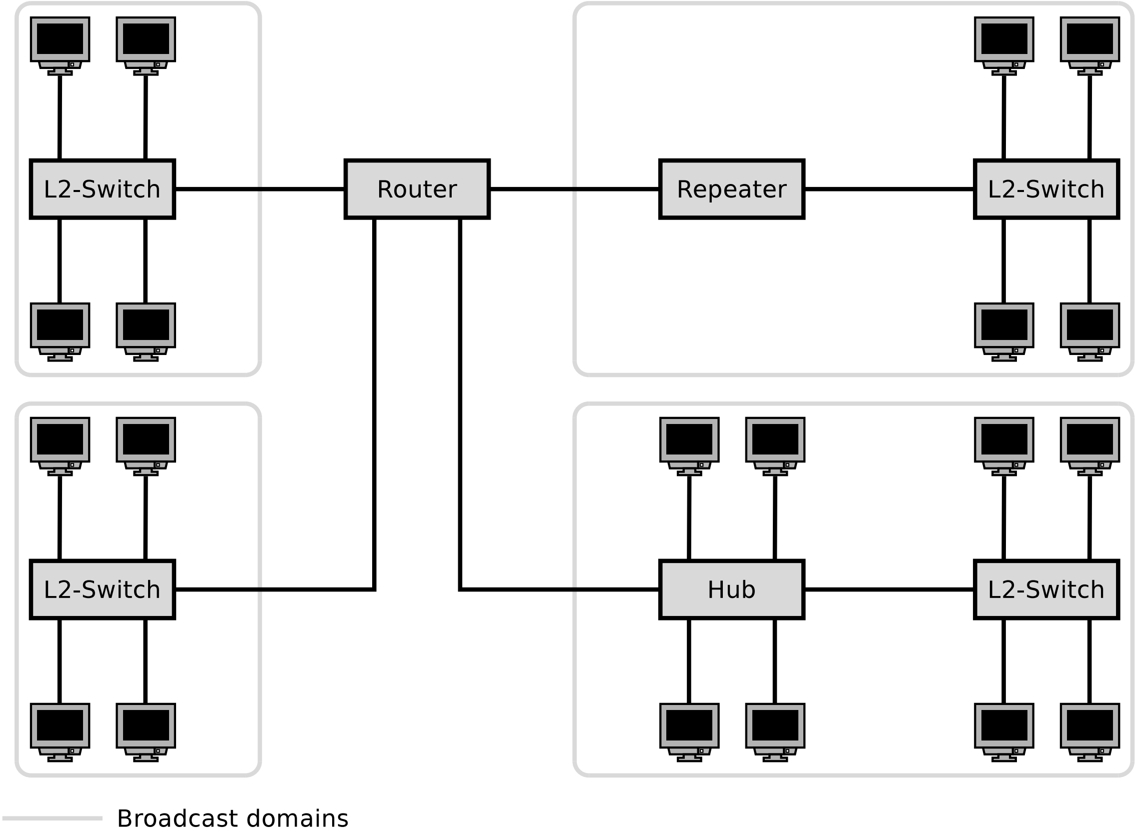

4 subnets are required because each port of a Router is connected to a different IP network. It is impossible to operate an IP subnet on multiple ports of a Router.

Exercise 9 (Private IP Address Spaces)

Name the three private IPv4 address spaces.

- 10.0.0.0/8
- 172.16.0.0/12
- 192.168.0.0/16

Exercise 10 (Fragmenting IP Packets)

4,000 bytes payload need to be transmitted via the IP protocol. The payload must be fragmented, because it is transmitted over multiple physical networks, whose MTU is $< 4,000$ bytes.

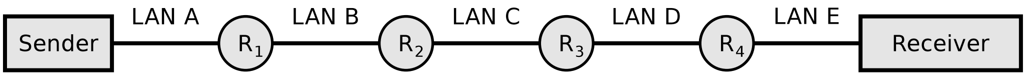

	LAN A	LAN B	LAN C	LAN D	LAN E
Network technology	Ethernet	PPPoE	ISDN	Ethernet	WLAN
MTU [bytes]	1,500	1,492	576	1,400	2,312
IP-Header [bytes]	20	20	20	20	20
max. payload [bytes] in theory	1,480	1,472	556	1,380	2,292
Multiple of 8	yes	yes	no	no	no
max. payload [Bytes] in practice	1,480	1,472	552	1,376	2,288

Display graphically the way, the payload is fragmented, and how many bytes of payload each fragment contains.

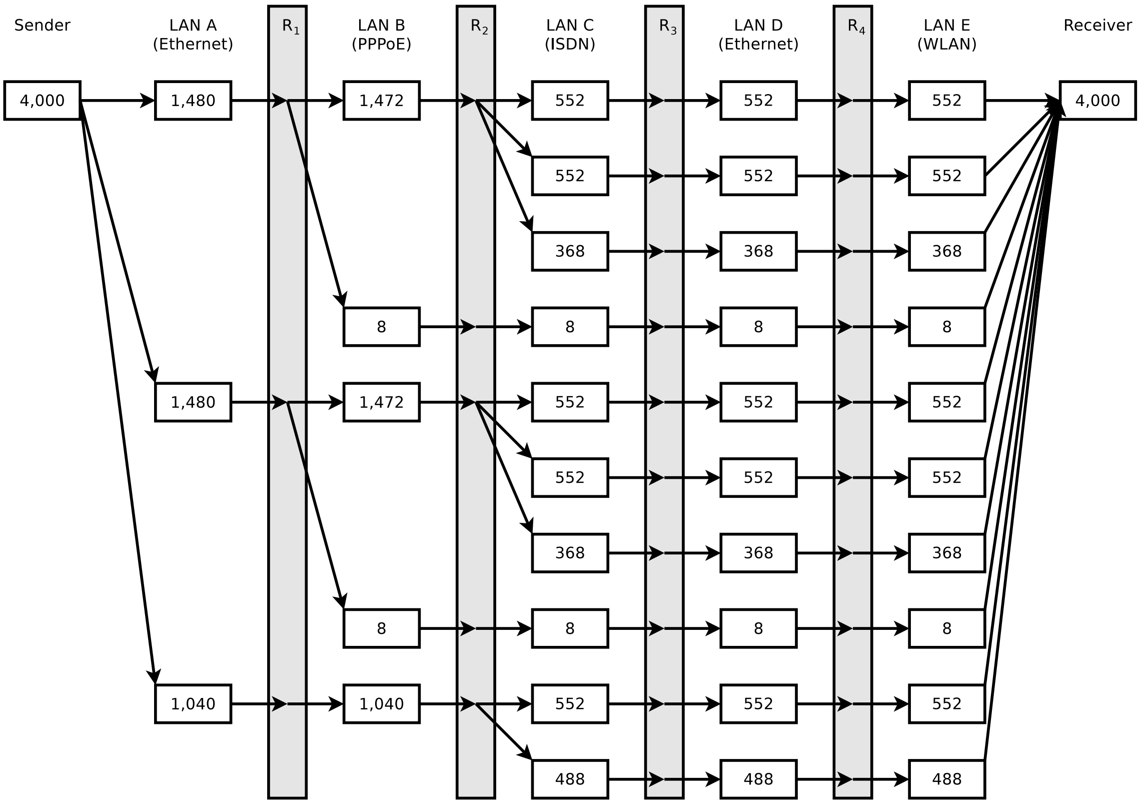

Exercise 11 (Forwarding and Path Calculation)

1. Name the two major classes of routing protocols.

   Distance Vector Routing Protocols and Link State Routing Protocols.

2. Name the algorithms for best path calculation, the routing protocol classes from subtask **??** do implement.

   Distance Vector Routing Protocols implement the Bellman-Ford algorithm.

   Link State Routing Protocols implement the Dijkstra algorithm.

3. Explain what an autonomous system is.

   Each AS consists of a group of logical networks, which use the Internet Protocol, are operated and managed by the same organization (e.g. an Internet Service Provider, a corporation or university) and use the same routing protocol.

4. The Border Gateway Protocol (BGP) is a protocol for...

   ☐ Intra-AS routing ☒ Inter-AS routing

5. Name the routing protocol class from subtask **??** that does the BGP implement.

   None. It implements path-vector routing, which has some similarities with distance-vector-routing.

6. Open Shortest Path First (OSPF) is a protocol for...

   ☒ Intra-AS routing ☐ Inter-AS routing

7. Name the routing protocol class from subtask **??** that does the OSPF implement.

   Link state routing.

8. The Routing Information Protocol (RIP) is a protocol for...

   ☒ Intra-AS routing ☐ Inter-AS routing

9. Name the routing protocol class from subtask **??** that does the RIP implement.

   Distance vector routing.

10. When RIP is used, each Router only communicates with its direct neighbors. Name advantages and drawbacks of this method.

    Advantage: The network is not flooded $\Longrightarrow$ *protocol causes little overhead.*

    Drawback: Long convergence time because updates propagate slowly.

11. When RIP is used, the path cost (metric) only depends on the number of Routers (hops), which need to be passed on the way to the destination network. Name a drawback of this method.

The metric hop count often results in routes, which are not optimal, because all network segments have an equal weight.

12. When OSPF is used, all Routers communicate with each other. Name advantages and drawbacks of this method.

    Advantage: Short convergence time.

    Drawback: The network is flooded $\Longrightarrow$ protocol causes strong overhead.

Exercise 12 (Dijkstra's Algorithm)

1. Calculate the shortest path from node A to all other nodes using Dijkstra's algorithm.

Source: Jörg Roth. Prüfungstrainer Rechnernetze. Vieweg (2010)

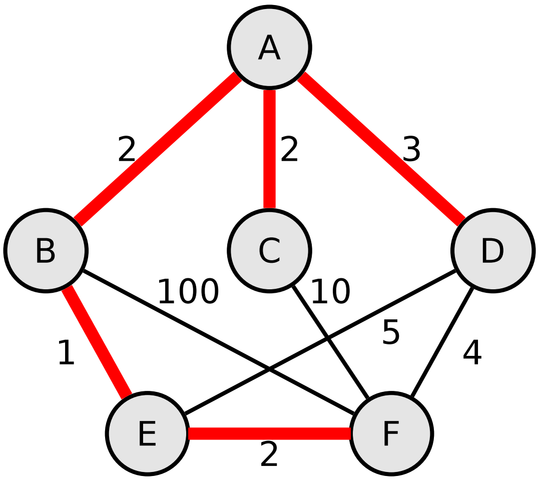

	Distance values					
	Initial	Step 1	Step 2	Step 3	Step 4	Step 5
d_A	**0** ← min	$\underline{0}$ ✓	0 ✓	0 ✓	0 ✓	0 ✓
d_B	∞	**2** ← min	$\underline{2}$ ✓	2 ✓	2 ✓	2 ✓
d_C	∞	2	**2** ← min	$\underline{2}$ ✓	2 ✓	2 ✓
d_D	∞	3	3	**3** ← min	$\underline{3}$ ✓	3 ✓
d_E	∞	∞	3	3	**3** ← min	$\underline{3}$ ✓
d_F	∞	∞	102	12	7	**5** ← min

The active node is underlined.

Nodes visited = {A, B, C, D, E, F}

Shortest paths = {A, A⟶B, A⟶C, A⟶D, B⟶E, E⟶F}

2. Calculate the shortest path from node A to all other nodes using Dijkstra's algorithm.

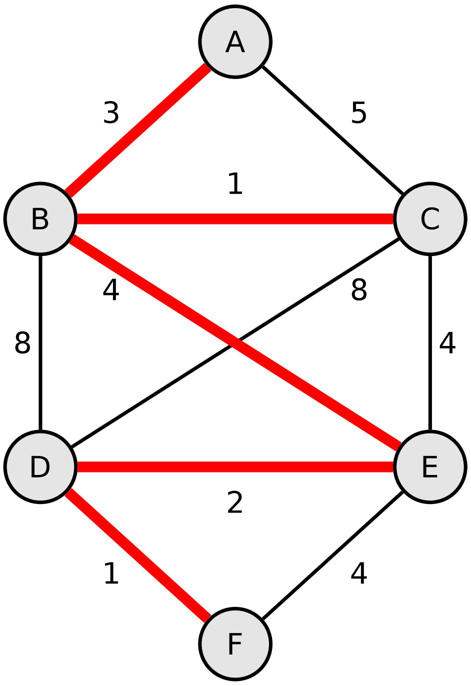

	Distance values					
	Initial	Step 1	Step 2	Step 3	Step 4	Step 5
d_A	**0** ← min	$\underline{0}$ ✓	0 ✓	0 ✓	0 ✓	0 ✓
d_B	∞	**3** ← min	$\underline{3}$ ✓	3 ✓	3 ✓	3 ✓
d_C	∞	5	**4** ← min	$\underline{4}$ ✓	4 ✓	4 ✓
d_D	∞	∞	11	11	**9** ← min	$\underline{9}$ ✓
d_E	∞	∞	7	**7** ← min	$\underline{7}$ ✓	7 ✓
d_F	∞	∞	∞	∞	11	**10** ← min

The active node is underlined.

Nodes visited = {A, B, C, E, D, F}

Shortest paths = {A, A⟶B, B⟶C, B⟶E, E⟶D, D⟶F}

Exercise 13 (Internet Control Message Protocol)

1. Explain the purpose of the Internet Control Message Protocol (ICMP).

   It is used for the exchange of diagnostic and control messages, as well as error messages.

2. Give two examples for command line tools, which use ICMP.

   ping, traceroute

Exercise 14 (IPv6)

1. Explain the concept of Scopes in IPv6.

   IPv6 not only distinguishes private and public addresses (like IPv4), but also several address scopes. Each IPv6 address has a so-called scope. The scope is the part of a network in which the associated address is considered valid and routed.

2. Explain what the Host Scope is.

   The Host Scope includes the loopback address `::1/128`.

3. Explain what the Link-Local Scope is.

   The Link-Local Scope includes Link-Local (Unicast) Addresses (LLA). Every network interface requires a Link-Local Address at any time. Link-Local Addresses fe80::/10 are only valid in the local network. Routers do not forward packages with these addresses.

4. Explain what the Unique-Local Scope is.

   The Unique-Local Scope includes Unique Local Addresses (ULA). Routers should not forward packages with these addresses outside the local administrative domain (organization or site). They are private addresses intended for local communication inside an administrative domain, but can be globally valid (unique) if they are assigned by a provider. Local generated ULA are very likely unique.

5. Explain what the Global Scope is.

   The Global Scope includes Global Unicast Addresses. Routers forward packages with these addresses.

6. Explain what the IPv6 address `::1/128` addresses.

   It is the loopback address.

7. Give the name of the scope of the IPv6 address `::1/128`.

   Host Scope.

8. Give the name of the scope of addresses that have the prefix `fe80::/10`.

   Link-Local Scope.

9. Give the name of the scope of addresses that have the prefix `fc00::/7`.

   Unique-Local Scope.

10. Give the name of the scope of addresses that have the prefix `2000::/3`.

    Global Scope.

11. IPv6 has no broadcast addresses but for some purposes, a broadcast-like functionality is required. Explain how IPv6 emulates the broadcast functionality.

    In IPv6, Multicast addresses are used to emulate the Broadcast functionality. The address `ff02::1` *has the Link-Local Scope and addresses all nodes in the local network.*

12. Give the prefix of Multicast addresses.

    Multicast addresses start with the first 8 bits set to value `11111111`*. Thus, they have the multicast prefix* `ff::/8`*.*

13. Name three ways of setting the Interface-ID.

    - *Static manual addressing*
    - *Stateless Address Autoconfiguration (SLAAC)*
    - *Setting the network configuration via DHCPv6*

14. Explain what Stable Privacy is and why it is used sometimes in the context of setting the Interface-ID.

    Stable Privacy is an optional extension of SLAAC (Stateless Address Autoconfiguration). Specifies the address generation without using a MAC address. A random secret key is generated and used for Interface ID generation. The secret key is a 128-bit long hexadecimal string that looks like an IPv6 address.

    One benefit (compared to SLAAC) is improved security because no MAC address is used for generation. The MAC address of the node is not exposed. This allows anonymity.

    Another benefit (compared to Privacy Extension) is that the address for the node is stable. Once generated, the Interface ID does not change anymore until reboot.

15. Explain what Privacy Extension is and why it is used sometimes in the context of setting the Interface-ID.

    Privacy Extension is another optional extension of SLAAC (Stateless Address Autoconfiguration). It uses the Interface-ID in a temporary manner. A new Interface-ID gets generated periodically. Old Interface-IDs can still be used for established connections.

    One benefit (compared to SLAAC) is improved security because no MAC address is used for generation. The MAC address of the node is not exposed. This allows anonymity. And because a new Interface-ID is generated periodically, the level of anonymity is even better compared to Stable Privacy.

    One drawback is that the address expires. It is not stable.

16. If a node has created an Interface-ID via SLAAC, it must validate that no other node in the network has the same Interface-ID. Explain how this is done in practice.

If a node has generated an IPv6 address for itself, it needs to validate that no other node in the network already uses this address. This procedure is called Duplicate Address Detection (DAD). The node sends a Neighbor Solicitation (NS) message to the address that it wants to use itself. Sender Address is the unspecific address (`::` $\Longrightarrow$ 128 zero bits). If a node in the local network already uses this IP address, it is a duplicate. The node will reply with a Neighbor Advertisement (NA) message sent to the Link-Local multicast address `FF02::1` (every node in the local network will receive this message). The node that was sending the Neighbor Solicitation (NS) message needs to generate a new address and carry out the Duplicate Address Detection again. If no Neighbor Advertisement (NA) message is received for some time, the address can be used ($\Longrightarrow$ no duplicate).

17. Give a short explanation for a concrete use-case of the ICMPv6 message Router Advertisement (RA) in practice.

    Routers periodically send (the time can be set in the UI) Router Advertisement (RA) messages into connected networks to inform others about their presence, the Network Prefix, the Prefix length, and, i.a., the MTU. Destination address in the IPv6 package is the Link-Local multicast address `FF02::1` to reach all nodes in the local network.

18. Give a short explanation for a concrete use-case of the ICMPv6 message Router Solicitation (RS) in practice.

    If a node does not want to wait for incoming Router Advertisement (RA) messages, it can request RA messages by sending RS messages. Destination address in the IPv6 package is the Link-Local multicast address `FF02::2` to reach all Routers in the local network.

19. Give a short explanation for a concrete use-case of the ICMPv6 message Neighbor Solicitation (NS) in practice.

    The Neighbor Solicitation (NS) message is the IPv6 alternative to an ARP Request when using IPv4. It is used to request the MAC address of a neighbor.

20. Give a short explanation for a concrete use-case of the ICMPv6 message Neighbor Advertisement (NA) in practice.

    The Neighbor Advertisement (NA) message is the IPv6-alternative to an ARP Reply when using IPv4. It is used to reply to a Neighbor Solicitation (NS) message (send the MAC address of an IPv6 address).

21. Give an explanation how a node learns if it is supposed to use a DHCPv6 server for requesting a address configuration (stateful address configuration) or if it is allowed to create an Interface-ID by itself (stateless address configuration).

    The periodically send Router Advertisement (RA) message i.a., includes the flag `managed`. If it is set, the client is supposed not to set the address stateless but to request the address configuration from a DHCPv6 server (stateful).

Exercise 15 (IPv6 – Simplification)

1. Simplify these IPv6 addresses:

   - `1080:0000:0000:0000:0007:0700:0003:316b`

     Solution: `1080::7:700:3:316b`

   - `2001:0db8:0000:0000:f065:00ff:0000:03ec`

     Solution: `2001:db8::f065:ff:0:3ec`

   - `2001:0db8:3c4d:0016:0000:0000:2a3f:2a4d`

     Solution: `2001:db8:3c4d:16::2a3f:2a4d`

   - `2001:0c60:f0a1:0000:0000:0000:0000:0001`

     Solution: `2001:c60:f0a1::1`

   - `2111:00ab:0000:0004:0000:0000:0000:1234`

     Solution: `2111:ab:0:4::1234`

2. Provide all positions of these simplified IPv6 addresses:

   - `2001::2:0:0:1`

     Solution: `2001:0000:0000:0000:0002:0000:0000:0001`

   - `2001:db8:0:c::1c`

     Solution: `2001:0db8:0000:000c:0000:0000:0000:001c`

   - `1080::9956:0:0:234`

     Solution: `1080:0000:0000:0000:9956:0000:0000:0234`

   - `2001:638:208:ef34::91ff:0:5424`

     Solution: `2001:0638:0208:ef34:0000:91ff:0000:5424`

   - `2001:0:85a4::4a1e:370:7112`

     Solution: `2001:0000:85a4:0000:0000:4a1e:0370:7112`