

Answers provided by Juha Takkinen, IDA, juha.takkinen@liu.se.

Question 1.

a)

Protocol = a set of rules for what messages to send (syntax), in what order (timing), and what they mean (semantics).

Network architecture = an architecture for how to organize protocols that provide network functionality to an application, usually consisting of layers, where a layer provides services to the layer above it.

(Reference model for network architectures = standardized network architecture, for example the OSI model with seven layers or the Internet model with four layers)

(Network design = how to organize the network equipment, such as computers, hubs, switches, and routers, so that it fulfills the performance requirements of the network users and applications)

b) The sender has 4 states (two for each sequence number, of which one is for DATA and the other one for ACK) while the receiver has 2 states (one for each sequence number).

The receiver does nothing when a packet is lost; it just ACKs all packets that are received.

Question 2.

[the question does not have any b) part]

a) Assume no other delays than those stated (RTT) and headers included in calculations.

Also assume that the ACK has no transmission time, only propagation.

i) The total transfer time = (handshake time) + (transfer of packet A--R1) + (transfer of fragments R1--B) + (transfer of ACK B--A) =

$$= (2 \times \text{RTT}) + ((1000 \times 8 \times 10^3) / (5 \times 10^6)) + (2 \times (250 \times 8 \times 10^3) / (5 \times 10^6) + 3$$

$$\times (250 \times 8 \times 10^3) / (5 \times 10^6)) + (\text{RTT} / 2) =$$

$$= 0.2 \text{ s} + 1.6 \text{ s} + 2 \text{ s} + 0.05 \text{ s} = 3.85 \text{ s}.$$

(note that the routers are performing work in parallel when data is divided into packets)

ii)

The average throughput = data amount / total transfer time = $(1000 \times 8 \times 10^3) / 3.85$
bps = 2078×10^3 bps = 260 KB/s

c)

The sender will move its window only one step to the right, covering positions 1-5, when it receives the four remaining ACKs because of the missing ACK for packet no. 1. The retransmission timers for the other packets will be stopped.

The retransmission timer for packet 1 will count down and then the packet will be retransmitted by the sender. The receiver will retransmit the ACK for packet 1. When the ACK for packet 1 is received by the sender, the sender can move the window to positions 5-10.

Question 3.

a) [For clarity's sake, the whole session is given below, with the lines from the question inserted with their number]

```
telnet smtp.liu.se 25
Trying 130.236.230.205...
Connected to smtp.liu.se.
Escape character is '^'.
ix) 220 munin.unit.liu.se ESMTP Postfix [the smtp server
presents itself and the protocol version]
ii) helo liu.se [the smtp client and its domain present itself
to the server]
x) 250 munin.unit.liu.se [the server acknowledges the helo command]
viii) mail from:<juha.takkinen@liu.se> [the client lists the
sender]
x) 250 Ok [the server acknowledges the mail command]
v) rcpt to:<lena.stromback@liu.se> [the client lists the
receiver]
x) 250 Ok
i) data [start of data]
354 End data with <CR><LF>.<CR><LF>
iii) mime-version: 1.0 [the message will be encoded using mime
so that more than 7-bit ascii can be accommodated, if any]
vii)Happy holidays! [the message]
iv) . [period] [the message is finished and will now be sent]
x) 250 Ok: queued as AA756279A9F
vi) quit [quit the session; no more e-mails from this domain]
221 Bye
Connection to smtp.liu.se closed by foreign host.
```

%

b)

The DHT data structure is commonly used as an index for storing and looking up information about hosts and what data they have stored. Data is stored as <key, value> pairs in the nodes of the DHT. The key in the data pair is, for example, the name of the music file to be searched, and the value is the ip address of the host that has the mp3 file in question. Both the node ID and the data key are hashed using the same hash function. In this way a simple match can be performed on the data key and the node ID in order to locate the node that has the key and subsequently look up the value (ip address).

c) The root server knows about TLD servers, while the authoritative server knows about hosts owned by the organization where the authoritative server is situated, including mail drops, if any.

Question 4.

a)

The timeout value is based on the sampleRTT, estimatedRTT and devRTT and three different constants, all calculated using the exponential moving weighted average. [a list of the relevant formulas follows]

b)

single ACK = transmit the next segment in the sending window, which has the exact sequence number of the ACK

double ACK = update the counter for duplicate ACKs

triple duplicate ACK = reset the counter for duplicate ACKs and go to fast retransmit of the lost segment with the same sequence number as the ACK. Then, halve the sending rate by setting the congestion window to half of its value and the threshold to this halved value.

Question 5.

a)

Ip packets can get lost, i.e., dropped in a router because they have the wrong checksum or the outgoing queue is full (another reason can be that the TTL field has reached zero).

b)

When an ip packet is fragmented each fragment is sent as an individual ip packet. They are reassembled at the receiver, based on the information stored in the second word in the packet header: the ID field has the same value for every fragment that belong to the

same original packet; the MF flag is set to 1 if the fragment is not the last fragment; and the Offset field tells the receiver where in the original packet the fragment belongs.

c)

In order to locate the subnets, isolate each interface that has an ip address, i.e., the interfaces belonging to routers and the hosts (see figure below). Concludingly, there are three subnets in the network: two subnets with two addresses and one subnet with nine addresses (not including the switch). Host A is behind R2, host F is behind R1 and the rest of the hosts are connected to the subnet limited by R1, R2 and R3.

If the network grows with one-third of hosts then there will be a need for 3, 3 and 12 hosts in each respective subnet. However, since there are two reserved addresses in each subnet (the broadcast all-ones and "this network" all-zeros) the no. of addresses needed are 6, 6 and 14, respectively. The two smaller subnets would then need a subnet mask of 29 bits (increased with one bit from before the expansion) while the larger subnet will need 28 bits, i.e., 255.255.255.232 (which actually is the same subnet mask as before the expansion).

Question 6.

a)

i) Two approaches are possible. The first approach is that host A uses the ip address and the subnet mask to initially find out that host G is in a different subnet, so it sends the packet to R2 using the ip address. R2 will in turn use ARP (assuming it does not know about host G yet) and broadcast an ARP request with a query for the MAC address of G belonging to G's ip address that R2 got from host A. Host G will respond with a unicast ARP packet to R2 containing G's MAC address.

The second approach uses a proxy-ARP in R2, which means that host A sends the packet with G's ip address in a MAC frame directly addressed to router R2. R2 then performs the same operations as in the first approach above.

ii)

Because the switching table is empty and the three first communications introduce new nodes, the procedure for these three will all look the same with regard to the steps performed: the switch will store the sending node's MAC address and also the link in the switching table and then flood the frame on every other link except the incoming one. The receiver will pick up the frame and respond directly to the sender, which will cause the switch to discover both the receiver's MAC address and link and put this information into the switching table. In the last communication, between E and R1, router R1 is already known to the switch, so the switch will put the frame directly on to the correct link, after storing host E's MAC address and link in the switching table.

b)

The sender reserves the channel by sending an RTS to the receiver. The receiver responds with a CTS, which is heard by all stations within range of the receiver and which are also hidden from the sender. Then the sender can transmit its data, after which the receiver concludes with an ACK. The ACK is also heard by the hidden stations, which now are free to try and access the channel.

Question 7.

[the algorithm actually builds the routing table, which in turn is used to derive the forwarding table]

a)

N	B	C	D	E	F
A	1, A	-	-	-	-
AB	done	2, B	-	-	-
ABC	done	done	3, C	-	-
ABCD	done	done	done	4, D	10, D
ABCDE	done	done	done	done	5, E
ABCDEF	done	done	done	done	done

The forwarding table in A consists of columns for destination, nextHop and cost as follows:

A - -
 B B 1
 C B 2
 D B 3
 E B 4
 F B 5

b)

The inter-AS protocol is concerned only with reachability and implements an organization's policy about its network's visibility to the outside world, what and how much traffic from other companies to carry, etc. The BGP protocol implements this.

The intra-AS protocol is concerned with least-cost paths and varying requirements connected to different-sized organizations. Therefore, several routing protocols are available here, such as RIP, OSPF and others by Cisco for example.

The interplay between the two types of protocols listed above is based on distributing varying information to all routers within ASs about which outside networks (foreign ASs) that are reachable from each AS and where a next-hop foreign router is along a path of ASs.

c)

Routing loops is a problem that can occur in routing protocols, which means that the

same node occurs more than once along a path.
Question 8.

a)

The handshake is needed because SSL must negotiate several security protocols (for encryption and hashing) before data can be transferred. Also, the communicating nodes must authenticate themselves, at the same time preventing insertion and playback attacks by using nonces and MACs.

b)

The operations are true because the message m is encrypted. This is accomplished using a shared secret key (session key, K_s) randomly created by the sender. In order for the receiver to be able to decrypt the confidential message, the session key is appended (+) to the sent message and encrypted using the public key of the receiver (K_b^+).

c)

Authentication in WEP is done using a shared secret key. A nonce is encrypted with this key in order to prevent playback attacks. This action is performed by both the wireless station and the access point.