Computer Networks - Final Exam

December 23, 2016

Duration: 2:15 hours, closed book.

- This is a closed-book exam.
- Please write your answers on these sheets in a readable way, in English or in French.
- Please do **not** use a red pen.
- You can use extra sheets if necessary (don't forget to put your name on them).
- The total number of points is 100.
- This document contains 19 pages.
- Good luck!

Full Name (Nom et Prénom):

SCIPER No:

Division: ☐ Communication Systems ☐ Computer Science
☐ Other (mention it): . . . .

Year: ☐ Bachelor Year 2 ☐ Bachelor Year 3
☐ Other (mention it): . . . .

Problem 1 (10 points)

For each question, please circle a single best answer.

1. Routers operate at the:
   (a) Application layer
   (b) Transport layer
   (c) Network layer
   (d) Link layer

2. Circuit switching:
   (a) Can support a higher number of users than packet switching.
   (b) Is simpler to implement than packet switching.
   (c) Offers a more efficient use of resources than packet switching.
   (d) Offers a more predictable performance than packet switching.

3. A datagram with size 1250 bytes is sent over a link with rate 10 Mbps, a length of 200 km, and a signal propagation speed of 200,000 km/s. The time it takes to transfer the datagram to the other endpoint (from first bit sent to last bit received) is:
   (a) 1 ms
   (b) 1.125 ms
   (c) 2 ms
   (d) 4 ms

4. If we double the transmission rate of a link:
   (a) We double the link's propagation delay.
   (b) We reduce the link's propagation delay by half.
   (c) We reduce any packet's transmission delay over that link by half.
   (d) We do not affect any delay component.

5. The only type of delay experienced by a packet that can be zero is:
   (a) propagation delay
   (b) queuing delay
   (c) transmission delay
   (d) processing delay

6. BitTorrent uses DHTs to:
   (a) Search for content based on user-specified keywords (e.g. ”Game of Thrones”).
   (b) Search for the addresses of peers that store a specific content based on a content identifier.
   (c) Store content on many peers so that it remains available if the peers leave the network.
   (d) Cache content on many peers to reduce download times.
7. Distance-vector routing algorithms suffer from the following scaling problem as the total number of routers in the network grows:
   (a) Large storage requirements, since each router needs to store the entire network graph.
   (b) Large number of open connections, since each router needs to exchange messages with all the other routers in the network.
   (c) Long convergence time, since the number of iterations needed to reach convergence increases as the number of routers in the network increases.
   (d) There are actually no such scaling problems.
8. TCP offers reliable data delivery. Why does the link layer sometimes also offer ACKs and retransmissions?
   (a) To reduce TCP’s complexity.
   (b) To increase the probability that data is successfully delivered.
   (c) To reduce the time it takes for data to be successfully delivered.
   (d) For historical reasons – there is really no point in offering ACKs and retransmissions at two layers.
9. The subnet mask for the network 128.178.0.0/15 is:
   (a) 255.255.128.0
   (b) 255.255.0.0
   (c) 255.254.0.0
   (d) 128.178.0.0
10. Go-back-N offers the following advantage compared to Stop-and-wait:
   (a) Higher throughput when there is little or no packet loss.
   (b) Fewer packets need to be retransmitted in case of packet loss.
   (c) The receiver can handle reordered packets better.
   (d) More reliable data delivery.

Problem 2 **(33 points)**

Setup:

Consider the network in Figure 1, consisting of:

- Hosts A, B, C and D
- File server E, DNS server F, web server G
- Router and NAT gateway R_1
- Router R_2
- Switches S_1 and S_2

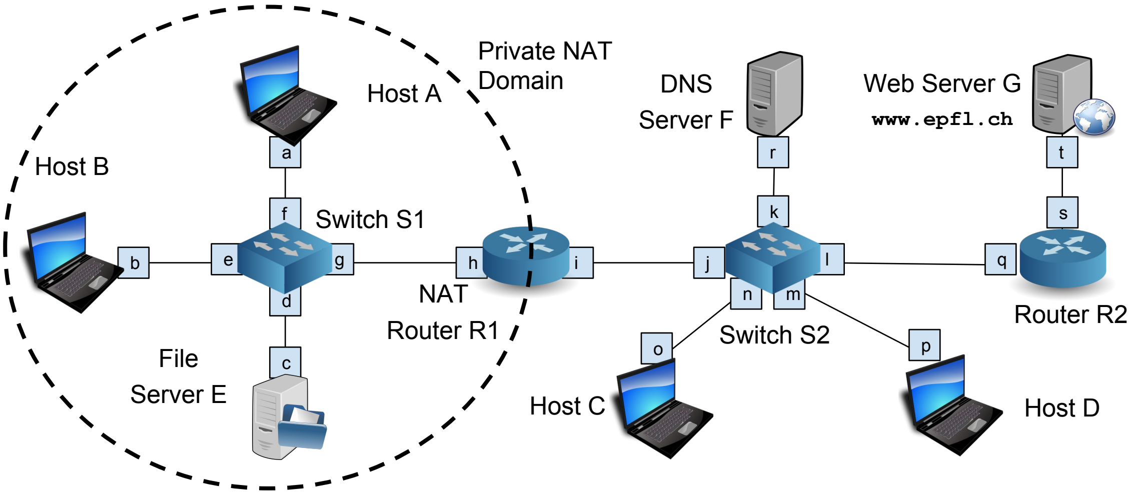

Figure 1: The Network Topology used in Section A

Interface h and all the interfaces to its left belong to a private NAT domain that uses R_1 as its NAT gateway.

Question 1 (10 points):

Allocate an IP prefix to each IP subnet and an IP address to each network interface that needs one, following these rules:

- All public addresses must be allocated from `2.2.2.0/24` (they should have the binary format `00000010.00000010.00000010.xxxxxxxx`).
- All private addresses must be allocated from `10.1.0.0/16` (they should have the binary format `00001010.00000001.xxxxxxxx.xxxxxxxx`).
- Each IP subnet must be allocated the smallest possible IP prefix.
- The first and last address in each IP subnet must be reserved for the network and broadcast address. These two addresses cannot be assigned to any interface.
- Router interfaces need IP addresses.

Use Table 1 to show the IP address allocated to each interface. Justify your answer on page 6.

Network interface	IP address
Example: x	1.2.3.4

Table 1: Allocation of IP addresses for the network interfaces in Figure 1

Question 2 (14 points):

A web browser running on Host A retrieves web page `index.html` from `www.epfl.ch` (web server G). Describe all the messages that are **received or forwarded by $\mathbf{R_1}$** until the browser displays the web page, by filling Table 2. Make the following assumptions:

- All hosts use DNS server F as their local DNS server.
- All caches on all devices are empty.
- `index.html` fits in one packet and does not refer to any other objects.

Column "Trans." refers to the transport-layer protocol. Column "App." refers to the application-layer protocol. When you want to refer to the IP address of interface x, write "IP of x". When you want to refer to the MAC address of interface x, write "MAC of x". If a field is not applicable, indicate that with a "–".

#	Src MAC	Dst MAC	Src IP	Dst IP	Trans.	Src Port	Dst Port	App.	Purpose
ex	MAC of x	MAC of y	IP of w	IP of v		5000	6000	HTTP	Request: `image.png`
1									
2									

Table 2: Messages received or forwarded by R_1, in Question 2

Question 3 (3 points):

Show the NAT translation table of R_1 right after Host A has retrieved `index.html`, by filling Table 3. Use the first row of the table to specify the meaning of each column.

Table 3: NAT translation table of router R_1

Question 4 (3 points):

Show the forwarding tables of switches S_1 and S_2 right after Host A has retrieved `index.html`, by filling Table 4. Use the first empty row of the table to specify the meaning of each column.

Switch $\mathbf{S_1}$		**Switch $\mathbf{S_2}$**	

Table 4: Forwarding tables of switches S_1 and S_2, in Question 3

Question 5 (3 points):

How can Host B prevent Host A from getting `index.html` by sending a single packet? When should Host B send that packet and what should it contain?

Problem 3 (24 points)

Question 1 (12 points):

In each of the following scenarios, Host A wants to communicate with Host B and achieve some security property. In each scenario:

i. Identify an existing problem or weakness and, if applicable, describe an attack that exploits it.

ii. Provide a solution that fixes the weakness (i.e. what should A send instead). Make sure to provide enough detail for your solution to be understandable. For example, if you say "A should use a MAC for authentication", but you do not explain how the MAC should be computed, then your answer is not complete.

Scenarios:

a. A wants to send one message m to B, ensuring confidentiality and authenticity. For this, A sends: $H(K, m)$.

b. A wants to send one message m to B, ensuring confidentiality and authenticity. For this, A sends: $K\{m\}$.

c. A wants to send one message m to B, ensuring confidentiality and authenticity. A knows B's true public key K_B^+. A sends: $K_B^+\{m,\ K_B^+\{H(m)\}\}$.

d. A is sending messages to B. Whenever B receives a message m, it should be able to verify that A indeed sent a message with the same content as m at least once. For this, A appends $H(K)$ to each message it sends.

e. A wants to send one message m to B, ensuring authenticity and data integrity. For this, A cuts m into two pieces, m_1 and m_2, sends first $[m_1,\ H(K)]$, then $[m_2,\ H(K)]$.

f. A and B are friends that want to have a sensitive online conversation, ensuring confidentiality, authenticity, and data integrity (A receives exactly the sequence of messages sent by B and vice versa). For this, they use SSL as described in class: B sends a nonce and a certificate with its public key to A, A sends a nonce and a master key to B (encrypted with B's public key), and from the master key and the nonces, they both derive other keys that they use for confidentiality and authenticity. Hint: In the example we saw in class, B was an online store. In this question, A and B are friends having a sensitive conversation.

where:

- H is a cryptographic hash function that is globally known to everyone.
- K is a symmetric key, shared between A and B.
- K_B^+ is B's public key.

(*Extra answer page for Question 1.*)

Question 2 (12 points):
Hosts A and F secure their communication with the following protocol:

1. A sends a "hello" message with a nonce n_A and a certificate containing its public key (K_A^+)
2. F responds with a nonce n_F and a certificate containing its public key (K_F^+)
3. After this exchange, to communicate a message m_i, A sends: $K_F^+ \left\{m_i,\ K_A^- \left\{H\left(n_F, m_i\right)\right\}\right\}$
4. Similarly, to communicate a message m_j, F sends: $K_A^+ \left\{m_j,\ K_F^- \left\{H\left(n_A, m_j\right)\right\}\right\}$,

where K_A^-, K_F^- are A's and F's private keys and H is a globally known cryptographic hash function.

Questions:

a. Does this protocol guarantee confidentiality?
If yes, explain why. If not, describe an attack and provide a solution.

b. Assume a man-in-the-middle, who records all the messages sent by A, and, when the communication between A and B ends, she resends them to F, trying to impersonate A. Will his attack be successful (or not) and why?

c. Is this protocol vulnerable to any attack(s) other than the ones described above?

If yes, briefly describe the attack(s) and provide the changes that make the protocol completely secure. If necessary, you can add new steps in the protocol, but **not** modify the existing ones.

Extra answer page for Question 2.

Problem 4 (33 points)

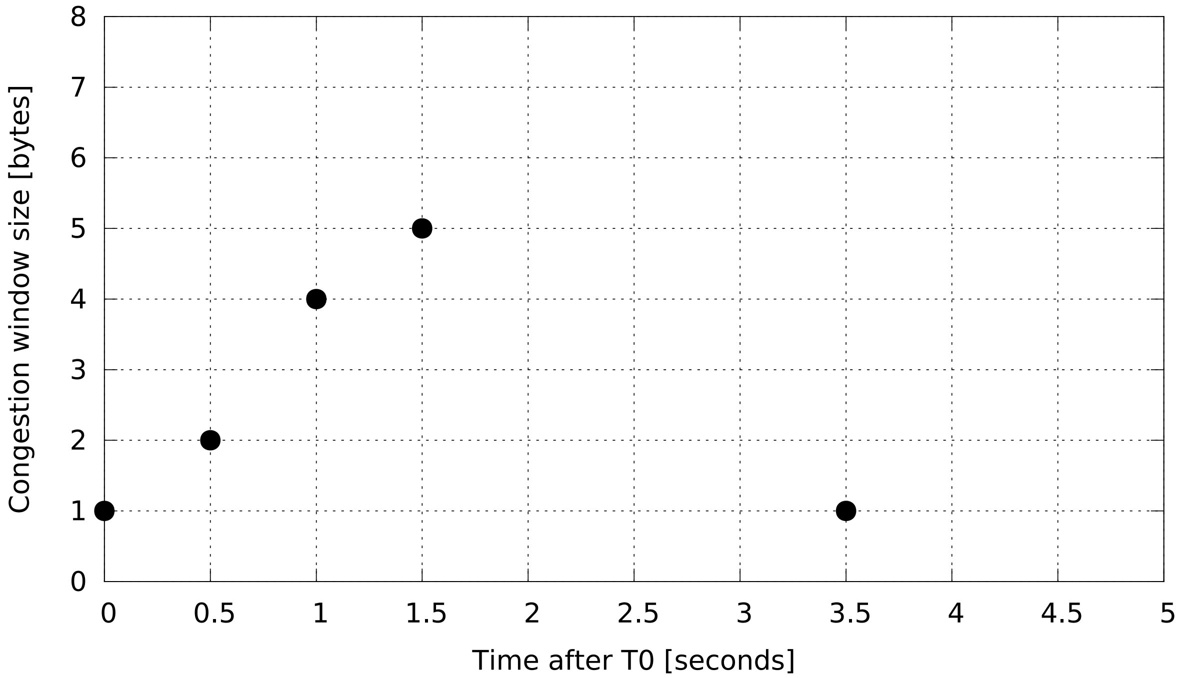

Figure 2: Congestion window of Alice over time

Alice has opened a persistent TCP connection to Bob. At time T_0, Alice starts sending to Bob, over this connection, a file of size 12 bytes in segments of MSS = 1 byte.

Figure 2 shows how the congestion window of Alice, $cwnd$, changes over time after T_0 and until the file transfer completes. Each of the $five$ points in the graph shows the time a change in $cwnd$ took place and $cwnd$'s value after the change.

Make the following assumptions:

- Transmission delay is negligible.
- Bob sends an ACK for each segment it receives.
- The first segment that Alice transmits after T_0 has sequence number 10.
- Fast-retransmit is disabled.
- Only one segment gets lost after T_0, and it is a segment sent by Alice.

Question 1 (15 points):

a. What is the RTT between Alice and Bob?

b. What is the retransmission timeout used by Alice?

c. What was the size of Alice's congestion window, $cwnd$, the last time a packet was lost before T_0?

d. Complete the diagram on the next page that shows what happens after T_0 and until the file transfer completes:

- All segments exchanged between Alice and Bob.
- The sequence numbers sent by Alice and the acknowledgment numbers sent by Bob.
- The state of Alice's congestion-control algorithm.
- The size of Alice's congestion window, `cwnd`, in bytes.
- The value of Alice's slow-start threshold, `ssthresh`, in bytes.

cwnd [bytes]	ssthresh [bytes]	State of the congestion control algorithm for Alice	Sequence number diagram	
			Sequence number	Acknowledgement number
			Alice	**Bob**

Figure 3: Sequence diagram to be completed for Question 1

(*Extra answer page for Question 1.*)

Question 2 (6 points):

a. How long does the file transfer take? Assume that the file transfer completes once Alice has received the final ACK for file data.

b. Now assume (just for this question, i.e., Question 2b) that fast-retransmit is enabled. Does this change the duration of the file transfer and how/why?

Question 3 (7 points):

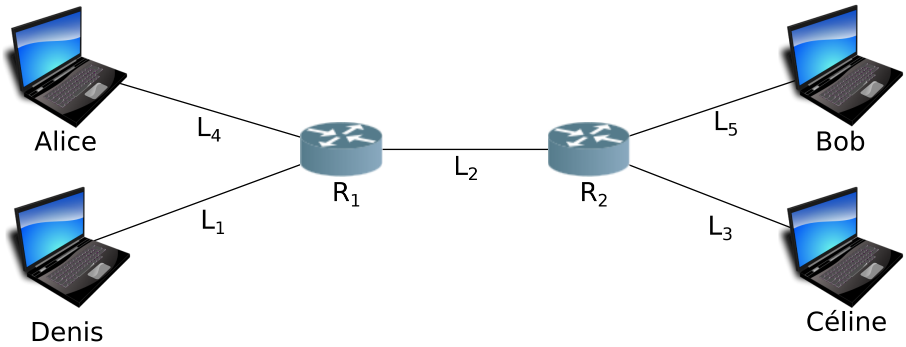

Figure 4: The Network Topology used in Question 3

In Figure 4, Alice is sending a large file to Bob using TCP. Denis tries to disrupt their communication by sending traffic to Céline. No other hosts send any traffic.

a. Describe the simplest attack strategy that achieves Denis's goal. What condition needs to hold for the transfer rates of the links such that this strategy works?

b. How will the TCP connection between Alice and Bob be affected by this attack? Draw a simple diagram that shows how Alice's congestion window, $cwnd$, evolves over time during the attack. You do not need to provide specific time values on the x-axis, just show the trend (e.g., does $cwnd$ increase monotonically?)

Question 4 (5 points):

Describe the attack strategy that achieves Denis's goal while minimizing the amount of traffic that Denis sends to Céline.

Hint: Denis does not need to send traffic at a constant rate.