ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Computer Networks - Midterm

October 31, 2014

Duration: 2:15 hours.

- This is a closed-book exam.
- Please write your answers on these sheets in a readable way, in English or in French.
- 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 17 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): . . . .

1 Short questions (10 points)

For each question, please circle a single best answer.

1. In a circuit switched network we can have
   (a) data loss, unpredictable delay.
   (b) data loss, predictable delay.
   (c) no data loss, unpredictable delay.
   (d) no data loss, predictable delay.

2. Alice is connected to the Internet using a 100Mbps connection, Bob is connected to the Internet using a 10Mbps connection. When Alice sends a 100Mb file to Bob, the transfer will take
   (a) less than 10 seconds.
   (b) exactly 10 seconds.
   (c) more than 10 seconds.
   (d) we cannot say precisely, all the above can be true.

3. Consider a switch with a buffer size of 20Mb and its output link with transmission rate of 10Mbps. The switch may introduce a maximum queuing delay of
   (a) 0 seconds.
   (b) less than 2 seconds.
   (c) more than 4 seconds.
   (d) ∞

4. A host uses HTTP to access a web page that consists of a base file and one picture, both stored on the same web server. How many round-trip-times (RTTs) are required to retrieve the entire web page?
   (a) exactly 2 RTTs
   (b) at most 2 RTTs
   (c) at least 3 RTTs
   (d) at most 3 RTTs

5. The following constitutes an attack against the DNS protocol:
   (a) impersonate the local DNS server.
   (b) denial-of-service against the root or TLD servers.
   (c) poison the cache of the local DNS server.
   (d) all of the above.

6. A peer-to-peer application
   (a) always uses UDP, because we don't have dedicated infrastructure for the server.
   (b) always uses TCP, because peers act both as clients and servers.
   (c) may use TCP or UDP as a transport layer protocol.
   (d) doesn't use any transport layer protocol.

7. A valid set of actions performed by a TCP client is:
   (a) create socket, wait for connection setup, receive service request, send reply, close the connection.
   (b) create socket, initiate connection setup, send service request, receive reply.
   (c) create socket, send service request, receive reply, close the connection.
   (d) create socket, send service request, receive reply.

8. A web server that implements persistent HTTP communicates using
   (a) a different socket for each request.
   (b) the same socket for all requests.
   (c) a different socket for each connecting client.
   (d) a different socket for each web page served.

9. For a Go-Back-N protocol, if the window size of the sender is N, then the window size of the receiver is
   (a) 1.
   (b) N.
   (c) $\frac{N}{2}$.
   (d) any value smaller than N.

10. An application uses UDP as a transport layer. If the application wants to have reliable transfer it can implement
   (a) a stop-and-wait protocol.
   (b) a selective repeat protocol.
   (c) a go-back-N protocol.
   (d) any of the above.

2 Application layer (30 points)

Setup:

Three users are logged into the workstations `user1.epfl.ch`, `user2.epfl.ch` and `user3.epfl.ch`, all located inside EPFL's network.

EPFL offers a web server `www.epfl.ch`, a local name server `ns.epfl.ch`, and a web proxy `proxy.epfl.ch`.

A web server `www.example.com` is located outside EPFL's network.

Make the following assumptions:

- All the computers inside EPFL's network use `ns.epfl.ch` as their local DNS server.
- `ns.epfl.ch` is also the *authoritative* DNS server for the `epfl.ch` domain.
- All DNS queries are resolved *iteratively*.
- All DNS servers and workstations maintain DNS caches.
- Web browsers and web proxies perform caching.
- All caches are initially empty. This is true for each question.
- Application layer messages fit in one packet.
- For simplicity, assume nobody else is generating traffic on the Internet.

Question 1 (10 points):

User 1 browses `http://www.epfl.ch/index.html`. The only object referenced by the base file is the image `http://www.epfl.ch/image.png`. The web browser uses a *single non-persistent* connection.

List all the packets that are exchanged, including any connection setup packets. Show the source and destination of each packet, the application protocol, the transport protocol, and the purpose of the packet. Fill in the answer in Table 1.

Packet	Source	Destination	Application protocol	Transport protocol	Message
1					
2					

Table 1: Packets transmitted on the Internet, in Question 1.

Question 2 (12 points):

User 2 browses `http://www.epfl.ch/help.html` through the web proxy. This web page does not reference any objects. The browser uses a *single non-persistent* connection.

List all the packets that are exchanged, including any connection setup packets. Show the source and destination of each packet, the application protocol, the transport protocol, and the purpose of the packet. Fill in the answer in Table 2.

Packet	Source	Destination	Application protocol	Transport protocol	Message
1					
2					

Table 2: Packets transmitted on the Internet, in Question 2.

Question 3 (4 points):

User 1 is about to browse `http://www.example.com/`.

User 3 knows this and wants to intercept User 1's web traffic. For this purpose, User 3 wants User 1 to connect to `user3.epfl.ch` instead of `www.example.com`.

Explain how User 3 can achieve this and describe the DNS messages that are exchanged.

Question 4 (4 points):

User 3 wants to intercept the web traffic between all the computers inside EPFL (not just User 1) and the website `http://www.example.com/`.

For this purpose, User 3 wants them to connect to `user3.epfl.ch` instead of `www.example.com`.

Explain how User 3 can achieve this, and describe the DNS messages that are exchanged.

3 Network delays (30 points)

Suppose the following:

- End-nodes $1, 2, 3, 4, ..., N$ ($N \geq 5$) are part of a network, shown in Figure 1. The nodes are connected via switches S_1 and S_2, which perform packet switching using *store-and-forward*.
- The size of the packet buffer at every switch is infinite and processing delays can be ignored.
- Each of the physical links (solid lines) has length ℓ meters and propagation speed c.
- The transmission rate of the links that connect the end-hosts to the switches is R (in both directions).
- The transmission rate of the link that connects S_1 and S_2 is $\frac{R}{k}$, $k > 1$ (in both directions).
- All nodes are organized in a *one-way* circular DHT (shown with dashed lines in Fig. 1). Each node is aware only of its subsequent neighbor; e.g., Node 2 knows only about Node 3.
- The one-way circular DHT is implemented on top of UDP.
- Node 1 stores a *movie* of size F bits and Node 2 is responsible for storing information about this movie.

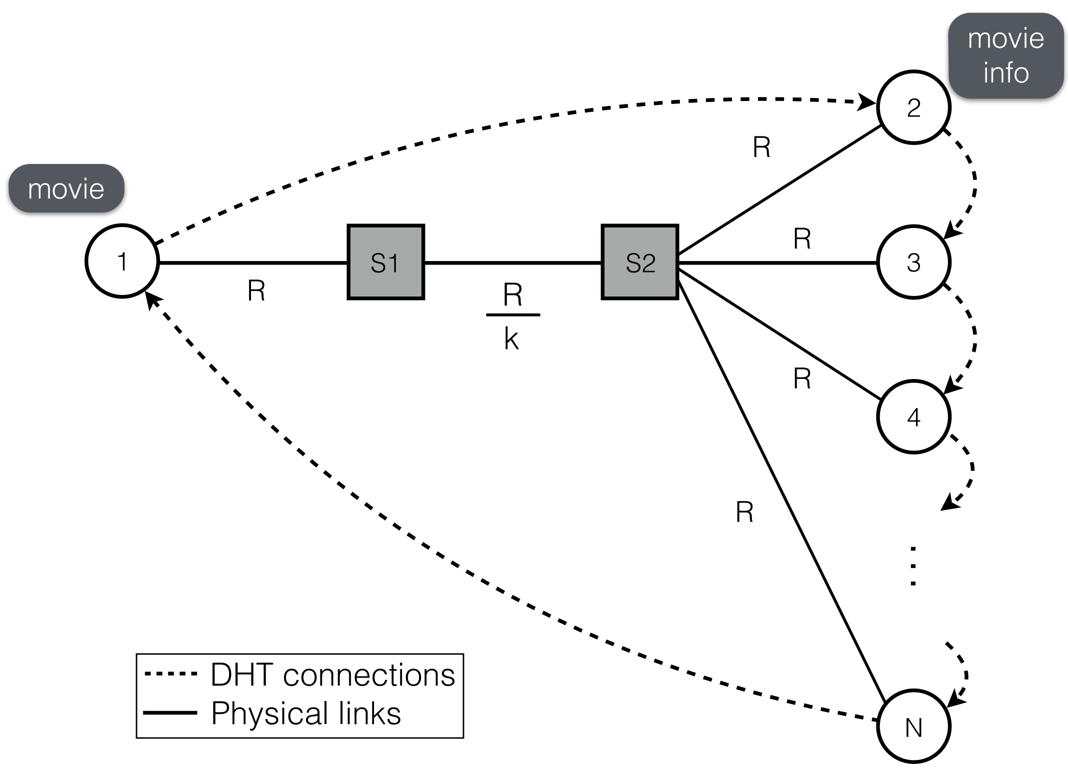

Figure 1: The setup for the exercise on network delays.

Question 1 (5 points):

Suppose that any DHT query/request message fits into one packet of size q bits. Compute the total time that a node i needs to learn where the movie is stored, for the following cases:

a. $i = N$

b. $i = 3$

Question 2 (3 points):

What mechanism would you use to reduce the average delay experienced by the nodes in learning where the movie is stored? Give a mechanism example and explain how would it affect the delay experienced by Node 3?

Question 3 (12 points):

After querying the DHT, all Nodes $2, \ldots, N$ are aware of where the movie is stored and they request it from Node 1. We are at the point where Node 1 has collected all the requests. Now, Node 1 splits the movie into P packets and then starts sending the packets using the *Client-Server* approach.

For this question, you may ignore propagation delay, but no further approximations are accepted.

a. Suppose all P packets have the same size.

   i. How long will it take for the movie to be delivered only to one node (e.g. Node 2)?

   ii. How long will it take for the movie to be delivered to *all* nodes $2, \ldots, N$?

b. Suppose now that all P packets have random sizes. What would be the best sending strategy for Node 1 to improve its delay performance: send them in a descending order (i.e. send the biggest packet first) or an ascending order (i.e. send the smallest packet first)? Does it make a difference? Justify your answer.

Question 4 (10 points):

Consider now that Node 1 uses the *P2P approach*. It, again, splits the movie into P equally sized packets and distributes the packets with the help of its peers.

Assume the following redistribution scheme: Node 1 sends packets to Node 2 only. Upon receiving a packet, Node 2 redistributes it to its subsequent neighbor (Node 3). Similarly, when the i^{th} node receives a packet it redistributes it *only* to the $(i+1)^{th}$ node. When a node has sent the last packet of the file it stops redistributing.

For this question, you may ignore propagation delay, but no further approximations are accepted.

a. How long will it take for the movie to be delivered to all the nodes?

b. Is there a value k for which the Client-Server approach has better performance than the P2P approach? Justify your answer.

4 Transport layer (30 points)

Question 1 (10 points):

End-host A communicates with end-host B using a Go-Back-N protocol with sender window size $N = 3$. The link may drop packets, but it can neither reorder nor corrupt packets. A is trying to send a file to B. It does so by splitting the file in 8 packets with sequence numbers from 0 to 7.

In the entire duration of the file transfer, end-host A receives *only* the following ACKs (and in that order): $ACK1$, $ACK1$, $ACK4$, $ACK6$, $ACK7$

Show which packets have been transmitted by A and B, including the ones which have been dropped. Answer this question by completing the sequence diagram in Figure 2.

Your answer should include the following:

- All packet transmissions, including ACKs.
- Lost packets and ACKs.
- Timeout events.
- Changes to the sender and receiver window, when packets (including ACKs) are received by either end-host.

We have already completed some of the information to help you get started. End-host A has sent data packets 0 and 1, which have reached end-host B. B sends $ACK0$ and $ACK1$, but only $ACK1$ reaches A.

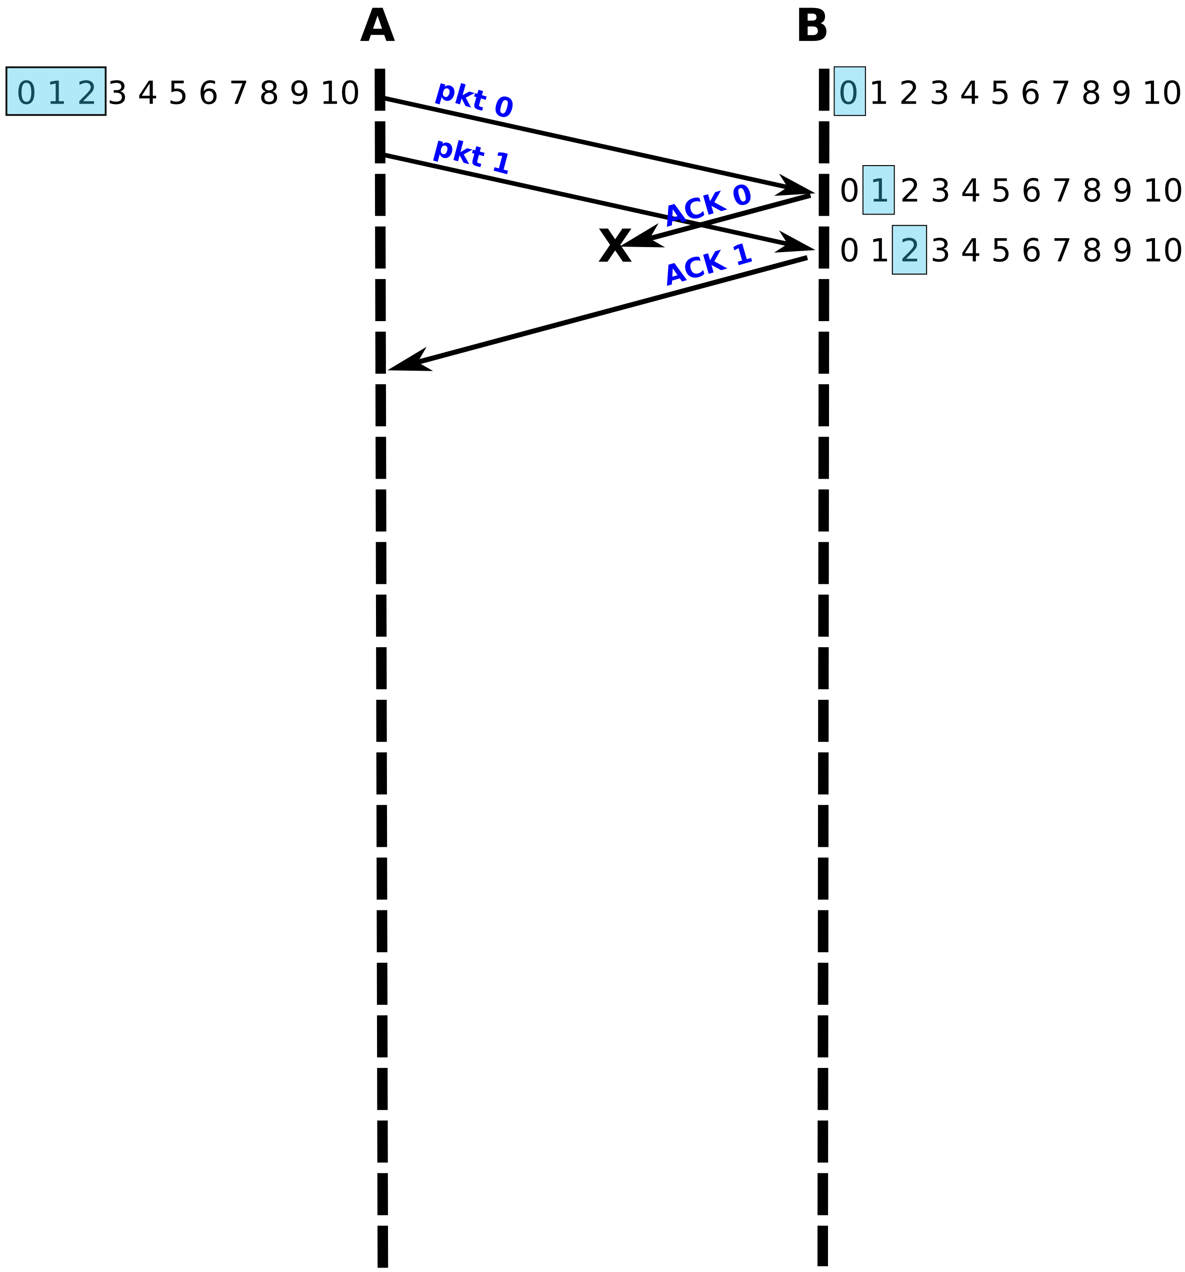

Figure 2: Sequence diagram to be completed for Question 1.

Question 2 (20 points):

Table 3 describes 5 different network channels between a sender A and a receiver B. E.g., in channel (3), there may be packet corruption only on the path from A to B, and we need to do pipelining. There is no packet reordering in any of the channels.

channel ID	errors in $A \rightarrow B$	errors in $B \rightarrow A$	support for pipelining?
(1)	none	none	yes
(2)	corruption	none	no
(3)	corruption	none	yes
(4)	corruption &loss	none	yes
(5)	corruption & loss	corruption & loss	yes

Table 3: Network channel types for Question 2.

For each channel, specify which of the mechanisms in Table 4 are necessary for providing reliable data delivery (no data loss and no data duplication).

You have the following constraints:

- Specify a mechanism only if it is necessary. E.g., if one can provide reliable data delivery over a given channel without checksums, you should not specify checksums for that channel.
- Do not specify timeout-based retransmissions if NACK-based retransmissions are sufficient.

Fill in Table 4 by writing "yes" or "no" in each cell. Then justify your choice for each channel by explaining why the mechanisms you picked are necessary and sufficient.

channel ID	sequence numbers	checksums	NACK-based retransmissions	timeout-based retransmissions
(1)				
(2)				
(3)				
(4)				
(5)				

Table 4: Reliability mechanisms to be selected for Question 2.