



# Midterm - Computer Networks

Prof. J.-P. Hubaux and Dr. M. H. Manshaei

November 3, 2009

Duration: 1:45 hours, closed book.

Please write your answers on these sheets, at the end of each question;  
use extra sheets if necessary (put your name on them).

You may write your answers in English or in French.

The total number of points is 40.

This document contains 19 pages.

**Student First name:**

**Last name:**

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

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

*(answers to the questions are shown in italic and blue)*

## 1 Short questions

(5 points)

*For each question, please circle a single best answer.*

1. What advantage does a circuit-switched network have over a packet-switched network?
  - (a) It can support wireless connections.
  - (b) It can be implemented over fiber-optic link.
  - (c) It can provide the same services with less overhead.
  - (d) It can guarantee a certain amount of end-to-end bandwidth. *CORRECT*
  - (e) All of the above.
  
2. A tier-1 ISP:
  - (a) is connected to all other tier-1 ISPs. *CORRECT*
  - (b) is connected to only a few of tier-1 ISPs.
  - (c) is a customer of one or more tier-2 ISPs.
  - (d) Both (a) and (c)
  
3. What information is used by a process running on one host to identify a process running on another host?
  - (a) The data rate between the source and the destination in bits/sec.
  - (b) IP address of the destination host.
  - (c) Port number of the destination socket.
  - (d) Both (b) and (c). *CORRECT*
  - (e) All of the above.
  
4. FTP sends control information *out of band*. This means that:
  - (a) FTP usually uses more bandwidth than HTTP or SMTP.
  - (b) FTP uses two UDP connections simultaneously.
  - (c) FTP uses two TCP connections simultaneously. *CORRECT*
  - (d) Both (a) and (b).
  - (e) Both (a) and (c).

5. Suppose that you want to do a transaction from a remote host to a server as fast as possible. Which transport protocol would you use?
- (a) FTP.
  - (b) HTTP.
  - (c) UDP. *CORRECT*
  - (d) TCP.
6. With the *Selective Repeat* protocol:
- (a) the receiver sends individual ACK packets. *CORRECT*
  - (b) the sender has a timer for the oldest unacked packet.
  - (c) the receiver sends cumulative ACK packets.
  - (d) Both (a) and (b).
7. Which parameter is used to calculate the timeout interval in TCP?
- (a) Sample Round Trip Time (RTT).
  - (b) RTT variation.
  - (c) Average congestion window size.
  - (d) Both (a) and (b). *CORRECT*
8. After a loss event detected by triple duplicate ACK, TCP Reno
- (a) cuts its congestion window to 1 Maximum Segment Size (MSS).
  - (b) goes to fast recovery phase. *CORRECT*
  - (c) goes to slow-start phase.
  - (d) Both (a) and (c).
9. Consider a server socket object: `socket = ServerSocket(12345);`  
What does the invocation of `socket.accept()` return?
- (a) `true` if there is a new TCP segment in the socket's buffer; `false` otherwise.
  - (b) A new TCP segment from the socket's buffer (blocks if no segments are available).
  - (c) `true` if a new TCP connection request has arrived; `false` otherwise
  - (d) A socket associated with a new TCP connection (blocks if no connections are available).  
*CORRECT*
10. What information a TCP socket does **not** need to keep track of?
- (a) Local port.
  - (b) Number of open connections. *CORRECT*
  - (c) Sequence numbers.
  - (d) Congestion window size.

## 2 Network Delays

(8 points)

**Question 1:** Consider two hosts, A and B, connected by a single link of rate  $R$  bps. Suppose that the two hosts are separated by  $D$  meters, and suppose the propagation speed along the link is  $V$  meters/sec. Host A is to send a packet of size  $L$  bits to host B.

- a. Express the propagation delay,  $d_{prop}$ , in terms of  $D$  and  $V$ .

$$d_{prop} = D/V \text{ seconds.}$$

- b. Determine the transmission time of the packet,  $d_{trans}$ , in terms of  $L$  and  $R$ .

$$d_{trans} = L/R \text{ seconds.}$$

- c. Ignoring processing and queuing delays, obtain an expression for the end-to-end delay.

$$d_{end-to-end} = (D/V + L/R) \text{ seconds.}$$

- d. Suppose  $V = 2.5 \times 10^8$ ,  $L = 120$  bits, and  $R = 56$  Kbps. Find the distance  $D$  so that  $d_{prop}$  equals  $d_{trans}$ . Assume that 1 Kbps = 1,000 bps.

$$D = V \frac{L}{R} = (2.5 \times 10^8) \frac{120}{56 \times 10^3} = 536 \text{ km.}$$

**Question 2:** Consider two hosts, A and B, separated by three links. These three links are connected by two packet switches. Let  $d_i$ ,  $s_i$ , and  $R_i$  denote the length, propagation speed, and the transmission rate of link  $i$ , for  $i = 1, 2, 3$ . Each packet switch introduces a constant processing delay  $d_{proc}$  for each packet. Consider sending a packet, of length  $L$ , from A to B.

- a. Assuming no queuing delays, what is the total end-to-end delay for the packet, in terms of  $d_i$ ,  $s_i$ ,  $R_i$ , ( $i = 1, 2, 3$ ),  $L$ , and  $d_{proc}$ ?

*The first end system requires  $L/R_1$  to transmit the packet onto the first link; the packet propagates over the first link in  $d_1/s_1$ ; the packet switch adds a processing delay of  $d_{proc}$ ; after receiving the entire packet, the packet switch connecting the first and the second link requires  $L/R_2$  to transmit the packet onto the second link; the packet propagates over the second link in  $d_2/s_2$ . Similarly, we can find the delay caused by the second switch and the third link:  $L/R_3$ ,  $d_{proc}$ , and  $d_3/s_3$ . Adding these delays gives  $d_{end-to-end} = L/R_1 + L/R_2 + L/R_3 + d_1/s_1 + d_2/s_2 + d_3/s_3 + d_{proc} + d_{proc}$ .*

- b. Suppose now the packet consists of 1,500 bytes, the propagation speed on all three links is  $2.5 \times 10^8$  m/s, the transmission rates of all three links are 2 Mbps, the packet switch processing delay is 3 msec, the length of the first link is 5,000 km, the length of the second link is 4,000 km, and the length of the last link is 1,000 km. For these values, what is the end-to-end delay? Assume that 1 Mbps = 1,000 Kbps = 1,000,000 bps.

*We simply plug the values into the above equation to get  $6 + 6 + 6 + 20 + 16 + 4 + 3 + 3 = 64$  msec.*

**Question 3:** Consider a router buffer preceding an outbound link. In this problem, you will use Little's formula, a famous formula from queueing theory. Let  $N$  denote the average number of packets in the buffer plus the packet being transmitted. Let  $a$  denote the rate of packets arriving at the link. Let  $d$  denote the average total delay (i.e., the queueing delay plus the transmission delay) experienced by a packet. Little's formula is  $N = a \times d$ . Suppose that on average, the buffer contains 10 packets, and the average packet queueing delay is 10 msec. The link's transmission rate is 100 packets/sec. Using Little's formula, compute the average packet arrival rate, assuming there is no packet loss.

*The total number of packets in the system includes those in the buffer and the packet being transmitted.  
 $N = 10 + 1 = a \times d \Rightarrow 11 = a \times (d_{\text{queueing}} + d_{\text{trans}}) \Rightarrow 11 = a \times (0.01 + 1/100) = a \times (0.01 + 0.01) \Rightarrow a = 550$  packets/sec.*

**Question 4:** Suppose two hosts, A and B, are separated by 15,000 km and are connected by a direct link of rate  $R = 3$  Mbps. Suppose the propagation speed over the link is  $2.5 \times 10^8$  m/s. Consider sending a file of 700 Mbytes from A to B. Suppose the file is sent continuously as one large message. What is the maximum number of bits that will be in the link at any given time? Assume that 1 Mbps = 1,000,000 bps and 1 Mbyte =  $(1,024)^2$  bytes.

*The maximum number of bits is  $R \times d_{prop} = 180,000$  bits.*

### 3 Packet Segmentation

(5 points)

**Question 1:** Consider sending a large file of  $F$  bits from host  $A$  to host  $B$ .  $A$  and  $B$  are connected by a path of  $\ell$  links (as shown in Figure 1), and links are uncongested (i.e., no queuing delays). Each link has a transmission rate of  $R$  bps. Host  $A$  divides the file into segments of  $S$  bits (we assume that  $F = k \times S$ ). Host  $A$  then adds  $H$  bits of header to each segment and makes the packets of  $S + H$  bits. Find the value of  $S$  that minimizes the delay of sending the file from host  $A$  to host  $B$ . In this problem we neglect propagation delay. For simplicity assume that  $k$  and  $S$  can be non-integer.



Figure 1: There are  $\ell$  links between host  $A$  and host  $B$ .

The delay to send the first segment is  $\ell \frac{S+H}{R}$  second. The delay for the remaining  $F/S - 1$  packets is  $\frac{H+S}{R}$  (In other words, every  $\frac{H+S}{R}$  second a new segment arrives at the host  $B$ ). Thus the total delay is:

$$\frac{S+H}{R} \times (\ell + F/S - 1)$$

Now we take the derivative with respect to  $S$  and setting it to 0:

$$(\ell - 1) - \frac{FH}{S^2} = 0$$

then the value  $S$  which minimize the delay is

$$S = \sqrt{\frac{FH}{\ell - 1}}.$$



Note that the second derivative is  $\frac{2FH}{S^3}$  which is always positive (convexity condition). Therefore  $S = \sqrt{\frac{FH}{\ell-1}}$  will minimize the delay.



No. .	Time	Source	Destination	Protocol	Info
113	3.513465	128.178.151.105	74.125.43.106	HTTP	GET / HTTP/1.1
114	3.569630	74.125.43.106	128.178.151.105	HTTP	HTTP/1.1 302 Found (text/html)
115	3.571624	128.178.151.105	74.125.43.103	HTTP	GET / HTTP/1.1
121	3.682999	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 200 OK (text/html)
123	3.759971	128.178.151.105	74.125.43.103	HTTP	GET /extern_js/f/CgJlbhICY2grMA04OywrMA44BywrMBY4ECwrMbc4AywrMBg4BCwrMBk4CSwrMBI
141	3.897220	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 200 OK (text/javascript)
142	4.005004	128.178.151.105	74.125.43.113	HTTP	GET /generate_204 HTTP/1.1
143	4.043497	74.125.43.113	128.178.151.105	HTTP	HTTP/1.1 204 No Content
146	4.111207	128.178.151.105	74.125.43.103	HTTP	GET /cs1?v=3&s=webhp&action=&e=17259,18167,21766,21920,22107,22194,22217&ei=mc_
147	4.166620	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 204 No Content
363	16.909073	128.178.151.105	74.125.43.103	HTTP	GET /setprefs?sig=0_5gnLALef_LbMGLUAc7mBmFR_s9g=&hl=de HTTP/1.1
365	16.952305	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 302 Found (text/html)
366	16.957871	128.178.151.105	74.125.43.103	HTTP	GET / HTTP/1.1
371	17.027366	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 200 OK (text/html)
373	17.104635	128.178.151.105	74.125.43.113	HTTP	GET /generate_204 HTTP/1.1
375	17.143281	74.125.43.113	128.178.151.105	HTTP	HTTP/1.1 204 No Content
376	17.161764	128.178.151.105	74.125.43.103	HTTP	GET /cs1?v=3&s=webhp&action=&e=17259,18167,21766,21920,22107,22194,22217&ei=p8_
377	17.221356	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 204 No Content

Frame 114 (483 bytes on wire, 483 bytes captured)

- ⊞ Ethernet II, Src: Cisco\_ff:fc:50 (00:08:e3:ff:fc:50), Dst: Us1\_6d:19:e3 (00:1a:6b:6d:19:e3)
- ⊞ Internet Protocol, Src: 74.125.43.106 (74.125.43.106), Dst: 128.178.151.105 (128.178.151.105)
- ⊞ Transmission Control Protocol, Src Port: http (80), Dst Port: netopia-vo3 (1841), Seq: 1, Ack: 617, Len: 429
- ⊞ Hypertext Transfer Protocol
  - ⊞ HTTP/1.1 302 Found\r\n
Location: http://www.google.ch/\r\n
Cache-Control: private\r\n
Content-Type: text/html; charset=UTF-8\r\n
Date: wed, 14 oct 2009 13:18:17 GMT\r\n
Server: gws\r\n
Content-Length: 218
X-XSS-Protection: 0\r\n
\r\n
- ⊞ Line-based text data: text/html

Figure 3: Wireshark traces - The first HTTP response from the web server

**Question 3:** What is the status code of the first HTTP response message (detailed in Figure 3) from the server? What does this status code indicate? Hint: you can infer the answer from the traces.

*The status code of the first HTTP response message from the server is 302 Found. This status code indicates that the response to the request can be found under another URL using a GET method. (Note: This is the most popular redirect code, but also an example of industrial practice contradicting the standard. In HTTP1.0 this status code required the client to perform a temporary redirect, but popular browsers implemented it as a 303 See Other. HTTP1.1 added status code 303 to distinguish between the two behaviors. However, the majority of Web applications and frameworks still use the 302 status code as if it were the 303.)*

**Question 4:** Why did the server respond with this status code? What happens as a consequence when the student's browser receives this HTTP response message?

*Based on the IP address of the received request, the server determines that the request has been sent from a machine with a geographical location in Switzerland. Thus, the server wants to serve this client requested webpage that is customized for visitors from Switzerland and that will give them the most relevant results. In this example, Google server redirects visitors from Switzerland to view the page [www.google.ch](http://www.google.ch) instead of the page [www.google.com](http://www.google.com). Another possibility is that the server redirects the student's browser based on the student's preferences which are stored in the cookie that was sent to the server associated with the HTTP GET request. Following the receipt of the HTTP response message with the status code 302 the student's browser issues a HTTP GET request to the host [www.google.ch](http://www.google.ch).*

**Question 5:** What is the IP address of the machine from which the request was sent? What are the IP addresses of the servers from which the content of the webpage was downloaded?

*The IP address of the user's machine is 128.178.151.105. The content of the webpage is downloaded from two servers with IP addresses 74.125.43.103 and 74.125.43.113.*



No. .	Time	Source	Destination	Protocol	Info
113	3.513465	128.178.151.105	74.125.43.106	HTTP	GET / HTTP/1.1
114	3.569630	74.125.43.106	128.178.151.105	HTTP	HTTP/1.1 302 Found (text/html)
115	3.571624	128.178.151.105	74.125.43.103	HTTP	GET / HTTP/1.1
121	3.682999	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 200 OK (text/html)
123	3.759971	128.178.151.105	74.125.43.103	HTTP	GET /extern_js/f/CgJlbhICY2grMAo4oywrMA44BywrMBY4ECwrMBC4AywrMBg4BCwrMBk4CSwrMBk4
141	3.897220	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 200 OK (text/javascript)
142	4.005004	128.178.151.105	74.125.43.113	HTTP	GET /generate_204 HTTP/1.1
143	4.043497	74.125.43.113	128.178.151.105	HTTP	HTTP/1.1 204 No Content
146	4.111207	128.178.151.105	74.125.43.103	HTTP	GET /csi?v=3&s=webhp&action=&e=17259,18167,21766,21920,22107,22194,22217&ei=mc_
147	4.166620	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 204 No Content
363	16.909073	128.178.151.105	74.125.43.103	HTTP	GET /setprefs?sig=0_5GnLA1ef_LBMGLUAc7mBMfr_S9g=&hl=de HTTP/1.1
365	16.952305	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 302 Found (text/html)
366	16.957871	128.178.151.105	74.125.43.103	HTTP	GET / HTTP/1.1
371	17.027366	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 200 OK (text/html)
373	17.104635	128.178.151.105	74.125.43.113	HTTP	GET /generate_204 HTTP/1.1
375	17.143281	74.125.43.113	128.178.151.105	HTTP	HTTP/1.1 204 No Content
376	17.161764	128.178.151.105	74.125.43.103	HTTP	GET /csi?v=3&s=webhp&action=&e=17259,18167,21766,21920,22107,22194,22217&ei=p8_
377	17.221356	74.125.43.103	128.178.151.105	HTTP	HTTP/1.1 204 No Content

```

Frame 365 (663 bytes on wire, 663 bytes captured)
Ethernet II, Src: Cisco_ff:fc:50 (00:08:e3:ff:fc:50), Dst: Usi_6d:19:e3 (00:1a:6b:6d:19:e3)
Internet Protocol, Src: 74.125.43.103 (74.125.43.103), Dst: 128.178.151.105 (128.178.151.105)
Transmission Control Protocol, Src Port: http (80), Dst Port: transact (1869), Seq: 17334, Ack: 2867, Len: 609
Hypertext Transfer Protocol
⊞ HTTP/1.1 302 Found\r\n
  Location: http://www.google.ch\r\n
  Cache-Control: private\r\n
  Content-Type: text/html; charset=UTF-8\r\n
  Set-Cookie: PREF=ID=9F9a8a9a78b6dceb:U=8d286f9a92a0d7dd:LD=de:NR=10:TM=1255525867:LM=1255526311:S=AVJi3AjkdJap6Vhm; expires=Fri, 14-Oct-2011
  Date: wed, 14 Oct 2009 13:18:31 GMT\r\n
  Server: gws\r\n
  Content-Length: 217
  X-XSS-Protection: 0\r\n
  \r\n
Line-based text data: text/html

```

Figure 5: Wireshark traces - HTTP response from the web server

**Question 7:** Explain what was the student's action that caused this GET request (packet 363 in Figure 4) and the further communication between his machine and the server (packets 363 to 377).

*The student has changed his language preference, by clicking on the link on the webpage. The click has caused the browser to send the GET request with the string **setprefs** in the URL, which indicates the change in the student's preferences. In the previous question we have seen that this GET request has caused the server to set a new cookie on the student's machine, thus also indicating that the student's preferences have been changed. Given the observed difference in the values of the two cookies and the string **hl=de** in the URL of the GET request, the student has changed his language preference (from English to German).*

**Question 8:** In this further communication (packets 363 to 377), the student again obtains a webpage. Has the student already seen this webpage or not (in the current browsing session given by the traces)?

*As the user has changed his language preference, a new webpage (in German) is served to the user for the same URL [www.google.ch](http://www.google.ch). Therefore, the user has not seen this page before.*

**Question 9:** The student deletes all the cookies in his browser and types the same *URL* as in the beginning of the browsing session. What webpage would the student download?

*The student types [www.google.com](http://www.google.com). However, his browser will be redirected to [www.google.ch](http://www.google.ch) based on the IP address of the student's machine. (In the case when there is no redirection based on the IP address, the student would download the [www.google.com](http://www.google.com) webpage.) The student's language preferences will not be respected as the server cannot "recognize" the student (thus it doesn't know the student's preferences) because the student has deleted all the cookies in his browser.*

**Question 10:** When the first student completes the browsing session (given by the traces) another student at EPFL starts a browsing session from a different machine. The second student enters the same *URL* in the browser as the first student. Assume that the EPFL network has a local DNS server and a shared proxy Web cache. Also, assume that the local cache on the second student's machine is empty and all the cookies were deleted before he started the browsing session. Explain from where the webpage is downloaded and why. Explain which application and transport layer protocols will be used to download the page.

*The second student also types [www.google.com](http://www.google.com). However, his browser will also be redirected to [www.google.ch](http://www.google.ch) based on the IP address of the student's machine. Even though the first student has downloaded the same webpage ([www.google.ch](http://www.google.ch)) before the second user, the webpage is not cached at the proxy Web cache because the "Cache-control" field in the header of the HTTP response message (Figure 3) has value **private**. The "Cache-Control" value "private" indicates that the server has generated a personalized response for this user and it can be cached in the users local cache but not in a shared proxy cache. Thus, the webpage is downloaded from Google web servers. Before downloading the page, the IP address of the server was obtained from the local DNS server at EPFL. Application layer protocols used are: **DNS** to resolve URL to the IP address and **HTTP** to download the webpage. Transport layer protocols used are: **UDP** (used by DNS) and **TCP** (used by HTTP).*

## 5 P2P

(6 points)

**Question 1:** Consider an overlay network with  $N$  active peers, with each pair of peers having an active TCP connection. Additionally, suppose that the TCP connections pass through a total of  $M$  routers. How many nodes and edges are there in the corresponding overlay network?

*There are  $N$  nodes in the overlay network. There are  $N(N-1)/2$  edges.*

**Question 2:** Assume a DHT consisting of 5 nodes. In the underlying physical network (Figure 6), if two nodes share a link, the length  $d$  of this link is bounded:  $1 \text{ km} \leq d \leq 5 \text{ km}$ . What is the maximum physical distance between two neighbors in the DHT?



Figure 6: A physical network where all nodes are aligned.

*Neighbors in the DHT are not necessarily neighbors in the physical network, so the maximum distance is 20 km, i.e., between the endpoints of the physical network.*



**Question 3:** Consider distributing a file of size  $F = 700$  Mbytes to  $N = 10,000$  peers. The server has an upload rate of  $u_s = 100$  Mbps, and each peer has a download rate of  $d_i = 5$  Mbps and an upload rate of  $u = 300$  Kbps. Compute the minimum distribution time for both client-server and P2P architectures. Assume that  $1 \text{ Mbps} = 1,000 \text{ Kbps} = 1,000,000 \text{ bps}$  and  $1 \text{ Mbyte} = 1,024 \text{ Kbytes} = (1,024)^2 \text{ bytes}$ .

*For calculating the minimum distribution time for the client-server distribution, we use the following formula:  $D_{cs} = \max\{NF/u_s, F/d_{min}\} = 163 \text{ h } 6 \text{ min } 43 \text{ s}$ .*

*Similarly, for calculating the minimum distribution time for the P2P distribution, we use the following formula:  $D_{P2P} = \max\{F/u_s, F/d_{min}, NF/(u_s + \sum_{i=1}^N u_i)\} = 5 \text{ h } 15 \text{ min } 42 \text{ s}$ .*

**Question 4:** What is the role of Skype relays?

*They allow establishing calls between hosts in home networks that use NATs.*

## 6 TCP

(6 points)

**Question 1:** Complete the missing sequence numbers (Seq), acknowledgment numbers (ACK), and segment length (LEN) in the following TCP connection. We assume:

- No timeouts occur at the receiver.
- The sender starts the timer at  $t_1$ .
- The connection is full duplex (bi-directional data flow in same connection).
- The sender and the receiver have always data to transmit.
- There are no delayed acknowledgements at the sender or the receiver.

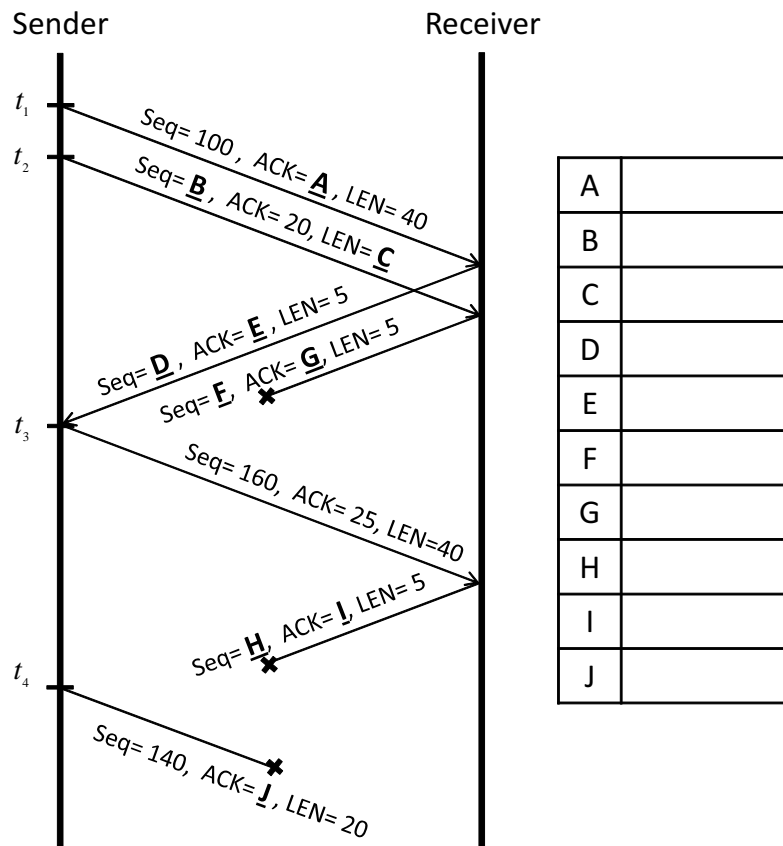


Figure 7: Sequence and ACK numbers for a TCP connection.

**Question 2:** Calculate the value of timeout for the segment with the sequence number 140.

As the sender restarts the timer at  $t_3$  and segment 140 is retransmitted at  $t_4$ , the timeout is  $t_4 - t_3$ .




Figure 8: Sequence and ACK numbers for a TCP connection.