

TDTS11 – Computer networks and Internet protocols (TEN1)
Final Examination: 8:00-12:00, Wednesday, June 5, 2013
Time: 240 minutes
Total Marks: 40
Grade Requirements: Three (20/40); four (28/40); and five (36/40).
Assistance: None (closed book, closed notes, and no electronics)
Instructor: Niklas Carlsson

Instructions:

- Read all instructions carefully (including these)!!!! Some questions have multiple tasks/parts. Please make sure to address *all* of these.
- The total possible marks granted for each question are given in parentheses. The entire test will be graded out of 40. This gives you 10 marks per hour, or six minutes per mark, plan your time accordingly.
- This examination consists of a total of 9 questions. Check to ensure that this exam is complete.
- When applicable, please explain how you derived your answers. Your final answers should be clearly stated.
- Write answers legibly; no marks will be given for answers that cannot be read easily.
- Where a discourse or discussion is called for, be concise and precise.
- If necessary, state any assumptions you made in answering a question. However, remember to read the instructions for each question carefully and answer the questions as precisely as possible. Solving the *wrong* question may result in deductions! It is better to solve the *right* question incorrectly, than the *wrong* question correctly.
- Please write your AID number, exam code, page numbers (even if the questions indicate numbers as well), etc. at the top/header of each page. (This ensures that marks always can be accredited to the correct individual, while ensuring that the exam is anonymous.)
- Answers can be provided in either English or Swedish. (If needed, feel free to bring a dictionary from an official publisher. Hardcopy, not electronic!! Also, your dictionary is not allowed to contain any notes; only the printed text by the publisher.)
- Good luck with the exam.

1) Question: Forwarding (6)

Show, illustrate, and explain the path of the first GET request *and* the first GET response between a Web client (browser) and a Web server. You can make the following assumptions:

- The client machine uses Ethernet, has a single interface with a MAC address AA:BB:AA:BB:AA:BB and an IP address 111.222.111.222
- The MAC and IP addresses of the Web server are DD:AA:DD:AA:AA:DD and 222.222.111.111. Similar to the client, the server has a single interface.
- The gateway router closest to the client has four interfaces. The first is the interface closest to the client and has MAC and IP addresses BB:BB:BB:BB:BB:BB:AA and 111.222.111.1. The second interface has MAC and IP addresses BB:BB:BB:BB:BB:BB:BB and 111.222.111.2. The third interface has MAC and IP addresses BB:BB:BB:BB:BB:BB:CC and 111.222.111.3. Finally, the fourth interface has MAC and IP addresses BB:BB:BB:BB:BB:BB:DD and 111.222.111.4.
- The gateway router closest to the server has four interfaces. The first is the interface closest to the server and has MAC and IP addresses CC:CC:CC:CC:CC:CC:AA and 222.222.111.1. The second interface has MAC and IP addresses CC:CC:CC:CC:CC:CC:BB and 222.222.111.2. The third interface has MAC and IP addresses CC:CC:CC:CC:CC:CC:CC and 222.222.111.3. Finally, the fourth interface has MAC and IP addresses CC:CC:CC:CC:CC:CC:DD and 222.222.111.4.
- The routing table at the gateway router closest to the client has many entries. However, for this question, the three most closely matching entries for each interface states 222.222.0.0/24 (over interface 2), 222.222.64.0/18 (over interface 3), and 222.222.0.0/18 (over interface 4).
- The routing table at the gateway router closest to the server has many entries. However, for this question, the three most closely matching entries for each interface states 111.222.192.0/16 (over interface 2), 111.222.0.0/17 (over interface 3), and 111.222.128.0/18 (over interface 4).

In addition to the above answers, you should also draw a picture of the topology and clearly state any assumptions you make about the topology (including parts of the networks not explained above) or anything else needed to solve the question. As with all your answers it is important that you also explain how you derived your answer. For example, why was the packet taking this particular route?

2) Question: Encapsulation (4)

Consider the same scenario as in question 1 (above). Show and illustrate the link-layer frame (and encapsulated information) for the first GET response message when it reaches the network interface card (link layer) of the Web server. You do not have to show all the details of the different headers; however, you should (i) specify what protocols the different headers are associated with, and (ii) provide the address information associated with the source and destination fields for each of the different headers (contained within this frame).

3) Question: HTTP and replication (6)

Performance and personalized service are important aspects of building good Web services. Please draw a picture illustrating the communication sequence when a client downloads a Web page with four embedded images. The main document and two of the embedded images can be found on the original Web server `www.aa.com` and the final two embedded images can be found on the server `www.bb.com`. Your picture should illustrate the client, its local DNS server, all involved Web servers, and the communication sequence. Your picture and accompanying explanation should assume that both the servers and the client communicate with pipelined HTTP (rather than non-persistent HTTP/1.0). Your picture should also clearly show connection establishment and teardown messages, as well as any other messages needed for the file transfer. You can assume that the HTML page and each of the images are 4KB, and the MSS is 1.5KB. Also, how many packets are being sent between each of the two servers and the client?

4) Question: TCP slow start (4)

Consider two machines A and B which are located 200ms apart. Assume that A is requesting a file from B using HTTP. Draw a figure and explain the entire communication sequence, including TCP handshake and connection teardown. You can assume that the payload is 16 packets and each packet can be sent in 1ms. You can assume that payload packet two (2) and eleven (11) are lost. For simplicity, you can assume that the TCP version is implementing fast retransmit, initially have $ssthresh = 8$, and the timeout period is constant at 500ms.

5) Question: TCP throughput (6)

Consider the throughput of two users (A and B) with a shared bottleneck link. Assume that they both downloading (very) large files from different servers, but that their bandwidth bottleneck is the shared link. Assume A uses one Transmission Control Protocol (TCP) connection, and B uses two TCP connections. The round trip time (RTT) for the connection of clients A is 125ms and the RTT for client B is 200ms. Furthermore assume that the total bandwidth of the link is 100Mbps. Show how the window size and throughput may vary over time for the two users. Also, please estimate the download rate of the different clients?

6) Question: Streaming (2)

In practice, TCP is typically used for video streaming? Please clearly and briefly explain how HTTP and caching plays an important role in reducing the load of the origin servers in such streaming systems.

7) Question: Distance vector routing (4)

Consider a node A with neighbors B, C, and D. Node A currently has the distance table below. (Note that the network currently is not in a very good shape, and a few routing table updates will be needed to get the network back into shape.) Assume that it receives an updated distance vector from neighbor B which looks as follows $[\infty, 0, 3, 7, 2, 11, 2]$ and an updated distance vector from neighbor C which looks as follows $[\infty, 5, 0, 5, 3, 8, 5]$. First, update the table below, including A's own distance vector. Second, assume that poison reverse is implemented, and show and explain exactly what information the node sends to each its neighbors (after the table has been updated).

Destination	Costs			
	A (via)	B	C	D
A	0 (A)	∞	∞	∞
B	1 (B)	0	4	2
C	1 (C)	2	0	3
D	1 (D)	3	5	0
E	? (?)	7	5	4
F	? (?)	2	12	9
G	? (?)	12	9	6

8) IP fragmentation (4)

Consider a 4,800 byte IPv4 datagram which reaches a link with a MTU of 1,500. Please show and explain how the datagram is fragmented. For each datagram of consideration (both before and after fragmentation), please indicate the value of each of the following header fields: (i) offset, (ii) ID, (iii) morefrag, and (iv) length. Finally, explain where and why the original datagram is being reassembled?

9) End-to-end delay (4)

A common mistake when designing distributed systems or applications is to ignore the network delay. Consider ten (10) messages, each of size $L = 1.5\text{KB}$ that will be sent and routed back-to-back between two machines in a distributed system using a regular store-and-forward network, such as the Internet. For simplicity, you can assume that the message do not experience any queuing delays (e.g., the network is empty or there is minimal competing traffic). Along their path the messages will traverse three routers, each requiring a packet processing delay $t_p = 1\text{ms}$ per packet. Each link i ($i = 1,2,3,4$) on the path can be assumed to have a propagation speed $s_i = 2.5 \cdot 10^8$ m/s, length $l_i = 2,500$ km, and transmission rate $R_i = 10$ Mbps. Please (i) draw the topology, (ii) list the different delays associated with the end-to-end path, (iii) derive an expression for the end-to-end delay, (iv) calculate an estimate of the end-to-end delay, and (v) the time between the first byte is sent and the last byte is fully received.

Good luck!!