

TDTS06 – Computer networks (TEN1)

Final Examination: 14:00-18:00, Thursday, August 29, 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 10 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:AA and 111.222.111.1. The second interface has MAC and IP addresses BB:BB:BB:BB:BB:BB and 111.222.111.2. The third interface has MAC and IP addresses 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: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:AA and 222.222.111.1. The second interface has MAC and IP addresses 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 and 222.222.111.3. Finally, the fourth interface has MAC and IP addresses 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 four most closely related entries for each interface states 222.222.0.0/24 (over interface 2) 222.222.64.0/18 (over interface 3), 222.222.0.0/18 (over interface 4), and 4) 222.222.0.0/16 (over interface 2).
- The routing table at the gateway router closest to the server has many entries. However, for this question, the four most closely related entries for each interface states 111.222.192.0/16 (over interface 2), 111.222.0.0/17 (over interface 3), 111.222.128.0/18 (over interface 4), and 111.222.0.0/24 (over interface 2).

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 and not some other 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 request 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 five 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 6KB, and the MSS is 1.5KB. Also, how many packets are being sent between each of the two servers and the client? Which of the two servers is likely to be active for the longest period of time?

4) Caching and end-to-end delay (4)

HTTP caches play an important role in networks. Assume that an ISP installs a cache in its network. Estimate how much improvement in the average RTT the clients will see after such installation. You can assume that (i) the clients in the ISP's network have an average RTT of 150ms to a representative sample set of Webpages that these clients access, (ii) after the initial warmup period (when the cache initially is empty) the cache will see an average hit rate of 35%, and (iii) the clients on average are located 5ms (single-direction delay) from the new cache. Also, can you reason about how much bandwidth the cache will save the ISP? For this part, you can first assume that the hit rate of all file sizes are the same, and then (perhaps more realistically) assume that the hit rate is smaller for large files than small files.

5) Question: TCP fairness (4)

Assume a bottleneck link with four users behind it A, B, C, and D? Assume that they are all downloading large files from different servers, but that their bandwidth bottleneck is the shared link. Assume A uses two connections, B five connections, C two connections and D three connections. The round trip time (RTT) for the connections of clients A and B are 100ms and the RTT for clients C and D are 150ms. Furthermore assume that the total bandwidth of the link is 100Mbps but under the above conditions only can operate at 90% utilization. Estimate the download rate of the different clients?

6) Question: TCP slow start (4)

Consider two machines A and B which are located 100ms 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 20 packets and each packet can be sent in 1ms, and that the fifth (5) payload packet is 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. You should include a figure that clearly show when each packet is sent and received.

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 $[\infty, 0, 2, 7, 5, 8, 3]$ and an updated distance vector from neighbor C which looks as follows $[\infty, 3, 0, 6, 2, 3, 5]$. 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	7
D	1 (D)	3	5	0
E	? (?)	7	5	4
F	? (?)	2	12	9
G	? (?)	12	9	7

8) IP fragmentation (4)

Consider a 2,000 byte IPv4 datagram which traverses 10 links on its path from A to B. Assume that links 1, 2, 6 and 8 have an MTU of 3,000 and the other links have an MTU of 1,500. Please show and explain when and into how many fragments the datagram is fragmented/defragmented, as it traverses each link of the path from A to B. For each datagram of consideration (both before and after fragmentation), your figure should also show the value of each of the following IP header fields: (i) offset, (ii) ID, (iii) morefrag, and (iv) length.

9) HTTP-based Adaptive Streaming (4)

Explain what the main difference between typical (non-adaptive) HTTP-based streaming (as used by YouTube, for example) and HTTP-based adaptive streaming. Your answer should illustrate the file structures and describe the difference in player actions under a simple scenario with varying bandwidth availability.

Good luck!!