TDTS11/TEN1 – Computer networks and Internet protocols
TDDD93/TEN1 – Large-scale distributed systems and networks

Final Examination: 14:00-18:00, Saturday, March 21, 2015
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+1=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.)
- Please answer in English to largest possible extent, and try to use Swedish or ”Swenglish” only as needed to support your answers. In general, questions are designed to be answered using figures/tables.
- 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)

Please carefully show, illustrate, and explain the path of (i) the *first HTTP request* and (ii) the *first HTTP response* between a Web client (browser) and a Web server. You can make the following assumptions:

- The client is located in Sweden and the server in USA.
- The client machine uses Ethernet, has a single interface with a MAC address AA:AA:AA:AA:AA:AA.
- The GET request is for a webpage: www.aa.com/index.html.
- The client has obtained a dynamic IP address 101.101.101.101 from a DHCP server, which is running on the closest gateway router.
- The client uses a local DNS server with IP address 101.101.1.1 and MAC address EE.EE.EE.EE.EE.EE.
- The MAC and IP addresses of the Web server are DD:DD:DD:DD:DD:DD and 197.222.111.111. 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 B1:B1:B1:B1:B1:B1 and 101.101.101.1. The second interface has MAC and IP addresses B2:B2:B2:B2:B2:B2 and 101.101.101.2. The third interface has MAC and IP addresses B3:B3:B3:B3:B3:B3 and 101.101.101.3. Finally, the fourth interface has MAC and IP addresses B4:B4:B4:B4:B4:B4 and 101.101.101.4.
- 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 C1:C1:C1:C1:C1:C1 and 196.222.111.1. The second interface has MAC and IP addresses C2:C2:C2:C2:C2:C2 and 196.222.111.2. The third interface has MAC and IP addresses C3:C3:C3:C3:C3:C3 and 196.222.111.3. Finally, the fourth interface has MAC and IP addresses C4:C4:C4:C4:C4:C4 and 196.222.111.4.
- The forwarding table at the gateway router closest to the client has many entries. However, for this question, the four most closely related entries states: 196.0.0.0/6 over interface 2, 197.0.0.0/8 over interface 3, 197.196.0.0/10 over interface 4, and 197.222.0.0/19 over interface 2.
- The forwarding table at the gateway router closest to the server has many entries. However, for this question, the four most closely related entries states: 101.101.192.0/18 over interface 2, 101.101.0.0/19 over interface 3, 101.101.64.0/18 over interface 4, and 101.101.0.0/16 over interface 2.

For this question you should also draw a picture of the topology, clearly indicate the path taken by the packet in this 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 (6)

Consider the same scenario as above (Q1). Please use figures to show and illustrate the link-layer frame and encapsulated information for:

1. the first *DHCP discovery* message when it is at the *outgoing* network interface card (link layer) of the *client*;
2. the first *DNS query/request* message when it is at the *outgoing* network interface card (link layer) of the *client*; and
3. the first *HTTP request* message when it reaches the *outgoing* network interface card (link layer) of the *gateway router closest to the server*.

You do not have to show all the details of the different headers; however, you should provide a figure that clearly (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 these two frames.

3) Question: TCP slow start (4)

Consider the download of the above file *index.html*, located on the server hosting *www.aa.com*. Assume that the client and server are using non-persistent HTTP/1.0, and that there is a 125ms round trip time (RTT) between them. Please draw a figure and explain the entire communication sequence associated with the file download, including the TCP handshake and connection teardown. You can assume that the payload is 20 packets, that each packet can be sent in 1ms, and that the eleventh (11) payload packet is lost. You can also assume that the TCP version is implementing fast retransmit and fast recovery, that the initially *ssthresh* value is 4, and that the timeout period is (approximately) constant at 500ms. You should include a figure that clearly show when each packet is sent and received.

4) 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 the client above downloads a smaller version of the above file *www.aa.com/index.html*, which in this question is only 2.5KB, but the webpage also include eight embedded images. The main document and four of the embedded images can be found on the original Web server *www.aa.com* and the final four embedded images can be found on the server *www.bb.com*. You can assume that no proxy cache is used and the client communicates directly with the server *www.aa.com using non-persistent HTTP* and directly with the server *www.bb.com using pipelined HTTP*. Your picture should illustrate the client, all involved Web servers, and the communication sequence. 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 2.5KB each, the MSS is 1.5KB, and both webservers have the same RTT and loss rates to the client. Also, how many packets are being sent between each of the two servers and the client? Which connections are likely to terminate last?

5) 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 [∞, 0, 1, 2, 5, 5, 5] and an updated distance vector from neighbor C which looks as follows [∞, 2, 0, 3, 7, 2, 7] . 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).

	Costs			
Destination	A (via)	B	C	D
A	0 (A)	∞	∞	∞
B	1 (B)	0	4	5
C	1 (C)	2	0	9
D	1 (D)	3	5	0
E	? (?)	7	8	4
F	? (?)	8	12	3
G	? (?)	4	5	7

6) Question: BGP routing (4)

Consider the following information visible to a Border Gateway Protocol (BGP) router in an Autonomous System (AS) A.

- There are four neighboring ASes: B, C, D, and E.
- AS A's internal path costs to the closest gateway routers of each of these three ASes are 2, 3, 4, and 5, respectively.
- A is a customer of B; A is a customer of C; C is peering with D; D is a customer of E; A is peering with D; A is peering with and E; and D is peering with Q.
- For prefix 11.22.0.0/20 AS A sees the following route announcements: B Y Z; C F Z; D Q P Z; and E R Z.
- The last web request that a client located within AS A did to a server located in AS B, C, D, E had a RTT of 10ms, 15ms, 30ms, and 25ms, respectively.
- The last traceroute that somebody in AS B performed to a server or router in AS Z required 12 router hops. The corresponding numbers for traceroute queries between ASes C-Z, D-Z, and E-Z are 13, 15, and 14, respectively.

Please (i) draw a picture of the topology, as seen by AS A, (ii) clearly *list and order* the information that is being used when determining this routing table entry, (iii) clearly *rank the four path alternatives* (B, C, D, E), and (iv) show and carefully explain how the particular entry for the above prefix would look in the router's forwarding table. Note that the order should clearly reflect the priority (and order) used in making the routing decision (and similar routing decisions).

7) Question: TCP fairness (4)

Consider three users (A, B, and C) behind a common bottleneck link. Assume that all three users are downloading large files from different servers, but that their bandwidth bottleneck is the shared link. Assume that user A uses one connection, B uses two connections, and C uses three connections. The round trip time (RTT) for the connections of clients B and C are 50ms and the RTT for client A is 150ms. Furthermore, assume that the total bandwidth of the link is 50 Mbps. Estimate the total download rate of the different clients?

8) IP fragmentation (4)

Consider a 3,400 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) Question: BitTorrent (2)

One of the more important mechanisms in BitTorrent is the rarest-first policy. Please clearly explain what this policy does, what it achieves, and why it has been important for the success of BitTorrent. A figure illustrating the policy will help your answer.

10) Bonus Question: End-to-end delay (4)

A common mistake when designing distributed systems or applications is to ignore the network delay. Consider a message of size L = 1,000 bytes that will be routed 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 its path the message will traverse three routers, each requiring a packet processing delay t_p = 1ms. 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 = 7,500 km, and transmission rate R_i = 100 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, and (iv) calculate an estimate of the end-to-end delay.

Good luck!!