

TDTS04 – Computer networks and distributed systems (TEN1)

Final Examination: 14:00-18:00, Thursday, August 20, 2014

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.)
- 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 (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 111.222.111.222 from a DHCP server, which is running on the closest gateway router.
- The client uses a local DNS server with IP address 111.222.1.1 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 196.222.111.111. 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 196.222.111.1. The second interface has MAC and IP addresses CC:CC:CC:CC:CC:BB and 196.222.111.2. The third interface has MAC and IP addresses CC:CC:CC:CC:CC:CC and 196.222.111.3. Finally, the fourth interface has MAC and IP addresses CC:CC:CC:CC:CC:DD 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 for each interface states 196.222.0.0/18 over interface 2, 196.222.110.0/23 over interface 3, 196.222.110.0/24 over interface 4, and 196.222.64.0/18 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 for each interface states 111.222.192.0/18 over interface 2, 111.222.0.0/18 over interface 3, 111.222.64.0/18 over interface 4, and 111.222.0.0/24 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 (4)**

Consider the same scenario as above (Q1). Please use a figure to show and illustrate the link-layer frame and encapsulated information for the first *HTTP request* message when it reaches the first traversed 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 this frame.

**3) Question: DHCP, ARP and DNS (4)**

Consider the same scenario as above (Q1+Q2). Please use a figure to show and explain (i) the sequence of packets, and (ii) their corresponding header information that were used by the client to obtain all the IP and MAC addresses that it uses when sending the HTTP GET request to the Web server. Please use a figure that clearly shows the timeline and indicates what address information is included in each of these messages. Your figure should include protocol and message information for how the following protocols are used: Dynamic Host Configuration Protocol (DHCP), Address Resolution Protocol (ARP), and Domain Name System (DNS).

**4) Question: TCP slow start (4)**

Consider the download of the above file *index.html*, located on *www.aa.com*. Assume that the client and server are using non-persistent HTTP/1.0, and that there is a 200ms round trip time (RTT) between them. 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, each packet can be sent in 1ms, and that the ninth (9) payload packet is lost. For simplicity, you can assume that the TCP version is implementing fast retransmit, initially have  $ssthresh = 4$ , and the timeout period is (approximately) constant at 500ms. You should include a figure that clearly show when each packet is sent and received.

**5) Question: HTTP and replication (6)**

Performance and personalized service are important aspects of building good Web services. To improve performance when downloading multiple objects from the same server, persistent connections and parallelism can be used. Please draw a picture illustrating the communication sequence when the client above downloads the Web page *www.aa.com/index.html* and three embedded images. You can also assume that both the servers and the client can use up-to five parallel persistent connections, *but* neither the client nor the servers allow pipelining. (In other words, your solution is expected to use persistent connections without pipelining.) The main document and one 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 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 5KB each, the MSS is 1.5KB, and both the contacted servers for the websites *www.aa.com* and *www.bb.com* 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 of the two servers is likely to be active for the longest period of time?

**6) 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  $[\infty, 0, 2, 7, 5, 9, 3]$  and an updated distance vector from neighbor C which looks as follows  $[\infty, 3, 0, 4, 2, 3, 5]$ . 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)	$\infty$	$\infty$	$\infty$
B	1 (B)	0	4	2
C	1 (C)	2	0	7
D	1 (D)	4	4	0
E	? (?)	8	5	4
F	? (?)	4	13	9
G	? (?)	12	9	7

**7) Question: Mutual exclusion and coordinator selection (4)**

Consider a simple scenario in which there are five nodes A, B, C, D, and E. Use a sequence of figures to illustrate and explain the message sequences and coordination between these nodes when (i) node A acts as a central coordinator for a shared memory resources (that all five nodes can use) and both nodes B and C almost at the same time decides that they want to write to the resource, (ii) the bully algorithm is used to elect a new coordinator when node A crash.

**8) Question: Lamport's clock (4)**

Assume that you have three processes p1, p2, and p3 which are implementing Lamport's clocks. There are many events that take place at these processes, including some messages being sent between the processes. In the figure below we use circles and arrows to specify in-processor events and messages being sent between processes, respectively. Please provide the logical timestamps associated with each event. You can assume that all three clocks start at zero, at the left-most point in time. (Also, explain how the processes would adjust their clocks if using Lamport's logical clocks.)



**9) 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, 5, and 7, respectively.
- A is a customer of B and C; A is peering with D and E.
- 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 25ms, 35ms, 45ms, and 55ms, respectively.
- The last traceroute that somebody in AS B performed to a server or router in AS Z required 13 router hops. The corresponding numbers for traceroute queries between ASes C-Z, D-Z, and E-Z are 12, 15, and 14, respectively.

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

**10) Bonus question: IP fragmentation (4)**

Consider a 4,000 byte IPv4 datagram which traverses 10 links on its path from A to B. Assume that links 1, 2, 5 and 7 have a maximum transmission unit (MTU) of 5,000 bytes and the other links have an MTU of 1,500 bytes.

- (a) Please use a figure and explain when fragmentations occurs, and into how many fragments the datagram is fragmented/defragmented as it traverses each link of the forwarding path from A to B.
- (b) Please use a figure and explain how this scenario will change if all routers on the path used IPv6. You can assume that the MTUs are the same as for the IPv4 case.

*Good luck!!*