

TDTS04 – Computer networks and distributed systems (TEN1)

Final Examination: 8:00-12:00, Friday, Aug. 26, 2016

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)

Examiner: 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.)
- 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)

Show, illustrate, and explain the path of (i) the *SYN packet* and (ii) the *third 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 Australia.
- Koalas are cute!
- The client machine uses Ethernet, has a single interface with a MAC address AA:AA:AA:AA:AA:AA.
- The GET request of interest here is for a website (*www.aa.com/index.html*) with two embedded images. You can assume that the images are hosted on the same website but under a directory *images/* and are names *imageAA.gif* and *imageBB.gif*. You can also assume that *imageAA.gif* is requested before *imageBB.gif*.
- The client has obtained a dynamic IP address 123.123.123.123 from a DHCP server, which is running on the closest gateway router.
- The client uses a local DNS server with IP address 123.123.1.1 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 123.123.111.1. The second interface has MAC and IP addresses B2:B2:B2:B2:B2:B2 and 123.123.111.2. The third interface has MAC and IP addresses B3:B3:B3:B3:B3:B3 and 123.123.111.3. Finally, the fourth interface has MAC and IP addresses B4:B4:B4:B4:B4:B4 and 123.123.111.4.
- 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 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, 197.0.0.0/9 over interface 2, 197.128.0.0/9 over interface 3, and 197.222.0.0/16 over interface 4.
- 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: 123.123.192.0/18 over interface 2, 123.123.0.0/18 over interface 2, 123.123.64.0/18 over interface 3, and 123.123.0.0/16 over interface 4.

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 figures to show and illustrate the link-layer frame and encapsulated information for:

1. the first *DNS query/request* message when it reaches the *outgoing* network interface card (link layer) of the *client*; and
2. the second *HTTP request* message when it reaches the *incoming* network interface card (link layer) of 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 (i.e., address information at all layers, including the application layer).

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 100ms 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 22 packets, that each packet can be sent in 1ms, and that the ninth (9th) 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 8, and that the timeout period is (approximately) constant at 300ms. 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 4KB, but the webpage instead include a total of six 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 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 *pipelined HTTP* and directly with the server *www.bb.com* using *persistent HTTP but without pipelining*. 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 4KB 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 $[\infty, 0, 1, 3, 3, 3, 3]$ and an updated distance vector from neighbor C which looks as follows $[\infty, 2, 0, 2, 4, 2, 4]$. 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	5
C	1 (C)	2	0	7
D	1 (D)	7	4	0
E	? (?)	6	8	7
F	? (?)	2	11	3
G	? (?)	10	5	1

6) Question: DNS (4)

Please explain the service provided by DNS? Also, show the typical route of a query which is routed through four kinds of DNS servers: local, authoritative, root, and TLD. What parts of the query is typically recursive and which part is typically iterative? (This should be clear from your picture.)

7) Question: Hot potato routing (2)

BGP is used for much of the routing on the Internet. Please explain how hot potato routing is used in this context.


8) Question: Transparency and multi-tier systems (6)

Transparency plays a central role in some distributed systems. Consider a simple multi-tier system with three levels: a user interface, an application server, and two replicated database servers. Assume these layers are implemented as a distributed cloud service at different geographic locations and that the average round trip time (RTT) between the machines used in the consecutive layers (starting with the top-tier layer) is 30ms and 20ms, respectively. Consider a workload (set of calls) with two different types of "jobs" (call types). The first type results in fully synchronized calls in which the application server requires 50ms total processing and the database requires 100ms processing to satisfy the request. The second type does not require any database access, is fully synchronized and requires 50ms processing at the application server.

- (a) For each of the two types of jobs, how much time is the client process locked from the moment it makes the request to the application server? You can assume that no large data is transferred between the layers such that the call and responses fits within a single package, and that messages do not need to be acknowledged. Please explain your answer and illustrate with a figure.
 - (b) Assuming 50% of the clients make each type of requests, what is the average response time (assuming no competing jobs or other reasons for queuing).
 - (c) Please give concrete examples of two types of transparency that are provided in the above example.
- Remember to explain your answers.

9) 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.)



Good luck!!

