

Information page for written examinations at Linköping University



Examination date	2019-08-30
Room (1)	<u>TER3(14)</u>
Time	8-12
Edu. code	TDDE35
Module	TEN1
Edu. code name Module name	Large-Scale Distributed Systems and Networks (Storskaliga distribuerade system och nätverk) Written examination (Skriftlig tentamen)
Department	IDA
Number of questions in the examination	9
Teacher responsible/contact person during the exam time	Andrei Gurtov
Contact number during the exam time	013-284702
Visit to the examination room approximately	ca. 10:00
Name and contact details to the course administrator (name + phone nr + mail)	Annelie Almqvist annelie.almquist@liu.se 013-282934
Equipment permitted	None (well, please see cover page)
Other important information	
Number of exams in the bag	

TDTS11/TEN1 – Computer networks and Internet protocols
TDDE35/TEN1 – Large-scale distributed systems and networks
TDTS06/TEN1 – Computer networks

Final Examination: 8:00-12:00, Aug. 30, 2019

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)

Examiners: Andrei Gurtov and 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 state assumptions and show/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.
- 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.
- Yet, some question(s) may be ambiguous or have contradicting information. If necessary, please clearly identify any such instance and clearly state any additional necessary assumptions needed in answering such a question.
- 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 and utilize figures and tables to the largest extent.
- 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.
- **Write down AID for later results checking**
- 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.
- Trump is the president in the 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 typical RTT between this particular server and client is 100ms, there are on average 15-20 routers that the packet need to traverse on this path, and typically involve 4-6 ASes. In this example, the current path goes through New York, NY.
- 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 197.222.111.1. The second interface has MAC and IP addresses C2:C2:C2:C2:C2:C2 and 197.222.111.2. The third interface has MAC and IP addresses C3:C3:C3:C3:C3:C3 and 197.222.111.3. Finally, the fourth interface has MAC and IP addresses C4:C4:C4:C4:C4:C4 and 197.222.111.4.
- 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 4, 197.0.0.0/8 over interface 3, 197.128.0.0/9 over interface 2, and 197.222.0.0/18 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: 123.123.0.0/18 over interface 4, 123.123.192.0/18 over interface 3, 123.123.64.0/18 over interface 2, and 123.123.0.0/24 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 reply* message when it reaches the network interface card (link layer) of the *client*; and
2. the *SYN* message when it reaches the (on the gateway) *outgoing* network interface card (link layer) of the *gateway 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 (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 150ms 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 seventh (7th) 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 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. To improve performance when downloading multiple objects from the same server, both pipelined HTTP and proxy caches can be used. Please draw two pictures 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 5KB, but the webpage also include 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*.

In the first figure you can assume that no proxy cache is used and the client communicates directly with the servers using *persistent connections*, but that *none of the requests are pipelined*. In the second figure, you can assume that all communication happens through the proxy, and that the proxy has the html file, but none of the images, stored locally. Furthermore, for the second figure, you can assume that all connections use *pipelined HTTP*. Your pictures should illustrate the client, its local proxy cache, 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. (Please also clearly state any assumptions

regarding the operation of the proxy.) You can assume that the HTML page and each of the images are 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 machines? 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 $[\infty, 0, 4, 4, 4, 4, 4]$ and an updated distance vector from neighbor C which looks as follows $[\infty, 2, 0, 5, 2, 5, 3]$. 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)	3	5	0
E	? (?)	7	8	4
F	? (?)	2	12	4
G	? (?)	11	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 and E; A is peering with C and D.
- 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 45ms, 75ms, 35ms, and 95ms, 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 19, 12, and 15, 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* (over B, C, D, and 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: Wireless (4)

Draw three pictures showing two stations sharing a channel using time, frequency and code multiple access (TDMA, FDMA, CDMA).

8) Question: DNS (4)

Please explain the service provided by DNS? (Be precise!) 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.)

9) Question: SDN (4)

Explain what Software Defined Networking (SDN) is using a picture that shows a controller and two switches using the OpenFlow protocol. Also, please show a typical path taken (a) by the first packet of a flow and (b) later packets. The controller is running a firewall application. Finally, clearly show the interfaces (APIs) and location of control and data planes.

Good luck!!