TDTS04 – Computer networks and distributed systems (TEN1)
Final Examination: 14:00-18:00, Thursday, March 14, 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 (9+1) 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 (4)

Show, illustrate, and explain the path of a SYN-ACK packet in the handshake 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 mail 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:BB:AA and 111.222.111.1. The second interface has MAC and IP addresses BB:BB:BB:BB:BB:BB:BB and 111.222.122.1. The third interface has MAC and IP addresses BB:BB:BB:BB:BB:BB:CC and 111.222.133.1. Finally, the fourth interface has MAC and IP addresses BB:BB:BB:BB:BB:BB:DD and 111.222.144.1.
- The gateway router closest to the server has three interfaces. The first is the interface closest to the client 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 222.222.122.1. Finally, the third interface has MAC and IP addresses CC:CC:CC:CC:CC:CC and 222.222.133.1.
- The routing table at the gateway router closest to the client has many entries. However, for this question, the three most closely matching entries for each interface states 222.222.0.0/16 (over interface 2), 222.222.192.0/18 (over interface 3), and 222.222.0.0/18 (over interface 4).
- The routing table at the gateway router closest to the server has many entries. However, for this question, the three most closely matching entries for each interface states 111.222.0.0/16 (over interface 1), 111.222.192.0/18 (over interface 2), and 222.222.111.0/24 (over interface 3).

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?

2) Question: Encapsulation (4)

Consider the same scenario as in question 1 (above). Show and illustrate the first link-layer frame for a SYN-ACK message when it reaches the 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 three other can be found on the server www.bb.com. Your picture should illustrate the client, it 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 2KB, and the MSS is 1.5KB. Also, how many packets are being sent between each of the two servers and the client?

4) 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 75ms 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?

5) 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 ninth (9) 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 constant at 500ms. You should include a figure that clearly show when each packet is sent and received.

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 [∞, 0, 5, 6, 3, 12, 3] and an updated distance vector from neighbor C which looks as follows [∞, 3, 0, 3, 9, 2, 8] . 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	2
C	1 (C)	3	0	3
D	1 (D)	3	5	0
E	? (?)	7	7	10
F	? (?)	2	12	7
G	? (?)	11	6	4

7) 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 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 40ms and 20ms, respectively. Consider a fully synchronize call in which the application server require 100ms total processing and the database require 120ms processing to satisfy the request.

(a) How long time is the client process looked from the moment it makes the request to the application server? You can assume that no data is transferred between the layers that the call and responses fits within a single package and messages do not need to be acknowledged. Please explain your answer and illustrate with a figure.

(b) Please give concrete examples of two types of transparency that are provided in the above example. Remember to explain your answers.

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.)

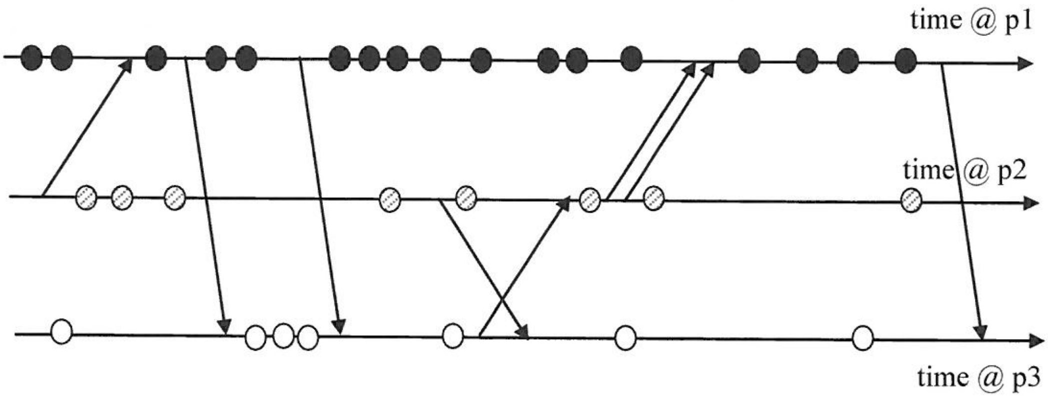

9) Question: End-to-end delay (4)

A common mistake when designing distributed systems is to assume that the network delay is zero. Consider a message of size L = 1,000 bytes that will be routed between two machines in your distributed system using a regular store-and-forward network (such as the Internet). For simplicity, you can (for now) 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 = 2,500 km, and transmission rate R_i = 10Mbps. 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.

10) Bonus Question: BGP routing (4)

Assume that we use Chord to implement a file sharing system. For simplicity, assume that we have a fully populated chord circle, in which there is peer node responsible for each identifier value and with a total of 128 identifier values (numbered from 0 to 127). Further assume that each node keeps track of its successor, its predecessor, and their full finger tables. Now, assume that a file "Blah blah" is made available by peer node 78. At this point, peer node 78 would want to add a pointer to its copy of the file and use a hash function to determine the responsible node. Assume that the hash function *h* resolves to *h*("Blah blah") = 67. Please draw a picture of the system in which you clearly show and explain the route that the insertion message would take to reach peer node 67, from peer node 78. Remember to please show your work and explain your answer.

Good luck!!