Department of Electrical and Information Technology

Final Exam – 2012 ht2

2012-12-19, 14:00 – 19:00

ETSF05/ETSF10 – Internet Protocols

If you have not passed any of the online quizzes, you need to take the whole exam and try to answer all questions. There are **19** questions in this exam, giving a total of **60** points. Minimum **36** points are needed to pass and get mark **3**. You get part of the points for a question if your answer is only partially correct.

For those of you who took part in the online quizzes, the exam questions are grouped into **3** sections. Each section corresponds to a particular online quiz. **For each online quiz you have passed, you are exempt from the corresponding section**; you need not answer the questions in that section. In other words, your exam will be assessed only on those sections you have not already passed online. The assessment will be section-by-section. Each section has a total of **20** points. Minimum **12** points are needed to pass a section.

For those of you who took part in the online quizzes, it will be possible to make use of your online quiz results in this term's retake final exams as well. However, it will not be possible to transfer partial results from this final exam to these retake exams. If, for example, you have passed Online Quiz 1 before and pass only Section 2 today, you will have to take both Section 2 and Section 3 in the upcoming retake final exam.

Please note that, even if you have passed some of the online quizzes, **you can choose to waive your quiz results and take the whole final exam instead**. In that case, your exam will be assessed according to the top paragraph. In case you choose to do so, don't forget to **state this at the beginning of your answers**.

Use all the time given to you. **Answer briefly and clearly.** Choose your words carefully in order not to write answers too long. Keep in mind that none of the questions requires an answer longer than 100 words. **Always motivate your answers.** Unclear, confused, and too generic answers, containing irrelevant information, will decrease your points!

Make sure that your handwriting is clear and readable. Unreadable answers cannot be marked!

If you want, you may use a pocket calculator and a notes page (one side of an A4-size paper, handwritten, which must be handed in with your answer sheets).

Best of success!

Kaan & Jens

Part A

Turn the page to begin…

Questions

1. Internet Routing (corresponds to Online Quiz 1)

1.1. **[6 points]** Perform RIP on the network topology shown in the diagram below, focussing on the routing tables of nodes D and E. (Note that the cost of each link is also given.) Assume that a new link with cost 3 is established between D and E. Show the updated version of E's table. Then, assume that E sends its updated table to D and show how D processes this information step-by-step and row by row to update its own table. Finally, show the updated version of D's table.

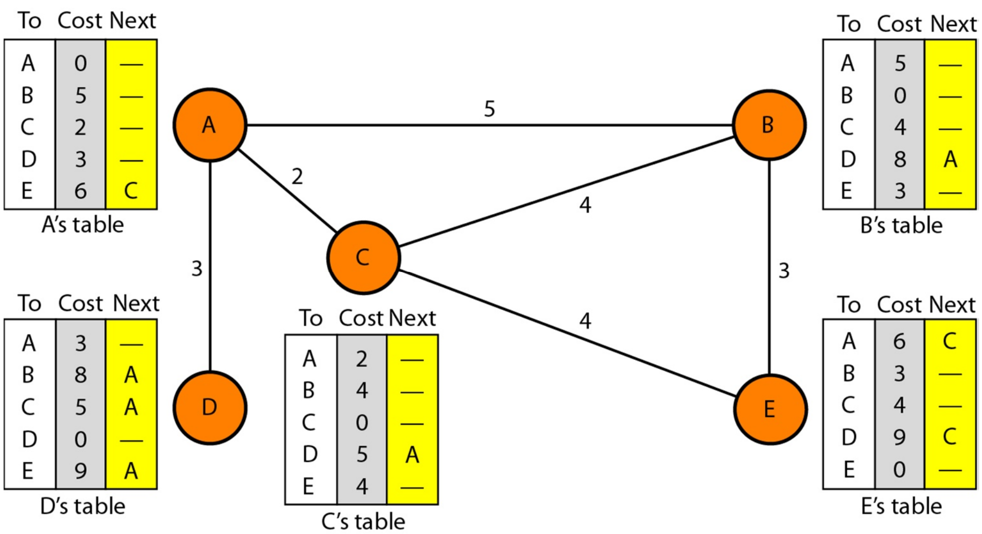

A's table

To	Cost	Next
A	0	—
B	5	—
C	2	—
D	3	—
E	6	C

B's table

To	Cost	Next
A	5	—
B	0	—
C	4	—
D	8	A
E	3	—

D's table

To	Cost	Next
A	3	—
B	8	A
C	5	A
D	0	—
E	9	A

C's table

To	Cost	Next
A	2	—
B	4	—
C	0	—
D	5	A
E	4	—

E's table

To	Cost	Next
A	6	C
B	3	—
C	4	—
D	9	C
E	0	—

Answer:

E's new table:

A	*6*	*C*
B	*3*	*-*
C	*4*	*-*
D	***3***	***-***
E	*0*	*-*

D's new table:

A	*3*	*-*
B	***6***	***E***
C	*5*	*A*
D	*0*	*-*
E	***3***	***-***

1.2. **[6 points]** Use once again the initial network topology shown in the diagram above. (Ignore all routing tables and the extra link added in the previous question.) Assume that the whole topology is known to all the nodes after the flooding of the individual "link state" knowledge. Perform the shortest path tree algorithm used in OSPF, draw the tree step-by-step and derive the routing table for node E.

Answer:

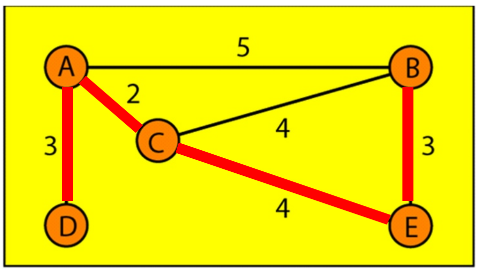

E starts by taking the least cost link to its neighbours. This is link E-B with cost 3. There are 3 candidate links for the next step, E-C with cumulative cost 4, B-A with cumulative cost 8, and B-C with cumulative cost 7. Again, the least cost link is selected, which is E-C. Now, the new candidate link C-A with cumulative cost 6 is the next link to be selected. Finally, A-D with cumulative cost 9 is added to the tree, it is the only option to reach D. E's routing table is identical to the figure at the top of the page.

Questions

1.3. **[3 points]** Given the forwarding table below, what is the next hop address for these destination addresses? (i) 130.235.128.100; (ii) 191.231.194.72; (iii) 100.100.12.192. Motivate your answers.

Net ID	Cost	Next Hop
130.235.0.0/16	5	81.12.32.4
191.231.194.0/26	2	129.100.1.1
84.24.0.0/22	1	181.14.62.5
100.100.12.160/27	3	4.235.17.9
191.231.194.0/24	1	73.32.56.123
0.0.0.0/0	1	112.123.89.1

Answer: *(i) 81.12.32.4 since mask /16 covers the range 0.0 – 255.255. (ii) 73.32.56.123 since mask /26 only covers the range 0 – 63 and the rest belongs to mask /24. (iii) 112.123.89.1 since mask /27 only covers 160 – 191 and there is no other match in the routing table.*

1.4. **[3 points]** It is not only routers that perform routing. Even a local host with only one interface (i.e. network card) has to do it, be it very primitive. Why? What does it do? (Tip: It has to do with ARP.)

Answer: *Local hosts need to find out whether a destination is located on the same LAN or not. If it is, then the packets are delivered directly. If not, they need to be delivered to a router (default gateway) to be forwarded towards the actual destination. Hosts use ARP requests for this purpose, i.e. to find out the MAC address "responsible" for a given IP address.*

1.5. **[2 points]** A source-based tree (SBT) is denoted (Source, Group) while a group-shared tree (GST) is denoted (*, Group). Using your knowledge on SBT and GST, explain why.

Answer: *In the source-based tree approach, each source in a group has its own multicast tree rooted at itself. So, there are as many trees as there are sources. In the group-shared tree approach, all sources in a group share the same tree, common for the whole group, rooted at a special node known as the rendezvous point.*

2. Transport Layer Protocols (corresponds to Online Quiz 2)

2.1. **[4 points]** UDP provides an unreliable, connectionless delivery service over IP. What is meant with (i) unreliable; (ii) connectionless? What can be said about TCP with regard to these issues, and why?

Answer: *(i) UDP is unreliable because there is no mechanism to retransmit lost datagrams. TCP is reliable because lost segments are retransmitted. (ii) UDP is connectionless because there is no session concept, and no connection establishment or termination. TCP is connection-oriented because it has a handshaking process for connection establishment or termination.*

2.2. **[3 points]** What are the two signs of congestion for the sender during a TCP session? How does TCP react to each of these signs? Explain using the terminology relevant to congestion control in TCP.

Answer: *Timeout and 3 duplicate ACK from the receiver. In the former case, the sender resets CWND and goes back to slow start. In the latter case, the sender halves CWND and remains in congestion avoidance.*

2.3. **[2 points]** What is the purpose of RTCP? Explain how it complements RTP, which provides real-time services over the Internet.

Answer: *RTCP controls the real-time multimedia session run by its sister protocol RTP. RTCP packets carry feedback information on the network conditions, transmission and reception statistics, packet loss, jitter, as well as source synchronisation in case there are multiple sources in the RTP stream.*

Questions

2.4. **[4 points]** A sender transfers a file of 5 GB to a receiver. Calculate the transmission time for the file. MSS is 1460 bytes. Standard TCP, IPv6 and Ethernet headers (i.e. without options) are added to each segment before the encapsulated frame can be sent out over a 100 Mbps link. What is the efficiency of this file transfer? Ignore retransmissions, congestion control and flow control in your calculations.

Answer: *5 GB = $5*10^9$ bytes. The number of segments sent is N_S = $5*10^9$ bytes / 1460 bytes = 3424658. Standard TCP, IPv6 and Ethernet headers consist of (20 + 40 + 18) bytes = 78 bytes. So, the size of each frame is S_F = (1460 + 78) bytes = 1533 bytes. The amount of data sent is S_F * N_s = 5250000714 bytes = 5,25 GB. Transmission time is T_T = (5250000714 bytes * 8 bits/byte) / (100000000 bits/s) = 420 s = 7'. Efficiency is the amount of useful data divided by the total: 1460/1533 or 5/5,25, which is 95%.*

2.5. **[4 points]** What is the total delay for a 65 KB (i.e., kilobytes) frame being sent on a network with 20 routers, each router having an average queuing time of 2.50 ms and average processing time of 1.25 ms. The links between the routers all have the same length, 2.5 km, and the speed of light inside the link is $5x10^7$ m/s. The links all have a bandwidth of 100 Mbps. Identify the 4 components of the total delay. Motivate your calculations.

Answer: *(i) Queueing delay: 2,50 ms * 20 = 50 ms. (ii) Processing delay: 1,25 ms * 20 = 25 ms. (iii) Transmission delay: (65000 bytes * 8 bits/byte) / (100000000 bits/s) = 5,2 ms. (iv) Propagation delay: (2500 m * 19) / ($5x10^7$ m/s) = 0,95 ms. Total delay is the sum of all these: 81,15 ms. The delays at the routers dominate the total delay.*

2.6. **[3 points]** What is traffic shaping? Why is it important for achieving good QoS in the network? What are the methods used for traffic shaping?

Answer: *Traffic shaping is about changing the traffic profile of data flow. Generally, bursty traffic is much more difficult to manage in a network since it is unpredictable and, thus, not easy to plan for in terms of network resources. QoS is about throughput, delay, jitter and loss. If we can shape the traffic profile from bursty towards VBR/CBR, it will be easier to maintain the network load at an acceptable level, resulting in better QoS. The methods we use for this purpose are leaky bucket and token bucket.*

3. Network Layer Protocols (corresponds to Online Quiz 3)

3.1. **[2 points]** What is the most important function provided by the network layer? How is this function implemented by IP? (Tip: Think of what is missing at the layers above and below.)

Answer: *End-to-end addressing. Routing protocols depend on this functionality. IP provides an address hierarchy with the concepts of network ID and host ID, and the address blocks are organised in such a way that it becomes possible for the routing protocols to generate forwarding tables without having to know individual host addresses.*

3.2. **[3 points]** Your organisation has been assigned the IPv4 address block 123.**123**.128.0/18. You need to create 500 subnets equal in size. (i) Find the subnet mask. (ii) Find the number of addresses in each subnet. (iii) Find the first and last addresses in the last subnet.

Answer: *(i) You cannot create exactly 500 subnets with binary digits. You need to create 512 = 2^9, which can be identified with an additional 9 bits on top of the original mask /18. So the subnet mask is /27 (ii) The remaining 32 – 27 = 5 bits allow us to create 2^5 = 32 host addresses in each subnet. (iii) We are looking for the address range of block 499. Bits 1-18 are already set to 0111 1011 0111 1011 10 = 123.123.128.0. Without touching the first two decimal numbers corresponding to the first 16 bits, we need to take the last 2 bits and join them with the first 6 bits of block 499 to find the third decimal number. Block 499 is 1 1111 0011 in binary. These are bits 19-27 in the address. Bits 28-32 change between 0 0000 and 1 1111, which is the range we are looking for. Putting everything together, we can now find the third and the fourth decimal numbers in the address. 1011 1110 = 190 is the third number. The fourth number changes between 0110 0000 = 96 and 0111 1111 = 127. So, the range of block 499 is from 123.123.190.96 to 123.123.190.127.*

Questions

3.3. **[2 points]** What are the two important additions/changes that make IPv6 more QoS-friendly than its predecessor IPv4?

Answer: *These are the traffic class and flow label fields in the fixed-sized header. These fields help the routers to (i) identify streams and treat them equally; (ii) classify packets and give them priorities.*

3.4. **[2 points]** Describe the main objectives of ICMP.

Answer: *ICMP is a support protocol for IP. Its objectives are (i) error reporting; (ii) querying the hosts. ICMP messages can be used for network diagnostics and debugging at the IP layer.*

3.5. **[3 points]** Describe the four stages of obtaining an IP address from a DHCP server for the first time. All the messages exchanged during this time are broadcast. Why not only some but all?

Answer: *The stages are as follows: (i) Client sends a DHCP-DISCOVER to get an offer from any server; (ii) Server sends a DHCP-OFFER and offers client an IP address; (iii) Client sends a DHCP-REQUEST to actually get the IP address offered; (iv) Server sends a DHCP-ACK to let client know that operation was successful. The first two messages are broadcast because client doesn't have an IP address yet. The last two are broadcast because client and server need to inform other servers in the same network.*

3.6. **[3 points]** Explain briefly the difference between iterative and recursive DNS resolution. Which of the two methods can benefit the most from caching? Why?

Answer: *In recursive resolution, each DNS server, if it is not the authority for the domain, asks the next server at the higher level until the name is finally resolved. The answer is then sent back via all the involved DNS servers to the host having originated the request. In iterative resolution, the DNS servers reply the host immediately even if they can't resolve the request, and give them the address of the next server at the higher level. It is then the host itself that needs to repeat the process until the name is resolved. Recursive resolution benefits more from caching because all intermediate servers get some useful information during the process. In iterative resolution, only hosts get the information.*

3.7. **[2 points]** Considering how Mobile IP works, explain why it is only suitable for macroscopic mobility (i.e. not continuous and high-speed but infrequent and slow mobility).

Answer: *In Mobile IP, a mobile host needs to find its home agent and foreign agent, and register with the foreign agent to set up packet forwarding. This needs to be done for every foreign site. All this happens with some delay and also creates a lot of overhead. If the mobile host changes its location constantly, and that at high speed, the registration process cannot cope with the speed at which the location changes occur.*

3.8. **[3 points]** What is the hidden terminal problem associated with wireless networks? What is the solution developed in response to this problem and how does it work? Explain with an example.

Answer: *Due to the limitations in transmission power (range) in wireless networks, node A can sense the medium, think it is idle, and start transmission, although it is in fact used by node C outside the range of A to send a message to node B inside the range of A. So, A destroys B's reception because it was not able to detect C's transmission. As a solution, A sends an RTS message to B before actual communication. B replies with CTS. Other nodes in A's and B's neighbourhood overhear either RTS or CTS (or both) and keep quiet during communication. RTS clears off A's neighbourhood and CTS B's. Problem solved.*