

TDTS06 Computer networks, October 22, 2008

Full answers to the written examination, provided by Juha Takkinen, IDA, juhta@ida.liu.se. These answers include motivations, that is, telling how and why one gets the answer, as well as the assumptions made, if any, and also detailed calculations step-by-step—in other words, the way the questions are to be answered according to the instructions given at the exam.

Question 1.

a) Addressing, layer 3, implemented in IP by the IP address, a 32-bit long identifier for networks and hosts.

Demultiplexing, layer 4, implemented in TCP by using the four-tuple consisting of the IP address and the port number of the destination and the source host, respectively, to locate the correct communication process at a host.

Flow control, layer 4, implemented by the Receive-window field in the TCP header.

b)

The missing parts consist of the last leg of the connection-termination phase of TCP, where the client already has closed but the server is still left to do a close.

X = Receive FIN / Send ACK, that is, the client has received the FIN packet from the server, which the client now ACKs.

Y = the server closes down / Send FIN, that is, the server does a close on itself and sends a FIN packet to the client telling this.

Question 2.

a) Routing through R1–R3 is cheaper than through R1–R2–R3 because the total sum of the largest delay (transmission delay) of the nodes involved is much smaller for R1–R3 than R1—R2—R3.

The minimal end-to-end delay (through R1—R3) is therefore 1.575 ms or 1575 microseconds. This value was calculated as follows:

$$\begin{aligned} & d(A\text{-nodal}) + d(R1\text{-nodal}) + d(R3\text{-nodal}) = \\ & = (50 \text{ microseconds} + 0 \text{ microseconds} + 100 \text{ microseconds} + 4 \text{ microseconds}) + \\ & + (20 \text{ microseconds} + 10 \text{ microseconds} + 660 \text{ microseconds} + 40 \text{ microseconds}) + \\ & + (20 \text{ microseconds} + 5 \text{ microseconds} + 660 \text{ microseconds} + 6 \text{ microseconds}) \end{aligned}$$

Answer: the minimal end-to-end delay is 1575 microseconds.

b)

The total delay at node D for 2 large packets: 1.708 ms or 1708 microseconds

This value was calculated as follows: $2 \times d(A\text{-nodal-L}) = 2 \times (50 \text{ microseconds} + 0 \text{ microseconds} + 800 \text{ microseconds} + 4 \text{ microseconds})$

Note. The definition of “nodal delay” only includes the queueing delay and the processing delay, not the transmission delay.

Answer: The total nodal delay at D is 1708 microseconds.

c)

The number of small packets sent in a second is 6493

This value was calculated as follows:

The delay at node A, that is, the time to put a small packet on the outgoing link, is $d(A\text{-nodal}) = d(A\text{-trans-S}) + d(A\text{-proc}) = 100 \text{ microsecs} + 50 \text{ microsecs} = 154 \text{ microsecs}$

The total amount of small packets in one second is then $(1 \text{ second}) / d(A\text{-nodal}) = \text{floor}((1 \times 10^6 \text{ microsecs}) / 154 \text{ microsecs}) = 6493 \text{ pkts}$

(By the way, this translates to a data rate of 5.2 Mbps; for large packets the data rate will be 9.36 Mbps, using the same calculation, which shows how much throughput is affected by the processing overhead for smaller packets; the delay values for A–R1 are actually based on a link speed of 10 Mbps.)

d)

Because the ACK for packet 2 is lost, the Send Window will only advance two steps (for packets 0 and 1) to positions 2–6, according to the definition of Selective-Repeat. Packets 3 and 4 are ACKed and all right.

After a while the retransmission timer will go off and the sender will retransmit packet 2, which means that the receiver will send an ACK for packet 2. When this ACK is received, then the window will finally advance to positions 5–9.

Question 3.

a)

DNS, the Domain Name System, is a distributed, hierarchic database of resource records containing mainly hostname-to-IP-address mappings but also information about mail-drops, for example. It is used by all networked applications.

b)

- i: False
- ii: False
- iii: True
- iv: True

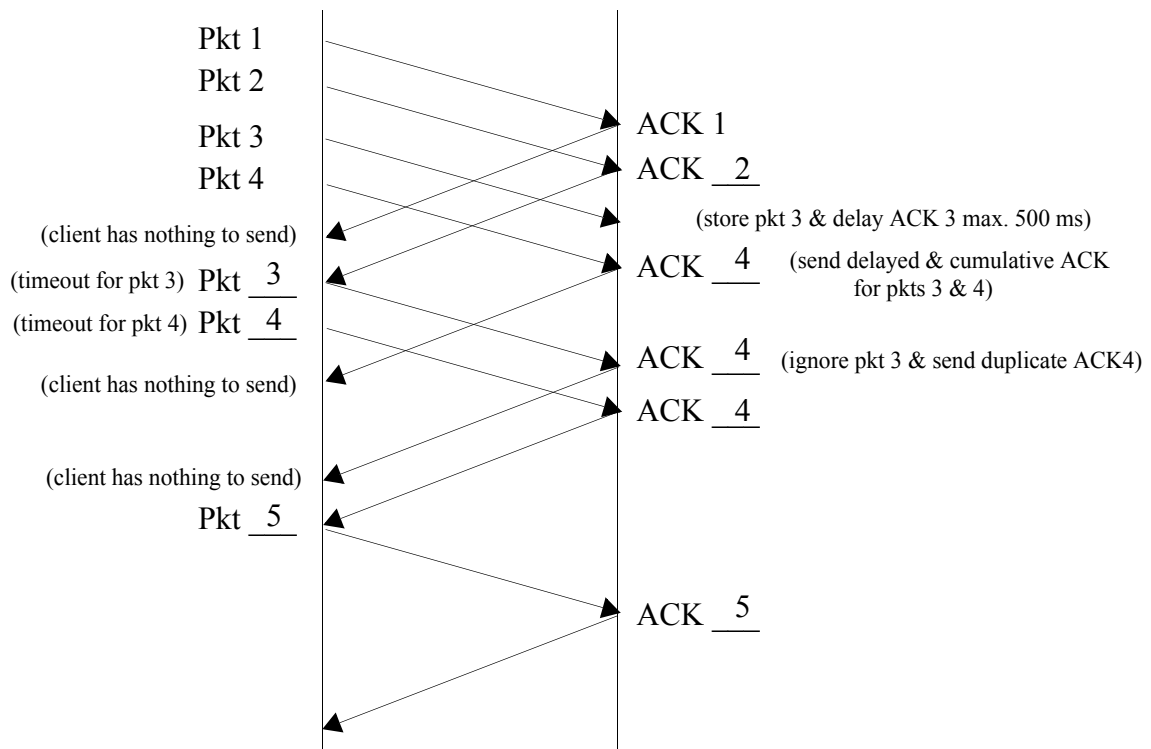
c)

i: The URL consists of a the following parts: the protocol, the host name and the resource (the directory and often also the file, name), that is, <http://gaia.cs.umass.edu/cs453/index.html>. This information can be gathered from the values after GET and Host:

ii: The connection is persistent because the HTTP version is 1.1 and Keep-Alive has a value.

Question 4.

a) i: One possible interpretation is as follows:



ii: Actually, all of them occur in the figure!

- a Delayed and Cumulative ACK is used in the figure for packet 4, which means that the receiving TCP entity deliberately waits a maximum of 500 ms to see if another segment arrives with that time period that can be ACKed too

- a Duplicate ACK is sent for packet 4 in the figure, which means that the receiving TCP entity already has received packet 4, which the sending side insists on retransmitting.

b)

A Triple duplicate ACK means that some packets are still received by the destination but some are also lost or delayed in the network. This is interpreted by the TCP sender as a network soon to be congested, so it backs off to half of the current congestion window (not as severe as a timeout, in other words).

Question 5.

a)

i: Four subnets means that we need four IP addresses in order to be able to address four different networks. However, because the IP addresses that have all-zeros and all-ones are reserved, we need an address space that can accommodate for at least six IP addresses (4 networks and the 2 reserved addresses). Moreover, because the address space must be a multiple of 2, we must use a minimum of 3 bits (2^3), which actually can give us up to 8 possible addresses, including the two reserved ones. Assuming that we will not need to expand our network at the router level anymore, we can select this solution.

For Organization 1, which has the address 200.23.28.0/23, this means that the first three bits after the netmask will be used for the internal addresses, that is, bits number 24–26. The subnet mask will thus be 255.255.255.192.

ii: [The IP address was corrected during the exam to 200.23.18.129]

R4 uses longest-prefix matching to find out that the packet should go to ISPs-R-Us instead of Fly-By-Night-ISP, even though the first 20 bits are identical between 200.23.16.0 and 200.23.18.0. In other words, the netmask for Fly-By-Night-ISP covers more bits.


The gateway router of Organization 1 finds out the correct subnet by applying its subnet mask, which is 255.255.255.192, to find out the subnet address of the packet. The subnet address in this case is 010 in binary or 2 in decimal for the packet.

The gateway router then delivers the packet to either a interface (if the subnet of the packet is on the same LAN as the router) or to another internal router in Organization 1 (if the subnet of the packet belongs to another LAN).

Question 6.

a) i: All of the stations participate in the MAC protocol and must be able to hear collisions, so therefore all the nodes belong to the collision domain.

ii: According to the figure below, node D's transmission (dark grey or red) is disturbed by B's transmission (light grey or yellow) at time 4 (approximately), while B's transmission is scrambled at time 5 by the signal sent by D. Because CSMA/CD is used, D aborts its frame. B also aborts its frame.



iii: [The last sentence was corrected to “Describe how *B* and *D* go about to retransmit their respective frames.”]

B will be the first one to discover a collision and will do binary exponential backoff, which means that the station selects a random number from $\{0, 1\}$ (calculated from the range $0, \dots, 2^2 - 1$) and multiplies it with 512 bit times to get a value for the backoff timer. Station B detects the collision later, after D in time, but does the same calculation because it is the first collision for it too. However, because of the randomization, B will probably not get the same value as D on its backoff timer.

Question 7.

a) Assumption: select the left-most if two columns have the same cost. The links 0–2, 0–4, 1–3, and 2–4 will not be used in the routing process. This is found out by examin-

ing the p-vector in the following table, after running Dijkstra's algorithm on the given network:

S	N'	D(0), p(0)	D(1), p(1)	D(2), p(2)	D(4), p(4)	D(5), p(5)
0	3	4, 3	9, 3		5, 3	3, 3
1	35	4, 3	9, 3	4, 5	5, 3	done
2	350	done	9, 3	4, 5	5, 3	
3	3502		9, 3	done	5, 3	
4	35024		6, 4		done	
	350241		done			

b)

The nodes' distance vector tables are initialized to the path costs found to each node's neighbours. All other fields in the tables are initialized to infinity (marked by a dash). After one exchange the nodes recalculate their tables by using the Bellman-Ford equation.

Step 1, initialization

Node x:

cost to x y z
from:

x 0 5 2

y - - -

z - - -

Node y:

cost to x y z
from:

x - - -

y 5 0 6

z - - -

Node z:

cost to x y z
from:

x - - -

y - - -

z 2 6 0

Step 2, one exchange

Node x:

cost to x y z
from:

x 0 5 2

y 5 0 6

z 2 6 0

Node y:

cost to x y z
from:

x 0 5 2

y 5 0 6

z 2 6 0

Node z:

cost to x y z
from:

x 0 5 2

y 5 0 6

z 2 6 0

The algorithm has finished since there are no new updates to be found in the next exchange.

c)

An AS is typically owned by an organization or a company and can run its own, internal routing protocol in its network. An AS also most often has an ID, which is also glo-

bally unique, assigned by ICANN. The ID is needed if the AS is going to participate in the BGP routing world.

BGP is used to determine reachability between different ASs.

Question 8.

a) i: The nonce is a random number generated by both the server and client and used to make sure that the other party is alive. The nonce is only valid for the current connection, which will prevent palyback attacks.

ii: The certificate is verified by contacting a third-party certificate authority (CA). By using the CAs public key on the certificate, the validity of the certificate can be confirmed.

iii: MAC, the Message Authentication Code, is a hash that contains a shared secret (known by both the server and the client); it is used to authenticate and confirm the integrity of the messages.

iv: Security is needed in many layers, depending on the service that needs to secured. User-specific information, for example, is accessible at the application layer in e-mail, so therefore if this information must be protected, the e-mail messages have to be encrypted.

At lower layers, security is also important. For example, IP packets may be altered by an intruder if not encrypted.

Question 9.

(24 students did the optional assignment in 2008 and 63% passed it and got 2–4 bonus points in the exam.)