

TDTS06 Computer Networks, October 22, 2007

Sketched answers to the written examination, provided by Juha Takkinen, IDA, juhta@ida.liu.se. ("Sketched" means that you, in addition to the below answers, need to show your calculations, assumptions and justifications.)

Question 1.

a) The missing event/action pairs are

A: SYN+ACK/ACK, which accomplishes the last leg of the three-way handshake to open a TCP connection.

B: FIN/ACK, which closes the client's side of the connections and starts the server's passive closing of the connection.

b)

Assumptions: the user clicks on a hyperlink described by the http protocol and is using a wireless computer at IDA. A typical protocol stack is then, top-down and including the protocols' goals:

http --> goal: to transfer http messages (objects) by "pulling" them from a server to a client

tcp --> goal to address processes and transfer segments reliably end-to-end

ip --> goal: to address hosts and transfer packets hop-by-hop in a network with "best effort"

mac or csma/ca (wlan) --> goal: to transfer frames over one link wirelessly

physical layer (spread spectrum) (wlan) --> goal: to represent bits so that csma/ca can do framing

Not shown: calls to resolve hostname and ip address, using dns and arp, respectively.

c)

A protocol implements a layer and offers services to the layer immediately above it, services such as open/close a connection and multiplexing. The interface is where the services are accessed by the upper layer protocol, for example the uniform protocol interface as defined in x-kernel.

Question 2.

a)

The differences between GBN and SR consist of:

- GBN uses cumulative ACKs whereas SR ACKs each individual packet
- SR buffers packets that are out-of-order on the receiver side, while GBN discards them.

(Another difference: When there is a timeout, GBN will retransmit all outstanding packets with sequence number greater or equal to i , while SR will only retransmit the i th packet.)

With a long-delay, high-bandwidth and almost-no-errors link, it depends on the order of the packets on the arriving side, which protocol would be the best:

- if packets mostly arrive in order, then GBN would be the best
- if packets almost never arrive in order, then SR would be the best.

b)

1. A router queue is full and operates according to FIFO with tail-drop, so the arriving packet is dropped.
2. The packet has become corrupt (checksum failed) at a router and the router therefore drops the packet.

c)

i)

The total delay is the sum of all the four different types of delays: transmission, processing, queueing, and propagation delay. The processing and queueing delays both occur at the outgoing ports of each router, that is, the ports directly connected to the relevant links.

Minimal total end-to-end delay is either A–R1–R2–B or A–R1–R3–B.

For the first path, the delay is:

$$(100 + 4 + 0 + 50) + (660 + 18 + 10 + 20) + (660 + 30 + 20 + 12) + (660 + 20 + 5 + 6) = \\ = 154 + 708 + 722 + 691 = 2275 \mu\text{s}$$

For the second path, the delay is:

$$(100 + 4 + 0 + 50) + (20 + 10 + 40 + 660) + (660 + 20 + 5 + 6) = \\ = 154 + 730 + 691 = 1575 \mu\text{s}$$

Answer: the minimal total delay is over the second path, that is, A-R1-R3-B, and the delay is 1575 μs .

ii)

Note that the routers are always store-and-forward. We assume that the packets are sent back-to-back and transferred along the path A-R1-R3-B. The transfer time is then:

$$\text{- for the first, large packet: } (800 + 4 + 0 + 50) + (5500 + 20 + 10 + 40) + (5500 + 20 + 5 + 6) = 854 + 5570 + 5531 = 11955 \mu\text{s}$$

- when the first packet has arrived, it has to wait for the second large packet and the last, small packet to arrive at B over the last link, which means an additional delay consisting of the transmission delay and propagation for each packet (we assume that processing and queueing delays are “hidden” by the retransmission of the previous packet in the router)¹:

$$(5500 + 6) + (660 + 6) = 5506 + 666 = 6172 \mu\text{s}$$

Then, the total transfer delay for two large packets and one small packet is:

$$11955 + 6172 = 18127 \mu\text{s}$$

Answer: The total transfer time is 18127 μs .

1. Including the processing and queuing delays gives a total transfer delay of 50 μs more, that is, $11955 + (5500 + 20 + 5 + 6) + (660 + 20 + 5 + 6) = 11955 + 5531 + 691 = 18177 \mu\text{s}$

Question 3.

a)

An overlay network is a peer-to-peer network where each node is a peer and each edge (link) a connection, typically created by using tcp.

b)

The statement is false. This is because DNS caching is indeed an important feature of the DNS. It is extensively exploited in order to improve the delay performance and to reduce the number of DNS messages required in order to resolve a hostname, etc., that are sent between servers in the DNS hierarchy.

c)

- i F [SMTP is persistent]
- ii T
- iii T
- iv F [FTP uses out-of-band channel for control packets]

Question 4.

a) The SampleRTT is measured by taking the time it takes to send a packet and receive the ACK for that packet. Time is not measured for retransmitted packets.

DevRTT is the variance of the RTT, measured on the difference between the Estimated-RTT and SampleRTT.

b)

TCP, as defined by breaking down the sentence:

- bytestream-oriented = numbers each byte
- transport layer = operates in the end systems and addresses ports
- reliable delivery = packets are not dropped, not corrupt, not out-of-order, and not duplicates
- segments = a variable number of bytes, determined by the TCP sending entity
- Internet = TCP was developed together with IP so therefore the reliability of TCP complements IP's best-effort service.

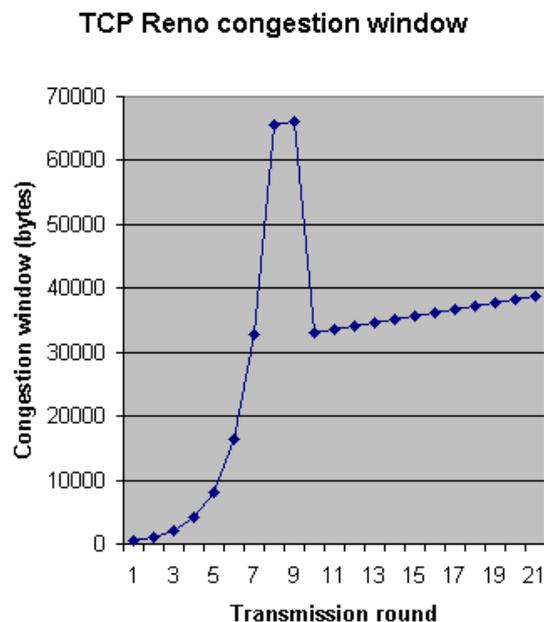
c)

i)

Round no.	CW size	Bytes
1	512	512
2	1024	1536
3	2048	3584
4	4096	7680
5	8192	15872
6	16384	32256
7	32768	65024
8	65536	130560
9	66048	196608
10	33024	229632
11	33536	263168

12	34048	297216
13	34560	331776
14	35072	366848
15	35584	402432
16	36096	438528
17	36608	475136
18	37120	512256
19	37632	549888
20	38144	588032
21	38656	626688
22		655360

The diagram is:



ii)

Slow start is operating in the interval 1–8, while congestion avoidance can be found in 8–9 and the interval 10–21.

Question 5.

a)

A subnet is a physical part of a larger network which is denoted by a single IP address. A subnet mask is used to discover the physical network to which a packet, which has been sent to the single IP address, should be forwarded.

b)

When Alice attempts to establish a TCP connection with Bob, she will send a TCP SYN packet with destination address 138.76.29.7 and a destination port number x. When the NAT receives this packet, it will not be able to know to which internal host in Bob's network it should direct the packet. This is because the NAT does not have an entry for a connection initiated from the outside. The NAT will simply drop the packet.

c)

Typically, one response is expected, because there can (must) be only one node with the matching IP address and MAC address since these addresses are considered unique.

d)

- i F [Tunneling occurs in the same layer]
- ii T
- iii F [IPv6 has no error detection]
- iv F [the field is used for fragmentation purposes]

Question 6.

a)

i)

The order of events is:

J, K (or A), H, D, E, F, A (or K), I, G, B

Note that event C is not used.

ii)

There is just one broadcast domain because all hubs and switches forward link-layer frames. In other words, a frame sent on the broadcast channel (frame destination address with all 1s) can be heard by all nodes.

iii)

There are three collision domains. They exist at each hub because the hubs copy and forwards bits only and operate in the physical layer and thus do not filter frames like the switches do.

b)

Error detection = The error-detection service is used so as to detect bit errors, which can be common because of signal attenuation and electromagnetic noise. A frame that has errors should not be forwarded so many link-layer protocols provide a mechanism to detect the presence of one or more errors. This is done by having the sender set error-detection bits in the frame and then having the receiver perform an error check. In a WLAN the powerful CRC method is used to detect bursts of errors that normally occur in a wireless setting.

Link access = The link-access service specifies the method for how to access the medium in order to transmit a frame. For this purpose, the MAC protocol is used, which can be rather simple or even non-existent (for point-to-point links) or relatively complex (for broadcast media). In the latter case the MAC protocol serves to coordinate the frame transmissions of several nodes. In a WLAN, which is a broadcast medium, the CSMA/CA variant of the MAC protocol is used, which is based on collision avoidance.

Question 7.

a)

Flooding is used by LS do distribute LS packets to all nodes in the network. The nodes need information about all other nodes in order to be able to run the spanning-tree algorithm.

The spanning tree algorithm used in LS is Dijkstra's algorithm, which calculates the least-cost path from all sources to all destinations when a node has received LS packets from all the nodes in the network. Dijkstra's algorithm is used both to find the least-cost path and to create a loop-free representation of the network.

b)

Each row in the table represents the knowledge about the network of one node. Initially, node x knows the distances to its neighbours:

DV table at node z		Cost to				
		u	v	x	y	z
From	u	n/a	n/a	n/a	n/a	n/a
	x	2	n/a	0	7	2
	y	n/a	n/a	n/a	n/a	n/a
	z	n/a	n/a	n/a	n/a	n/a

Next, the neighbours come in with their distance vectors, with information about links connected to each neighbour, and the distance vector table is updated to:

DV table at node z		Cost to				
		u	v	x	y	z
From	u	0	1	2	n/a	n/a
	x	2	3	0	5	2
	y	n/a	2	7	0	3
	z	n/a	5	2	3	0

Now, the information about links one hop away from each neighbour reaches node x:

DV table at node z		Cost to				
		u	v	x	y	z
From	u	0	1	2	3	4
	x	2	3	0	5	2
	y	3	2	5	0	3
	z	4	5	2	3	0

DV table at node z		Cost to				
		u	v	x	y	z
From	u	0	1	2	3	4
	x	2	3	0	5	2
	y	3	2	5	0	3
	z	4	5	2	3	0

c) An autonomous system is a separate routing area, owned by an organization or the like, which lowers the complexity of routing by adding another hierarchy. The owner can opt to run whatever routing algorithm he/she wants in the AS.

RIP is one example of an intra-AS routing protocol, while BGP is an inter-AS routing protocol.

Question 8.

a)

The given protocol does not implement authentication because anyone, such as Trudy, who has Alice's public key K_A^+ can remove the outer layer of encryption from the message. Trudy could then re-encrypt the message with another private/public key pair.

Alice should instead encrypt the message with her private key and then with the public key of the recipient, Bob, like so: $\{\{m\}K_A^-\}K_B^+$. Then only Bob can decrypt the message, using his private key; assuming of course that Bob's private key has not been compromised.

b)

i) SSL is Secure Socket Layer for TCP, a protocol layer that runs over TCP to provide authentication and encryption of connections. Also known as Transport Layer Security (TLS).

ii)

IEEE802.11i is the security framework for WLANs that replaces the WEP standard. It provides for much stronger encryption and includes a separate server to handle authentication as well as key generation and key distribution.

c)

The (application) gateway is an application-specific server through which all application data must pass. That is, it can be used to control access to specific applications, such as the web, e-mail and telnet.