TDTS06: Answers to exam October 23, 2010

Compiled by Juha Takkinen, juha.takkinen@liu.se (2010-12-03)

1. Protocols

   a i A protocol interface can be either the peer-to-peer interface for protocols working in the same layer but on different machines or the service interface provided by the lower-layer protocol to the upper-layer protocol. In the former case the protocol interface is designed around syntax, semantics and timing. In the latter case the services consist of for example primitives such as open and close.

   ii Demultiplexing is the process of extracting the specific data from a channel of mixed data that should go to a specific receiver, e.g. what tcp does based on the port numbers.

   b i The dashed transition handles lost or delayed packets by retransmitting a copy of them and restarting the retransmission timer.

   ii The protocol can handle lost or delayed packets (see i) above) and also packets out of order, because it can distinguish between different sequence number. The protocol can furthermore handle corrupt packets because of the checksum being included in each packet.

2. Networking basics

   a An 8-bit large field can represent a value in the range of 0-255, i.e. a maximum of 255 bytes in one window. The size of the pipe is determined by the bandwidth-delay product. The bandwidth is given so the delaly needs to be calculated, which is done as follows: 255 bytes / 1 Mbps = (255 * 8) / 1,000,000 bps = 0.00204 s or 2.04 ms. The largest delay for the pipe, if filled with 255 B of data, is thus 2.04 ms.

   b Four typical delays are transmission delay, propagation delay, processing delay and queueing delay. (Give examples of each.)

   c Store-and-forward menas that a node must receive the whole packet before retransmitting it. This is required because the node needs to examine each packet and, e.g., recalculate the checksum or read the address and look up the destination for the packet in a forwarding table.

3. Application

   a Assume pipelining and persistency, i.e. that several objects can be requested before receiving the response for the first one from the servera and that several objects can be sent over the same tcp connection. The total time consists of the tcp three-way-handshake, the request of the base document and then the request for the two objects referenced in the base document. Assume the handshake is 20 ms with the request for the base document piggy-backed on the last leg. Assume the base document is half the size of the objects, i.e. the time to transmit the base document is 250 KB / 10 Mbps = 2 Mbits / 10 Mbps = 0.2 s. The time to request and receive the two subsequent objects is 10 ms for the request and then 2 x 500 KB / 10 Mbps = 10 ms + 8 Mbit / 10 Mbps = 0.01 + 0.8 = 0.81 s. the total time in the scenario is then 0.02 + 0.2 + 0.81 s = 1.02 s. See also the figure below.

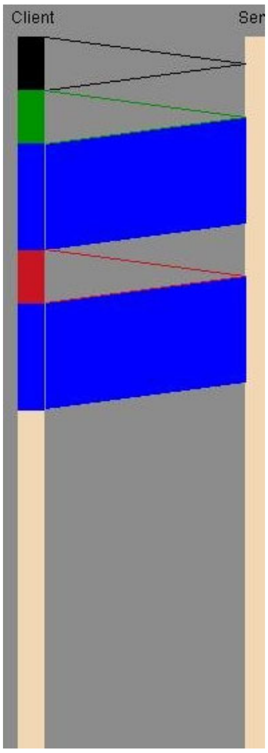

b i The question section lists the name, class and type. The default type is A and the value is nil (which will be filled in with the ip address in the answer section).

ii The hostname in the query, which is the name in the resource record, is static.ak.fbcdn.net. An alternative name for this is the value of the resource record type CNAME in the answer section, i.e. static.ak.facebook.com.edgesuite.com (for which the name in the query is the alias).

4. TCP

a i Assume K = 1024 and M = 1024 KB. Then it will take 11 rounds before the limit of 1 MB is reached, because the sender till transmit 1 KB in the first round, 2 KB in the second (total of 1 + 2 = 3 KB transmitted during over the connection), 4 KB next (total of 1 + 2 + 4 = 7 KB transmitted), etc. until 1024 KB is sent in the 11th round, accumulating a total of 2047 KB over the connection.

    - ii Assume that retransmissions are not counted, only original transmissions. According to the previous question, 2047 KB will have been transmitted after 11 rounds, after which the sender will stay at the maximum speed of 1035 KB per round. The file is 10 MB = 10240 KB large, i.e. 10240 - 2047 = 8193 KB remain to be transmitted after 11 rounds. This means that an additional 9 rounds is required in order to transmit the whole file, because 2047 KB + 9 x 1024 KB = 11263 KB ¿ 10 MB and 2047 + 8 x 1024 KB =10239 KB ¡ 10 MB. (1 KB will be transmitted in the last round.)

  - b
    - i The sequence number of the first acknowledged segment will be 200 + the length of the segment received in bytes = 200 + 2 KB = 200 + 2048 = 2248 (assuming K = 1024)
    - ii Assuming no other reads are performed when the two first segments are received, then the buffer will be completely full when the application process reads 1500 B. This means that there will be the equal amount of free space left in the buffer, i.e. the receiver will put 1500 as the value in its Advertised-window field of the tcp segment.

5. IP

  - a Each fragment will become a new ip packet with a new header, i.e., 20 bytes will be taken up by the header and 480 bytes of data. Assume the 3000 B are data only (no header included); then there will be 7 fragments because 3000 / 480 = 6.25. All fragments will have the same ID in the second word of the header, an offset value telling the receiver where in the original packet that the fragment belongs, and also the MF flag set to 1 (if not the last fragment) or 0 (if the last fragment).
  - b
    - i Machine A will address its ip packet to the web server with 128.119.40.186 as the ip address and 80 as the port number. The NAT router will intercept the packet and replace the source address with the ip address of the NAT router and a unique port number stored in the router's NAT table. The table contains mappings between the internal network's machines and their active port numbers, and the outgoing packets' ip addresses and corresponding port numbers. When the web server sends a response, addressed to the NAT router's ip address and the unique port number created by the router (or manually by the administrator), then the router will replace the address information with machine A's internal ip address and the original port number that machine A used when starting the communication.
    - ii Assuming that the last byte is used to address the hosts in the subnet, then the netmask should be 255.255.255.0.

6. LANs

  - a
    - i K is a random number in the interval $0 - 2^n - 1$, where n is the number of collisions. Because K = 3, there must therefore have been at least 2 collisions, including the latest.

    ii C will be retransmitting the frame after K * 512 bit times, i.e., 3 * 512 / 10 Mbps = 153.6 μs

  b The address contents of the IEEE 802.11 and ip headers in the respective packets are as follows: packet A contains the MAC address of the station as the source and the addresses of the both the access point and the router as the receiver in two other address fields. The source ip address is the station's ip address and the destination is the ip address of the machine at Uppsala University. In packet B, the source MAC address is the address of the access point and thedestination is the MAC address of the router, while the ip addresses are the same as for packet A.

7. Routing

  a An autonomous system (AS), as used in BGP, is a collection of routers with a common network prefix and administered by the same owner, organization or company. It implements a specific routing policy with regard to the other ASs on the Internet.

  b i

Step	N'	D(B), p(B)	D(C), p(C)	D(D), p(D)	D(E), p(E)	D(F), p(F)
0	A	2, A	4, A	7, A	-	5, A
1	AB	d	4, A	7, A	10, B	5, A
2	ABC		d	7, A	10, B	5, A
3	ABCF			6, F	10, B	d
4	ABCFD			d	10, B	
5	ABCFDE				d	

    ii The OSPF routing protocol makes use of Dijkstra's algorithm. In this example, node F will insert the distances to all of its neighbours into the OSPF advertisement packet, i.e. D(A) = 5, D(C) = 6, D(D) = 1 and D(E) = 6. In addition, the packet will also contain a sequence number and a unique ID.

8. Network security

  a WEP is an encryption protocol, originally developed for for 802.11 networks, using an RC4 stream to encrypt packets.

  b i The purpose of MD5 is to create a cryptographic hash of the message, which implements an integrity service. The purpose of IDEA is to create the confidentiality service by encrypting the message.

    ii A nonce is a random number that changes from session-to-session and used in order to counter replay attacks, i.e. avoid the same message being sent twice. PGP does not use nonces because the user can detect duplicate e-mail messages without help from PGP.