TDTS06 Computer networks, January 13, 2011

Answers to the written examination, provided by Juha Takkinen, IDA, juha.takkinen@liu.se.

Question 1.

a) Passive open is usually done by the server. It means that the server moves into the state of listening for incoming connections.

b) (Only the FSM of one party, either the sender or the receiver, is required in the answer.) Assume the channel corrupts messages but delivers them and in order. Furthermore, assume that the client must acknowledge the MD5 sum received from the server. Also, the MD5 is piggy-backed on ACKs.

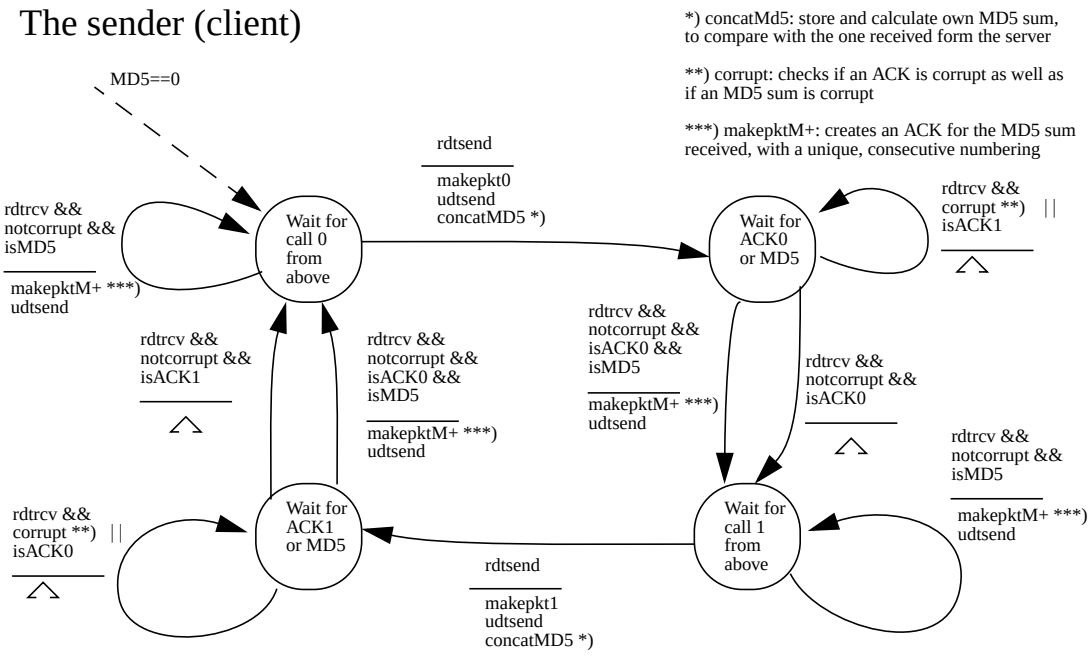

Question 2.

a)

i) The length of the the pipe corresponds to the delay of the link, which is distance / spped-of-light = 385,000 x 1,000 / 3 x 10^8 = 1,28333333 seconds. This is the time to transmit 1 bit of data from Earth to the lunar colony.

The width of the pipe corresponds to the bandwidth, which is 100 Mbps. This gives us a pipe to fill that is 100 x 10^6 x 1,28333333 = 128333333 bits large. This is the amount of data and the answer to the question.

ii) Because the link has many transmission errors and a long delay, the selective-repeat would be preferred. This is because then it is not necessary to retransmit a whole window's worth of data. Instead, only the missing packets need to be retransmitted, thus saving time and bandwidth.

b) Traceroute measures the time in ms between two Internet nodes, as well as the number of hops between them.

Question 3.

a) The statement is fales. FTP preserves state because the client establishes a Control Connection for the duration of an FTP session that typically spans multiple data transfers.

b) Persistent HTTP means that the same TCP connection can be used for several GET commands. Pipelining means that certain commands can be sent back-to-back without waiting for the response from the server.

c) A query is iterative if most of the querying is performed by the same server. A recursive query on the other hand is delegated ot other servers.

Question 4.

a) UDP protects the boundaries between messages in an application. This is because each message received from the application layer is put into a separate UDP datagram intact. TCP would most likely divide the message into several segments, depending on the receiver window size and the congestion window.

b) Slow start is a technique used when a TCP connection is opened, in order to quickly find the appropriate transmission speed. It increases the size of the congestion window exponentially until a threshold is reached (or a loss event).

Congestion avoidance increases the size of the congestion window linearly and is used when slow start reaches the threshold (see above). It stops its increase when a loss event occurs.

Question 5.

a) The sender will receive a response from each DHCP server that can hear the broadcast message from the sender. the sender will then select one of these servers to continue its negotation with.

b) The IP address addresses an interface on the Internet, positioned in the network layer. A router, for example, has one interface for each connected subnet. The hierarchical address space makes it possible to do hierarchical routing, from different network addresses and levels of the upper hierarchy (autonomous systems), down to local subnet and host levels. This decreases the size of routing tables in each router.

c) The statement is true. It describes the fragmentation that can occur at the sending host and when the directly connected network (most likely an ethernet) had a too small an MTU for the TCP segment to fit.

Question 6.

a)

i) One possible order is:

J, A, H, D, E, F, I, K, G, B (C is not used/applicable)

ii) There is one broadcast domain, covering all hosts. Both the hubs and the switch are transparent and can forward broadcast messages sent at the link-layer level.

b) Two examples of control packets are the Beacon frame and the Probe frame. Beacon is used by access points to announce their presence to potential hosts. Probe is used by hosts to listen for available access points.

Question 7.

a) Flooding is required in OSPF in order to distribute a node's link data about its neighbours to all other nodes in the subnet. RIP only needs to send its data to its neighbours, so flooding is not needed in RIP.

b) BGP distributes information about how to reach other autonomous systems (ASs) by using a path vector called AS-PATH (containing BGP attributes) where the AS numbers to reachable ASs are stored. There is also information about the routers and their addresses that can be used to reach the ASs listed in the AS-PATH.

c) One solution is to use poison reverse, i.e., give the cost of "infinity" to nodes from where the original cost to the lost node came from.

Question 8.

a)

i) A message digest is a hashed representation of a message, used for checking the integrity of the message.

ii) Symmetric encryption, as compared to asymmetric encryption, uses the same key for encryption as well as for decryption.

b) An active intruder can alter and retransmit messages that have been encrypted, while a passive intruder only can listen in on the encrypted communication.

c) No, it should not be possible to decrypt a hashed message. This is because the hashing function should be one way only; this is the definition of the function. However, it can be possible to get the same hash for two different messages (a collision) if the hasing function is poorly designed.