CS 421: Computer Networks **SPRING 2015**

MIDTERM
April 9, 2015
120 minutes

Name: ______________________________

Student No: ___________________________

Show all your work very clearly. Partial credits will only be given if you carefully state your answer with a reasonable justification.

Q1	
Q2	
Q3	
TOT	

1)

a) A user at node A uses FTP to upload four files to node B before terminating the FTP session.

   i. (4 pts) Which of the two nodes is the client and which one is the server in this scenario? Explain briefly.

   ii. (5 pts) How many TCP connections in total are established during the FTP session? Explain briefly.

   iii. (5 pts) If we want to draw an analogy between FTP and HTTP, in what sense FTP establishes persistent connections, and in what sense it establishes non-persistent connections? Explain briefly.

b) (6 pts) Suppose you click on a link at your Web browser to download a Web page. Suppose that the IP address for the associated URL is not cached in your local host, so that a DNS look-up is necessary to obtain the IP address. Suppose that your local DNS server uses iterative queries and it contacts N DNS servers before sending your host the IP address of the requested Web server. Each successive contact incurs a delay of RTT_i, $i = 1, \ldots, N$. Further assume that the Web page associated with the link contains exactly one object, a small HTML file with no other referenced objects. Let RTT_0 denote the delay between the local host and the Web server containing the object. Assuming that the transmission time of the object can be neglected, how much time elapses from the time when the client clicks on the link until the client receives the object?

c) (6 pts) TCP is a reliable transport layer protocol and the checksum field in a TCP segment is used towards ensuring reliability. But why does a UDP segment have a checksum field? Explain briefly.

d) (6 pts) Consider a TCP server-side socket that is used to communicate from server to client. Can data being read from the server-side socket have been sent by more than one client? If this were a UDP socket, can data being read from it have been sent by more than one client? Explain briefly.

2)

a) (12 pts) In the throughput analysis of various transport layer protocols, we have always assumed that the transmitter was saturated. We wish to remove this assumption now. More specifically, consider a Stop-and-Wait protocol with segment transmission time t_I and one-way propagation delay t_{prop}. Suppose that segments become available for transmission with probability $q < 1$ when the transmitter is ready to transmit a new segment. If there is no new segment to transmit when the transmitter is available for transmission, the transmitter rechecks the availability of a new segment every t_{check} time units until a new segment becomes available for transmission. (In other words, the number of trials it takes a new segment to become available at the sender for transmission after the reception of the ACK for the previous segment is **geometrically distributed** where t_{check} time units is spent between successive trials.) Calculate the maximum achievable throughput in this case assuming that there are no errors or loss, acknowledgement segments are negligibly small, and processing and queueing delays are ignored.

b) (12 pts) Consider a 1 Mbps channel with a 20 msec one-way propagation delay, i.e., 40 msec roundtrip propagation delay. We want to transfer a file of size 13500 Bytes. Each segment has a total size of 1625 Bytes including the 125 Bytes header, i.e., each segment contains 1500 Bytes of data. When there is data to be transmitted, each segment contains the maximum number of bytes. Assume that ACK segments are of 125 Bytes long and there is a processing delay of 1 msec after a segment is fully received at the receiver until the transmission of the corresponding ACK is started. **Go-Back-N** protocol is used with a window size of N = 4 segments. Assume that **every 6th segment** crossing the forward channel is lost while ACKs are not lost or corrupted. Assume that the processing delay at the sender after an ACK is received is negligible. How much time is required to complete the transfer of the whole file and receive the final ACK at the sender? Assume that the timeout for each window is 50 msec starting from the time when the sender sets the window.

c) (6 pts) Consider a transmission link running at 1 Gbps (1×10^9 bps). Assume that the speed of propagation is 3×10^5 km/s. What is the maximum distance between a TCP source and destination so that a window of size 100,000 Bytes can be fully utilized?

d) A sender has sent out a total of 8 data segments each containing 2000 Bytes using a TCP-like protocol. This protocol assigns the sequence and ACK numbers in the same way as TCP does, but has the capability of retransmitting segments without waiting for timeout or triple duplicate ACKs. Assume that the sender's sequence number for the first segment is 5001.

i. (3 pts) The sender subsequently receives two segments with ACK 9001. Given this event, which of the segments does the sender conclude might be lost and thus should be retransmitted?

ii. (3 pts) After retransmitting the segment in part i., the sender receives an ACK 15001 in response. Based on this, which of the original 9 segments can the sender conclude might be lost, if any?

3)

a) We would like to transfer a 6 MBytes ($6x10^6$ Bytes) file using TCP. Assume that the sender is in the slow start phase at the beginning of the file transfer and the slow start threshold is set to 10 MBytes. Each TCP segment has a size of 1000 Bytes, and recall that the sender starts with CongWin = 1000 bytes in the slow start phase. Assume further that TCP has been modified to allow very large windows so that the TCP receive window is always 4 MBytes ($4x10^6$ Bytes) during the entire connection. The distance between the sender and the receiver is 6000 km. Assume that the speed of propagation is 3×10^5 km/s. The transmission speed of the connection is 100 Mbps (100×10^6 bps). For simplicity, ignore processing and queueing delays.

   i. (8 pts) Suppose no segments are lost. How long does it take to transfer the entire 4 Mbytes file?

   ii. (6 pts) What is the average channel utilization, U, during the transfer of the file?

   iii. (8 pts) Redo part i. if the TCP receive window is always 1 MBytes ($1x10^6$ Bytes) during the entire connection.

b) (10 pts) Consider the case of two TCP connections sharing a single link with transmission rate of R bps, as shown in the figure below. Assume that two connections have the same MSS, but $RTT_1 = RTT_2 / 2$. Also assume that no other TCP connections or UDP datagrams use the shared link and no packets are lost except for the bottleneck router. Suppose that both TCP connections are operating in the AIMD (Additive Increase/Multiplicative Decrease) mode at all times.

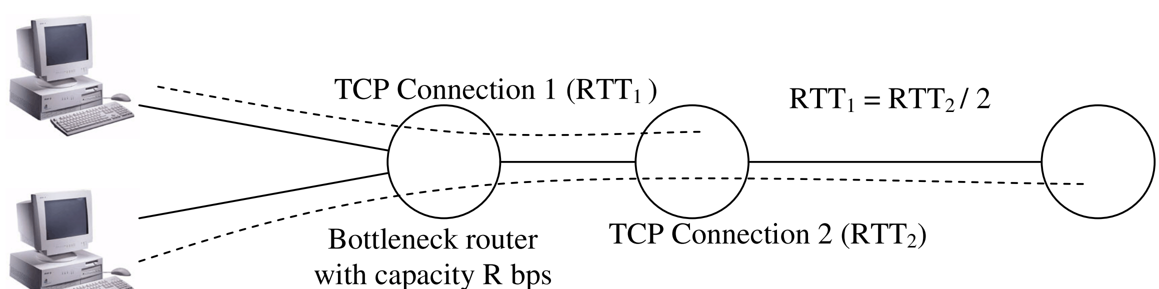

Similar to what we did in class, use the following figure in order to plot the time evolution of the joint throughputs realized by two connections assuming that the joint throughput initially starts at point A.