BIRLA INSTITUTE OF TECHNOLOGY & SCIENCE, PILANI
Department of Computer Science & Information Systems
Second Semester 2022-2023
COMPUTER NETWORKS (CS F303)
Mid Semester Examination (CLOSED BOOK)

Date of Exam: 13-03-2023 (AN1) **Time: 2:00 – 3:30 PM** **MM: 50** **Duration: 1.5 Hrs.**

Note: Attempt subparts of questions in sequence, one after the other.

Q.1 Consider the networks shown in the figure below. There are two user machines, *m1.a.com*, and *m2.a.com,* in the network *a.com*. Suppose the user at *m1.a.com* types in the URL *www.b.com/file.htm* into a browser to retrieve a 1Giga bit file from *www.b.com*. **[9+4=13M]**

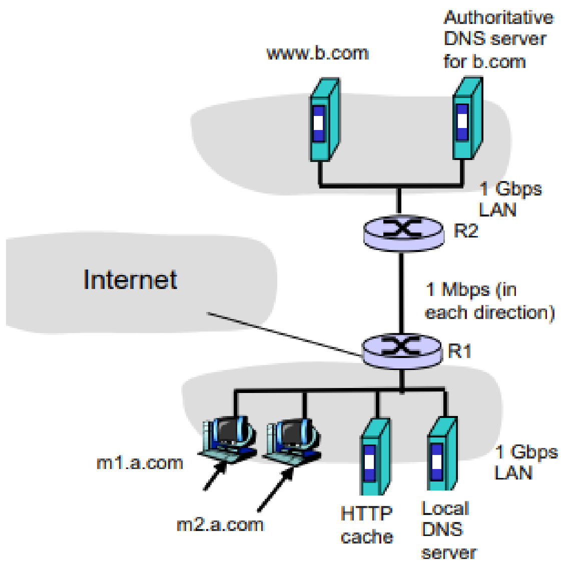

a) List the sequence of DNS and HTTP messages sent/received from/by m1.a.com as well as any other messages that leave/enter the a.com network that are not directly sent/received by m1.a.com from the point that the URL is entered into the browser until the file is completely received. For each message, indicate the source, destination, purpose, and time it takes to reach the destination. You can assume that every HTTP request by *m1.a.com* is first directed to the HTTP cache in *a.com,* that the cache is initially empty, and that all DNS requests are iterated queries at the local DNS server. You can also assume that as soon as the first few bits of the file arrive at R2, R1, and HTTP cache, R2, R1, and HTTP cache, forward them to the next link.

Consider the following assumptions for time/delay calculations:

- The packets containing any DNS commands and HTTP commands such as GET are very small compared to the size of the file, and thus their transmission times (but not their propagation times) can be neglected.
- Propagation delays within the LAN are small enough to be ignored.
- The propagation delay from router R1 to router R2 is 100 milliseconds.
- The propagation delay from anywhere in the *a.com* network to any other site on the Internet (except *b.com*) is 500 milliseconds.
- The propagation delay from anywhere in *a.com* to anywhere in *b.com* is 100 milliseconds.

Write your answer in the table format shown below. One sample entry is shown here.

S.No.	Message type with source and destination information	Purpose of the message	Amount of time it takes to reach the destination
1	*a.com* sends DNS requests to the local DNS server	To resolve the IP address of *m1.a.com*	100 milliseconds

b) Now assume that machine *m2.a.com* requests the same URL that *m1.a.com* made. List the sequence of DNS and HTTP messages sent/received from/by *m2.a.com* and any other messages that leave/enter the *a.com* network that are not directly sent/received by *m2.a.com* from the point that the URL is entered into the browser until the file is completely received. Write your answer in the table format used for (a).

Q.2 Consider a scenario in which Hosts A and B want to send messages to Host C. Hosts A and C are connected by a communication channel that can lose and corrupt messages. Hosts B and C are connected by another channel (independent of the channel connecting A and C) with the same properties.
The transport layer at Host C should alternate in delivering packets from A and B to the layer above (i.e., it should first deliver the data from a packet from Host A, then from a packet from Host B, and so on).
Design a stop-and-wait-like error-control protocol for reliably transferring packets from Hosts A and B to Host C, with alternating delivery at Host C as described above. Give FSM descriptions of Hosts A, B, and C and specify events and actions for each state. **[12M]**

Q.3 Two hosts in a voice-over-IP session are connected by a path of 4 routers. All links run at 1 Megabits/second, separating the hosts by 3000 Kilometers. All packets are of size 1500 Bytes. Assume the bit propagation speed is 2×10^8 meters/second. (1 Kilo Bytes = 1024 Bytes, but 1Megabit/second = 10^6 bits/second.) **[1+1+2+2+3=9M]**

a) What is the minimum round trip time (RTT), assuming no queueing delay and taking processing time at each host is negligible?

b) For this part only, let us assume that one router on the path has a steady queue occupancy of 5 packets. What is the one-way end-to-end delay in this case?

c) Now, let us assume the maximum queue occupancy for every router queue is 5 packets. What is the maximum one-way end-to-end delay?

d) For part (c), how long should the playback buffer be at the destination voice-over-IP client if each packet arrives successfully without being dropped? The playback buffer should never run empty to provide a pleasant playback experience. Express your answer in bytes.

e) For part (c), how long should the playback buffer be at the destination voice-over-IP client if each packet arrives successfully when first transmitted or is dropped the first time and then arrives successfully on the second attempt? Assume the sender retransmits a packet only after it finds out the packet was dropped, after twice the maximum one-way end-to-end delay. Express your answer in bits.

Q.4 Suppose Host A wants to open a connection to Host C. A proxy server B along the path can terminate A's connection to itself and open a connection to C. So, in this case, there are now two TCP connections, A to B and B to C. A thinks it is sending data to C, but B is processing the TCP segments and sending acknowledgments back to A, spoofed from B's IP address. Simultaneously, B opens a TCP connection to C, pretending to be A. **[2+2+3+3=10M]**

Consider the network shown in the figure, where the RTT from A to B and B to C is 50 milliseconds each, and there is no packetization, queueing, or processing delay, such that the RTT from A to C is 100 milliseconds. The maximum segment size is 1400 bytes. A sends an infinite stream of bytes, such that every segment is the maximum segment size.

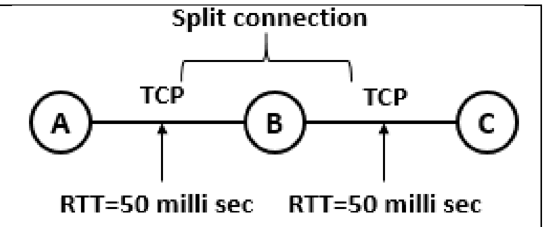

a) Suppose B does not split the TCP connection, so packets flow directly from A to C through B. The route between A and B drops 10% of data segments and does not drop acknowledgments, while the route between B and C do not drop any packet. What will the TCP throughput from A to C be?

b) Suppose B does split the connection, such that packets flow from A to B, terminate at B, then are forwarded in separate flow from B to C. The route between A and B drops 10% of data segments (no acknowledgments drop), while the route between B and C do not drop any packet. What will the throughput from A to C be?

c) Suppose B does split the connection, such that packets flow from A to B, terminate at B, then are forwarded in separate flow from B to C. The route between A and B drops 10% of packets, and between B and C also drops 10% of packets. What will the throughput from A to C be?

d) Finally, suppose B does not split the connection, so packets flow from A to B, passing through but not terminating at B. The route between A and B drops 10% of data segments, and the route between B and C also drops 10% of data segments. What will the throughput from A to C be?

Q.5 Consider that only a single TCP Reno connection uses one 10 Gigabits/second link, which does not buffer any data. Suppose this link is the only congested link between the sending and the receiving hosts. Assume that the TCP sender has a huge file to send to the receiver and the receiver's receive buffer is much larger than the congestion window. Each TCP segment size is 1500 Bytes. The two-way end-to-end delay of this connection is 150 milliseconds. The TCP sender detects all packet losses by three duplicate ACKs. **[2+2+2=6M]**

a) What is the maximum window size (in segments) that this connection can achieve?

b) What is the average window size (in segments) and average throughput (in bits/second) of this TCP connection?

c) How long (in seconds) would it take for this TCP connection to reach its maximum window size again after recovering from a packet loss?

xxx—00—xxx