Birla Institute of Technology & Science, Pilani
Department of Computer Science and Information Systems
First Semester 2023-24
Advanced Computer Networks (CS G525)
MID SEMESTER TEST (CLOSE BOOK)

Duration: 1.5 Hrs. **Date: 11-10-2023** **4:00 – 5:30 PM** **MM: 25**

Note: Answer sub-parts of a question (if any) at one place in the sequence.

Q.1 According to the end-to-end principle of system design, local implementations (i.e., lower layer implementation) of functions may enhance performance above that achievable using end-to-end implementations alone. Describe the performance benefits of localized implementation of "error control" (i.e., error detection and error correction) in the Internet. **[2M]**

Q.2 Assume host A sends a large file to host B over a TCP-Reno connection. The connection's roundtrip time (RTT) is 100 mili seconds, and the bottleneck link bandwidth is 10 Gigabits per second. Assume the path MTU is 1500 Bytes, and host A sends all packets with the maximum possible size. You can assume ACKs arrive within RTT. Since ACK packets are very small, you can ignore the bandwidth they consume. Also, ignore the link layer header in your calculations. *Note: Round down the non-integer numbers in your calculations.* **[2+1+4=7M]**

a) Assume an infinite initial threshold, no losses, and competing traffic; how long (in seconds) would the normal slow start mechanism take to achieve 100% utilization?

b) After reaching the full utilization of the link, how much time (in seconds) will the connection take to arrive once again at the full utilization of the link?

c) Assume another host, C, sends a large file to host B over a TCP-AIAD connection through the same bottleneck link of 10 Gigabits per second. The RTT and path MTU values are the same as the TCP-Reno connection described above. For the TCP-AIAD connection, every drop decreases the congestion window by 15000 packets (to a minimum of 1 packet). The increment value remains the same as in TCP-Reno. Both connections start together from a slow start and continue to a slow start until the bottleneck link capacity is exhausted. You should assume that each connection updates its congestion window after each RTT, decreasing it if any packet drops and increasing it otherwise. Packet loss occurs only when inflight data on the link exceeds the link capacity.

i) How many RTTs does it take for both connections to observe packet loss for the first time, and what are each connection's congestion window sizes (in packets) then?

ii) What are each connection's congestion window sizes (in packets) when they observe packet loss for the third time?

Q.3 Answer the following questions. **[1.5x4 = 6M]**

a) Does Open-Flow Software Defined Networking simplify the network switches and make them future-proof? Explain.

b) TCP suffers from retransmission ambiguity, while QUIC does not. Explain.

c) Explain the working of the BBR with the help of its state machine representation.

d) The NDN architecture inherently supports congestion control and multicasting, whereas Internet architecture does not. Explain with a suitable example.

Q.4 Answer this question based on Internet Architecture and Mobility aspects. **[1+1+2+1=5M]**

Consider an Internet connection between host A and server B using the standard Ethernet/IP/TCP stack, where A is served by a router R1.

a) Assume an ongoing TCP flow between host A and server B. If A were to move to a new network served by router R2, what part of the stack fails in its operation, and what is its main reason?

b) Between layering and end-to-end principles, which principle was violated in the design choice responsible for the lack of connection-oriented mobility and why?

c) Suggest a simple solution that fixes the mobility problem and still preserves the layering and end-to-end principles. In keeping with the end-to-end principle, you are allowed full freedom in the design of the endpoints but not in the network components (routers).

d) Which design principle of NDN architecture provides support for node mobility? Explain in brief.

Q.5 The topology shown in the below figure comprises three switches, S1, S2, and S3, and four Hosts, H1, H2, H3, and H4, having physical addresses MAC-H1, MAC-H2, MAC-H3, and MAC-H4 respectively. Host H1 is an HTTP server. What flow rules will be pushed at S1 to meet the traffic management requirements below? Assume S1 has two flow tables, one for policy rules and another for forwarding rules. Write the flow rules optimally (i.e., the minimum number of rules with the minimum number of header fields matching requirement) to meet the given policy and forwarding requirements. The switch S1 ports are labeled with eth-0, eth-1, eth-2, and eth-3. **[5M]**

Policy Requirements:

- H4 can only access the HTTP server at H1
- All types of communication are allowed between H2 and H3
- No other permitted communication.

Forwarding requirements:

- H1 to H4 communication should be through S3
- The rest of the communication (as the policy allows) happens on the shortest path between the communicating nodes.

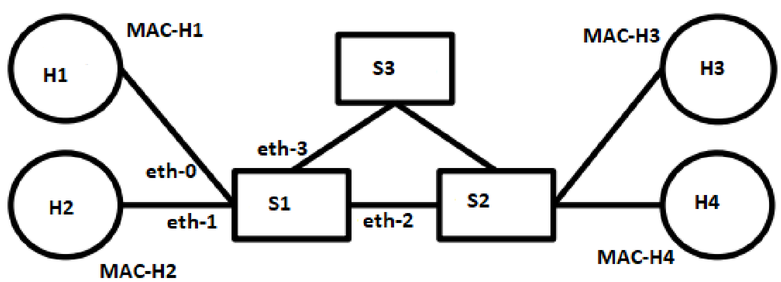

xx—00—xx