Birla Institute of Technology & Science, Pilani, Rajasthan
Department of Computer Science and Information Systems
First Semester 2022-2023
Advanced Computer Networks (CS G525)
Comprehensive Exam (CLOSE BOOK)

Duration: 3.0 Hrs. | **Date: 30-12-2022 (AN)** | **Maximum Marks: 35**

Note: Attempt subparts of a question in sequence.

Q.1 The overlay network shown in **Fig.1** is a Distributed Hash Table (DHT) with 16 nodes. Each node is labeled with the first value in its range values (e.g., B is responsible for hash values 100-199). The routing table for node **J** is shown in the figure. Node **J** has routes to the node that is 1, 2, 4, and 8 hops away. Assume that all nodes have routing tables that are configured similarly. Suppose the client shown in **Fig.1** sends a get request to node **H** with a key string of "**bitspilani**", and that hash("**bitspilani**") = 513. **[1+1+1+1=4M]**

a) List the servers through which this request would pass, assuming that the key string does not appear in any node's cache.

b) What servers would the request pass through if the key string appears in node M's cache?

c) Suppose each DHT node receives a get request for "**bitspilani**" about once per second. If the system is operated without caches, how many requests per second must the "responsible server" process?

d) Suppose each DHT node receives a get request for "**bitspilani**" about once per second and caching is enabled at all nodes, and each cache entry expires 60 seconds after being placed in the cache, approximately how often does the responsible server receive a get request from another server?

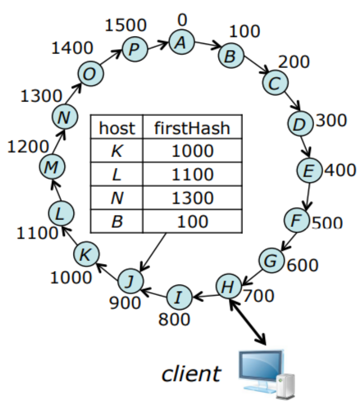

Fig.1

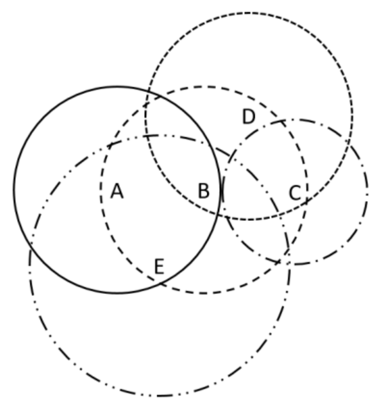

Fig.2

Q.2 Consider the topology of wireless nodes A, B, C, D, and E shown in **Fig.2**. Further, assume that losses only occur due to collisions. Each wireless node's transmission range is shown by a circle surrounding it as the center of the circle.

a) What are the hidden nodes for E to B transmission? **[1+1+1=3M]**

b) What are the exposed nodes for B to D transmission?

c) If A sends data to E and B sends data to D (as fast as they can) and nodes use CSMA protocol for channel access. What would be the throughput of their transfer as a proportion of their sending rate?

Q.3 Suppose two TCP connections share a path through a router **R**. The router's queue size is six segments. Each connection has a stable congestion window of three segments. Assume no congestion control is used by these connections. A third TCP connection is now attempted through **R**. The third connection does not use congestion control either. **[2+1+1=4M]**

a) Describe a scenario in which, for at least a while, the third connection does not get a share of available bandwidth, and the first two connections proceed with 50% of bandwidth each.

b) Does it matter if the third connection uses slow-start? Give reason.

c) How does congestion control on the part of the first two connections help to solve this?

Q.4 Suppose a router's drop policy is to drop the highest-cost packet whenever queues are full. The cost of the packet is calculated as the product of its size and the time remaining that it will spend in the queue. **[1+1.5+1.5=4M]**
a) What advantages and disadvantages might such a policy offer, compared to tail drop?

b) Give an example of a sequence of queued packets for which dropping the highest-cost packet differs from dropping the largest packet.

c) Give an example where two packets exchange their relative cost ranks as time progresses.

Q.5 The local retransmissions by the base station enlarge the RTT value and thus affects the RTT estimation for a TCP connection. **[1+3=4M]**
a) What are the consequences of RTT enlargement?

b) Suggest two solutions to protect the TCP source from the RTT enlargement problem. Express the limitations of the proposed solutions.

Q.6 This question explores the scalability of peer-to-peer systems, compared to client-server systems. Assume that the interior of the network has ample bandwidth and that the propagation delay is infinitesimal. The server has a 100 KByte file that it wants to distribute to a group of 250 receivers. All hosts have bidirectional 40 Kbit/second links to the Internet and the core of the Internet is not congested. **[1+1.5+1.5=4M]**
a) What is the minimum time for the server to transmit the data to all of the receivers in a client-server configuration?

b) Suppose that the receivers can upload data, too, but only after receiving an entire copy of the 100 KByte file. What is the minimum time for the server and the cooperating peers to transmit the data to all receivers in this configuration?

c) Suppose that a receiver can start uploading data to others after receiving the first 20 KByte chunk of data. How long does it take to deliver the data to all receivers?

Q.7 a) Contrast the use of a single receive buffer at the connection level and a separate receive buffer for each subflow for an MPTCP connection. **[1+1.5+2.5=5M]**

b) Contrast the static data distribution (divide application data across subflows in fixed byte sequences) and dynamic data distribution (application data can be sent to any subflow, i.e., no fixed binding between a subflow and byte sequences) over multiple subflows of an MPTCP connection.

c) Describe the behavior of the following MPTCP congestion control algorithm for the shared bottleneck (subflows share a bottleneck), disjoint subflows scenarios, and for meeting the MPTCP design objectives (three design objectives were discussed in the class).

For each ACK received in subflow r, we increase congestion window $\mathbf{w_r}$ of subflow **r** by the amount:

$$\min\left(\frac{\max_{i\in\mathcal{R}_u} w_i/\mathrm{rtt}_i^2}{\left(\sum_{i\in\mathcal{R}_u} w_i/\mathrm{rtt}_i\right)^2}, \frac{1}{w_r}\right)$$

Where $\mathbf{rtt_i}$ is the round-trip time on the path **i** and $\mathcal{R}_u$ is the set of all paths available.

When a loss event occurs for subflow **r**, we decrease the congestion window by the amount: $\mathbf{w_r/2}$.

Q.8 Answer the following questions. **[1+1.5+1.5+1+2=7M]**
a) What is the significance of pacing gain in BBR protocol? Explain.

b) Name Data Networking (NDN) architecture separates the Routing and Forwarding planes. What does it mean? How is it beneficial for meeting QoS requirements?

c) In IP-based Internet, web proxies or HTTP proxies cache the content for faster delivery to the user. Similarly, NDN architecture also allows content caching. Contrast both types of caching mechanisms.

d) In Dynamic Source Routing (DSR) protocol, how does a node decide the next hop for forwarding a data packet?

e) We can abstract the view of computer networks in terms of three interfaces, i.e., Host-Network, Operator-Network, and Packet-Switch. Contrast between Openflow enabled Software Defined Networks (SDN) and Multi-Protocol Label Switching (MPLS) enabled Internet based on above mentioned three interfaces.

xx—oo—xx