Lecture 11: TCP Implementation Quiz ANS

1. In real TCP, what is the primary unit used for sequence numbers and ACKs?
A. Packets
B. Bytes
C. Segments
D. Frames

ANS: B – TCP numbers individual bytes in the bytestream; packets (segments) are defined by the sequence number of their first byte and ACKs refer to byte numbers.

2. What does the sequence number field in a TCP header represent for a data segment?
A. The number of the last byte carried in the segment
B. The count of bytes transmitted so far on the connection
C. The number of the first byte of data in the segment
D. The total number of segments sent so far

ANS: C – By definition, the TCP sequence number of a segment is the sequence number of the first byte contained in that segment.

3. Suppose a TCP sender transmits a segment with sequence number j carrying B bytes. If all prior data has been received in order, what ACK value will the receiver send after correctly receiving this segment?
A. j
B. j + B − 1
C. j + B
D. j + 2B

ANS: C – Cumulative ACKs report the next expected byte, so after receiving bytes j through j+B−1, the receiver ACKs j+B.

4. In normal operation with one packet in flight, no loss, and fixed segment size B, how are the ACK number and next sequence number at the sender related?
A. ACK number is always smaller than the next sequence number
B. ACK number equals the next sequence number
C. ACK number is always larger than the next sequence number
D. There is no simple relationship

ANS: B – With a single in-flight segment and no loss, the receiver ACKs the next expected byte, which is exactly the sender's next sequence number.

5. Why does TCP choose a random Initial Sequence Number (ISN) for each new connection?
A. To minimize header overhead
B. To ensure that segments are aligned to MSS boundaries
C. To improve security and avoid confusion with stale segments from old connections

D. To simplify retransmission timeout estimation

ANS: C – Random ISNs reduce the risk that old duplicate segments will be mistaken for new data and mitigate some spoofing attacks.

6. If all bytes up to and including ISN+219 have been received correctly, which ACK value will the receiver send next?
A. ISN + 219
B. ISN + 220
C. ISN + 218
D. ISN + 221

ANS: B – The receiver ACKs the next expected byte, which is one more than the highest contiguous byte received, here ISN+220.

7. Which combination of TCP header fields is primarily responsible for implementing reliability?
A. Source port and destination port
B. Sequence number, acknowledgment number, and checksum
C. Advertised window and header length
D. Flags and urgent pointer
ANS: B – TCP uses sequence numbers and acknowledgment numbers to track which bytes have been sent and received, and a checksum to detect corrupted segments, together providing reliability.

8. If an ACK for byte 100 is lost but later an ACK for 120 arrives. What can Host A infer from the ACK=120?
A. Only bytes 100–119 were received; 92–99 must be retransmitted
B. All bytes up to 119 have been received, including 92–99 and 100–119
C. Only the segment starting at 120 was received
D. The connection should be reset

ANS: B – An ACK of 120 is cumulative and implies that every byte with sequence number less than 120 has been successfully received, covering the earlier lost ACK.

9. TCP connections are full-duplex. What does this mean in practice?
A. Data can only flow in one direction at a time, but the direction can switch
B. Both hosts can send and receive data simultaneously over the same connection
C. The connection uses separate ports for sending and receiving
D. Each segment can carry data or acknowledgments, but not both

ANS: B – Full duplex means each endpoint maintains its own bytestream and sequence numbers, allowing simultaneous bidirectional data transfer within one TCP connection.

10. In a full-duplex TCP connection, how many independent sequence-number spaces are used?

A. One, shared by both directions
B. Two, one for each direction A→B and B→A
C. Four, one per host and per port
D. One per segment

ANS: B – Each direction of the connection has its own sequence numbers and ACKs, so there are two independent sequence-number spaces.

11. What is the main purpose of the three-way handshake in TCP connection setup?
A. To exchange congestion window sizes
B. To synchronize IP addresses between client and server
C. To allow each host to learn the other's ISN and establish ISNs (initial sequence-numbers)
D. To negotiate MTU along the path

ANS: C – The SYN, SYN-ACK, ACK exchange communicates each side's ISN so both can initialize their send and receive sequence spaces.

12. During normal TCP connection termination, what does a FIN flag indicate?
A. The sender will neither send nor receive any more data
B. The sender is done sending but will continue to receive data
C. The sender is aborting the connection immediately
D. The sender is requesting a higher advertised window

ANS: B – A FIN says "I'm done sending" but the connection can remain half-open so the peer can still send data back until it also sends FIN.

13. How does the behavior of a TCP RST segment differ from a FIN segment?
A. RST must be acknowledged while FIN does not
B. RST closes only one direction; FIN closes both
C. RST immediately aborts the connection without requiring an ACK, discarding in-flight data
D. RST is used only during connection setup

ANS: C – RST is an abrupt termination: it does not require an ACK and indicates that the sender will no longer send or receive data, dropping in-flight segments.

14. In the TCP state machine, which state represents a connection that is fully open and exchanging data?
A. SYN_SENT
B. LISTEN
C. ESTABLISHED
D. TIME_WAIT

ANS: C – The ESTABLISHED state is where regular data transfer occurs with both sending and receiving.

15. In TCP's byte-based sliding window, a sender maintains a single retransmission timer rather than one timer per segment. Under what condition does the sender retransmit without waiting for the timer to expire?
A. When any ACK is received for a byte inside the current window
B. When three duplicate ACKs for the same byte are received
C. When the advertised window drops to zero
D. When the congestion window exceeds the receiver's advertised window

ANS: B – TCP uses two loss-detection triggers: timer expiration and three duplicate ACKs. Receiving three copies of the same cumulative ACK signals that a later segment arrived but an earlier one is missing, so the sender immediately retransmits the first unacknowledged segment (leftmost byte of the window) without waiting for the timer.

16. What is "piggybacking" in the context of TCP?
A. Sending data without any ACKs
B. Delaying data segments until the congestion window grows
C. Combining ACK information with outgoing data in the same segment
D. Using multiple connections over the same port

ANS: C – Piggybacking means carrying the ACK for received bytes in the header of a data segment being sent in the opposite direction, instead of sending a pure ACK.

17. Why are SYN-ACKs always piggybacked in TCP?
A. Because data must always be sent with SYN segments
B. Because the OS, not the application, performs the handshake and can easily combine SYN and ACK in one control segment
C. To reduce the number of sequence-number spaces needed
D. To avoid triggering congestion control

ANS: B – During the handshake, the OS sends a segment that both acknowledges the client's SYN and carries its own SYN; this is a natural piggybacking of control information.

18. In the TCP sliding-window scheme (measured in bytes), what does the window [first_unacked, first_unacked + W) represent?
A. All bytes that have been successfully sent and acknowledged
B. The set of bytes that are in flight but not be acknowledged yet
C. The total buffer space at the receiver
D. The advertised congestion window only

ANS: B – The window is the contiguous range of W bytes starting at the first unacked byte; these are the only bytes the sender may have outstanding in the network.

19. Which events cause the sliding window to move (slide) forward?
A. Any ACK that is received for any byte within the window

B. ACKs that acknowledge the leftmost unacked bytes, thereby advancing the first unacked byte
C. Timer expirations only
D. Arrival of out-of-order data at the receiver

ANS: B – The window boundary is tied to the first unacked byte; only when those earliest bytes are cumulatively ACKed does the window slide right.

20. TCP sets its effective send window W as the minimum of two quantities. Which two?
A. Congestion window and MSS
B. Receiver advertised window and MSS
C. min(RWND, CWND), Congestion window and receiver advertised window
D. RTT estimate and MSS

ANS: C – The usable window is min(RWND, CWND): RWND prevents overrunning receiver buffers, and CWND limits injected traffic to avoid network congestion.