ECSE-6600: Internet Protocols

Informal Quiz #07

Shivkumar Kalyanaraman:
GOOGLE: "Shiv RPI"
shivkuma@ecse.rpi.edu

UDP/TCP: Informal Quiz

UDP/TCP

- ☐ ☐ UDP provides congestion control services to applications using it.
- ☐ ☐ UDP provides multiplexing and demultiplexing services to applications using it.
- ☐ ☐ UDP provides error correction services to applications using it.
- ☐ ☐ UDP is minimally necessary because IP does not provide port-number space for applications.
- ☐ ☐ UDP is a connection-oriented protocol
- ☐ ☐ UDP uses a 32-bit CRC as its error detection method.
- ☐ ☐ TCP can re-assemble IP fragments
- ☐ ☐ TCP provides a reliable packet-stream service
- ☐ ☐ TCP provides reliability by using strong FEC codes
- ☐ ☐ Today's versions of TCP use ACKs, duplicate ACKs and SACKs
- ☐ ☐ TCP can recover from both packet loss and packet duplication in the network
- ☐ ☐ TCP uses a 32 bit sequence number for its reliability mechanisms
- ☐ ☐ TCP uses explicit NAKs in its window-based recovery and flow-control scheme
- ☐ ☐ Even the 3-way handshake in TCP is robust to packet losses and duplication.
- ☐ ☐ Path-MTU refers to the procedure of finding the minimum MTU of the path to reduce the probability of fragmentation.

TCP

- ☐ ☐ The link MTU is equal to the TCP MSS after the path-MTU procedure.
- ☐ ☐ TCP is connection-oriented because it explicitly sets up a path from source to destination via a signaling protocol (I.e. the SYN/ACK mechanism)
- ☐ ☐ Selective repeat is significantly more efficient than go-back-N only for occasional packet losses and large window sizes.
- ☐ ☐ Go-back-N requires out-of-order buffering support at receivers.
- ☐ ☐ TCP handles variable RTT by setting its timeout variable to a large, fixed value
- ☐ ☐ TCP is called “self-clocking” because the source sends traffic whenever it likes
- ☐ ☐ TCP by default uses a selective retransmission policy
- ☐ ☐ The RFC 793 RTT estimator could only tolerate variances of upto 30%
- ☐ ☐ The TCP congestion control algorithm is stable because it detects congestion reliably and its rate of window decrease is faster than its rate of window increase
- ☐ ☐ TCP’s use of cumulative acks reduces the need for any timeout/retransmission of acks
- ☐ ☐ Karn’s algorithm would be triggered often on a wireless or radio link which is very lossy
- ☐ ☐ A two-way handshake is sufficient for the robust setup of a half-duplex connection, but a three-way handshake is necessary for the robust setup of a full-duplex connection

TCP

- ☐ If timeouts are not used, in general, packet or ack-losses cannot be recovered from
- ☐ A duplicate ack gives the same information as a NAK, but it presumes the notion of a sequence number
- ☐ Sequence numbers allow the detection of duplicate packets, but the sequence number space must be sized sufficiently large compared to the window size depending upon the retransmission algorithm (go-back-N or selective-repeat) used.
- ☐ In a lossless network, window-based transmission can achieve full utilization
- ☐ TCP sets its RTO to an average RTT measure + 4*mean deviation of RTT, based upon Chebyshev's theorem
- ☐ Retransmission ambiguity would not occur if timestamps were used on packets.
- ☐ Self-clocking of TCP can be a liability in asymmetric networks where the reverse path can artifically constrain the forward path.
- ☐ Self-clocking can also lead to burstiness if the reverse path is congested, and/or the receiver uses a delay-ack time to suppress ACKs.
- ☐ The end-to-end congestion control model is the only one that can guarantee avoidance of congestion collapse.
- ☐ The notions of efficiency and fairness define an equilibrium point to which congestion control algorithms attempt to converge.

TCP…

- A stable congestion control algorithm converges to its equilibrium point.
- TCP uses additive-increase and multiplicative-decrease to probe for the equilibrium point.
- In the (w,α) notion of fairness, $\alpha = 1$ leads to max-min fairness.
- In equilibrium, TCP attempts to conserve packets and operate at high utilization.
- TCP does not guarantee low queueing delays because it depends upon packet loss for congestion detection
- Fast retransmit refers to the procedure of using three duplicate acks to infer packet loss
- TCP Tahoe sets its window to 1 after every loss detection
- TCP Reno may timeout quickly in a multiple packet loss scenario
- TCP SACK uses selective retransmit, and like NewReno, it does not reduce its window more than once per window of packets
- With a 28kbps reverse link & 1500 byte packets & regular TCP behavior, the forward link throughput is at most around 2 Mbps
- Header compression and link level ack suppression/regeneration could help in asymmetric bandwidth scenarios