TCP Congestion Control Review

cwnd – size of the sender's congestion window
threshold – used to determine if slow start or congestion avoidance algorithm
Flight Size – the amount of data that has been sent but not yet ACK'd

Original Version

- Start in Slow Start phase (cwnd = 1 MSS) and stay in this phase while cwnd < threshold
  - For each packet acknowledged, cwnd =cwnd + 1 MSS
- When cwnd > threshold, enter Congestion Avoidance phase
  - For each RTT, cwnd = cwnd + 1 MSS
  - An RTT is an ACK but it may acknowledge multiple packets (slow start would increase by 1 MSS for each packet acknowledged)
- Receipt of a duplicate ACK, set
  - threshold = max{Flight Size / 2, 2 * MSS}
- Timeout, go to Slow Start phase
  - cwnd = 1 MSS
  - threshold = max{Flight Size / 2, 2 * MSS}

Tahoe Version

- Add Fast Retransmit
  - Upon receipt of third duplicate ACK, retransmit the first missing segment
- Timeout OR a fast retransmit, go to Slow Start phase

Reno Version

- Add Fast Retransmit (same as Tahoe)
- Add Fast Recovery
  - Receipt of the third duplicate ACK, do NOT go to Slow Start phase, instead
    - cwnd = threshold + 3 MSS
      - 3 duplicate ACKs, so assume 3 missing segments, so should get 3 ACKS when retransmitted segment is received and ACK'd
  - Receipt of each additional duplicate ACK
    - cwnd = cwnd + 1 MSS
  - Record the highest outstanding unacknowledged packet sequence number (in variable *recover)*
  - When receive new ACK
    - cwnd = threshold
    - Return to congestion avoidance phase

New Reno Version

- Add Fast Retransmit and Fast Recovery
- Handle Fast Recovery differently
    - *Partial ACKs:* An ACK that acknowledges some but not all the segments that were outstanding at the start of fast recovery. NewReno interprets this as an indication of multiple loss.
    - If partial ACK received (dup ACK)
        - Re-transmit the next lost segment *immediately*
        - cwnd = cwnd + 1MSS
    - If new ACK received (ACK's sequence number in variable *recover*)
        - cwnd = min{threshold, FlightSize + 1MSS} [FlightSize is amout of outstanding data transmitted]
        - exit Fast Recovery

TCP Variation	**Response to 3 dupACK's**	**Response to Partial ACK of Fast Retransmission**	**Response to "full" ACK of Fast Retransmission**
Tahoe	Do fast retransmit, enter slow start	cwnd = cwnd + 1MSS	cwnd = cwnd + 1MSS
Reno	Do fast retransmit, enter fast recovery	Exit fast recovery, deflate window (cwnd = threshold), enter congestion avoidance	Exit fast recovery, deflate window (cwnd = threshold), enter congestion avoidance
NewReno	Do fast retransmit, enter modified fast recovery	Fast retransmit and deflate window (cwnd = threshold), remain in modified fast recovery	Exit modified fast recovery, deflate window (cwnd = threshold), enter congestion avoidance