CSC411: Advanced Networks

TCP Congestion Control

Note: **This class lecture will be recorded!**
If you do not consent to this recording, please do not ask questions via your video, audio or public chat; send your question to the instructor using the private chat.

Dr. Lisa Frye, Instructor
frye@kutztown.edu
Kutztown University

Timeout and Retransmission

- Varying transmission time
- Adaptive retransmission algorithm
- Sample round trip time

Round-Trip Time

- SampleRTT – time from segment sent until ACK received
- Why will the SampleRTT change from segment to segment on same TCP connection?

Average RTT

- averageRTT = (α * Old_ averageRTT) + ((1 – α) * SampleRTT)
- Value for α
- Why do you think the averageRTT is weighted the way it is?

Timeout value

- What is a good general (larger than, smaller than, etc) value for the timeout value?

- Timeout = β * averageRTT

Calculation of RTT values

- Why does the way TCP really works make the calculation of a sample round trip time non-trivial?

- Acknowledgement ambiguity
- Solution?

Karn's Algorithm

- Only deal with unambiguous ACKs
- Timer backoff strategy
  - Timeout event – increase timeout value
  - new_timeout = γ * timeout

- Variation in delay
  - Average RTT and variance

TCP Congestion Control

- End-end congestion control
- Network-assisted congestion control

- Congestion collapse

Congestion Control Variables

- CongWin
- Threshold

- Allowed window = min{CongWin, RcvWin}

TCP Congestion Control Algorithm

- Three main components
  - Additive-increase, multiplicative-decrease
  - Slow start
  - Reaction to timeout events

Slow start phase

- Start CongWin = 1 MSS
- Receive ACK, CongWin = CongWin + 1
- Increases exponentially
- Congestion avoidance
  - Only increase CongWin if all segments ACK'd

Multiplicative Decrease Congestion Avoidance

- Lost segment
  - CongWin reduced by half
  - Minimum value is 1 MSS

- Additive Increase, Multiplicative Decrease (AIMD)

Reaction to Timeout Events

- Enter slow start phase
- Grow exponentially until ½ value before timeout
- Threshold value
  - Initially very large (65KB)
  - Lost segment: Threshold = ½ * CongWin

- Why have different ways to handle congestion control?