University of Massachusetts Amherst

ECE697AA – Lecture 11

Congestion Control: TCP-Friendly UDP and RED

Tilman Wolf
Department of Electrical and Computer Engineering
10/09/08

TCP and UDP

- What happens if TCP and UPD flows share a link?

TCP and UDP

- What happens if TCP and UPD flows share a link?
- How does UDP react to congestion?
- How can UDP detect congestion?
- UDP is typically non-reactive
  - No congestion detection
  - No congestion control
  - No fairness with TCP
- How should UDP react to congestion?
  - “TCP-friendly”
  - Same throughput as if it was a TCP connection

Throughput of TCP

- What is the throughput of a TCP connection?
  - Congestion avoidance mode
    - » Window varies from W/2 to W
    - » One packet is lost during each cycle (why only one?)

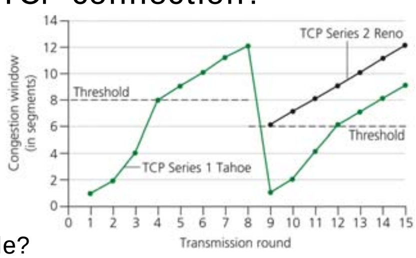

  - Loss rate
    - » How many packet sent per cycle?

$$\frac{W}{2}+\left(\frac{W}{2}+1\right)+\cdots+W=\sum_{i=W/2}^{W} i$$

$$\sum_{i=W/2}^{W} i=\frac{W}{2}\left(\frac{W}{2}+1\right)+\sum_{i=0}^{W/2} i$$

$$\frac{W}{2}\left(\frac{W}{2}+1\right)+\sum_{i=0}^{W/2} i=\frac{W^2}{4}+\frac{W}{2}+\frac{W/2(W/2+1)}{2}=\frac{3}{8}W^2+\frac{3}{4}W$$

Throughput of TCP

- Loss rate (cont.)
  - » One packet per cycle

$$L = \frac{1}{\frac{3}{8}W^2 + \frac{3}{4}W}$$

- Window size
  - » Solve for W
  - » For large W, $\frac{3}{8}W^2 >> \frac{3}{4}W$:

$$L \approx \frac{3}{8}W^2 \qquad \text{and} \qquad W \approx \sqrt{\frac{8}{3L}}$$

- Throughput
  - » incl. MSS because we consider packets

$$T = \frac{3}{4}\sqrt{\frac{8}{3L}}\frac{MSS}{RTT} = 1.22\frac{MSS}{RTT\sqrt{L}}$$

TCP-friendly UDP

- TCP friendly UDP requirements
  - Consider packet loss
  - Perform congestion control
  - Use algorithm that generates throughput with

$$T \leq 1.22\frac{MSS}{RTT\sqrt{L}}$$

  - If L=0, then maximum rate
- Note
  - Requires maintenance of RTT and MSS
  - RTT estimates should be updates once per RTT

TCP in wireless networks

- What happens to TCP in wireless networks?
- Wireless networks
  - Packet loss due to corruption on link
- TCP behavior
  - Packet loss is considered sign of congestion
  - More chances for timeout
  - Throughput

$$T = 1.22 \frac{MSS}{RTT\sqrt{L}}$$

» 100x increase in loss causes 10x reduction in throughput

TCP oscillations

- What happens if many TCP streams have same RTT?
  - Synchronization possible
  - Example:

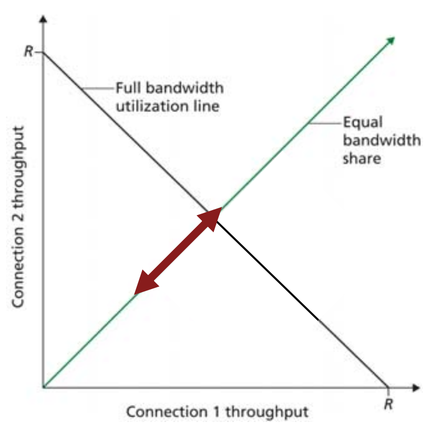

Congestion handling on routers

- Design goal of router
  - “Work conserving” in the broadest sense
  - Idle output is bad in time of congestion
  - Large queues can avoid “underflow”
- How much buffer space is needed?
  - Depends on RTT and link bandwidth
  - Rule of thumb: RTT*BW
  - What happens if buffer is larger?
- Better idea: avoid oscillation
  - Is there a way of preventing congestion?

Random early detection

- Start throttling sources before queue is full
  - Throttle sources by dropping packet
- RED
  - If queue above threshold, drop random packet from queue
  - TCP flows will slow down before buffer space is exhausted
  - Adapt drop probability with increasing queue length
- RED is based on local view
  - No end-system support required
  - Incremental deployment possible
- “Active queue management”

Random early detection

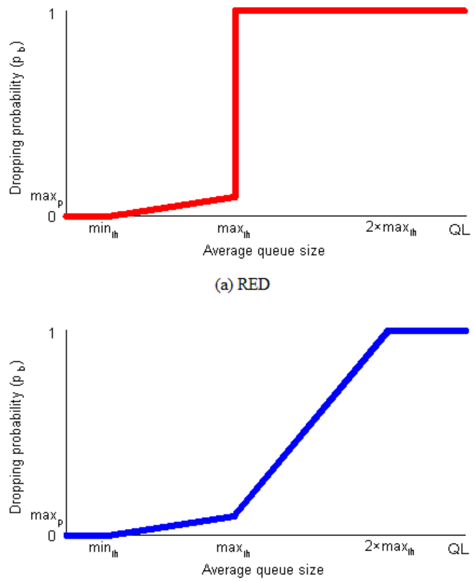

(a) RED

(b) RED with the "gentle option"

Assignments

- Read
  - Kurose & Ross: Chapter 8
- SPARK
  - Assessment quiz