

# The Congestion Game

Jeff Chase  
Duke University



# Review

- What are the limits to throughput for TCP?

# Limits to Throughput

- Host limitations/overhead
- Flow limits: how fast can app can produce or consume data?
- Wire speed
- Congestion limits available bandwidth.
- Efficiency
  - Leaving network idle
  - Sending duplicate data
  - High ratio of control to data
  - What causes loss of efficiency?

# TCP Efficiency

- Retransmission timer
  - Underestimate RTT causes unnecessary retrans
  - Overestimate RTT leaves bandwidth idle
- Window sizes
  - Too small? Leaves bandwidth idle.
- Congestion window estimation
  - Slow start
  - Sawtooth
  - Must be conservative...but not too conservative
- Not conservative enough? **Congestion collapse.**

# Sending Too Fast

- What happens if the sender goes too fast?

# Sending Too Fast

- Overflow at receiver? Receiver drops packet.
- Overflow network link? NIC drops packet.
- Overflow router? Router drops packet.
- Faster than fair share?
  - Pro: *you win*
  - Con: *somebody else loses*
  - *TCP is a game*

# Max-Min Fairness Criteria

- "Fair sharing"
- But flows have differing demands...
- Flows demanding less than their share get as much as they need.
- Flows demanding more than their share split the surplus.
- Generalizes to **proportional sharing**

# Trust and Rate Control

- Is gaming TCP a security problem?
- How should the network deal with this?
- Whose responsibility is it?
- What incentive does anyone have to play the game by the rules?
  - Good Samaritan?
  - Rodney King: "Can't we all just get along?"
  - Judge Judy?
  - Adam Smith?



# TCP Game (RFC 2581)

- Ack each segment with highest seqnum received.
- Acks drive actions at sender: **self-clocking**
- Below ssthresh, double window on each ack.
  - "Slow start"
- Above ssthresh, increment window on each ack.
  - Additive increase == "congestion avoidance"
- Loss? Congestion! Cut window in half.
  - Multiplicative decrease == "congestion control"
- Triple dup ack? Loss! "Fast retransmit". (Tahoe)
- After fast-retransmit loss, back to ssthresh, every ack adds to  $1/cwnd$  to window, even dups.
  - "Fast recovery" (Reno)

# Savage TCP (Daytona)

- Attack: "Ack early, ack often".
  - Three variations on a theme.
  - Ack early hides congestion loss.
  - "Big ack attack"
- Defense:
  - Don't make hidden assumptions.
    - One ack per segment? Uh uh.
  - Remove incentives to cheat.
  - Trust but verify.
    - Nonces and cumulative nonces.