

**CSC2209**  
**Computer Networks**

---

**TCP Vegas and TCP Nice**

**Stefan Saroiu**  
**Computer Science**  
**University of Toronto**

# TONS Presentation Today

---

- M. Ammar, from Georgia Tech in SF1105 from 2-3pm
  - Disruption tolerant networks
  - Class of mobile networks

# TCP Vegas

---

# Questions Today about TCP Vegas

---

- What Reno's problems is Vegas fixing?
- How is it fixing them?

# TCP Vegas

---

- Pro-active approach to congestion detection
- Small queues at routers

# TCP Vegas

---

- Early congestion detection using RTTs
- Additive decrease on early congestion
- Restrict # of packets maintained per flow at bottleneck router
- Small queues, yet good utilization

# TCP Vegas Techniques

---

# TCP Vegas Techniques

---

- More accurate RTT measurement/ timeouts
- New retransmission
- Congestion avoidance mechanism
- Slow-start



# RTT Measurement

---

- Reno: coarse-grained timer
- Vegas: fined-grained timer
- Better for calculating timeouts, could be more precise in timeout expiration, faster reaction to loss

# Retransmissions

---

- Reno: when receiving 3 duplicate ACKs
- Vegas:
  - When receiving duplicate ACK, checks if timeout expired, and if so, retransmits
  - When receiving a non-duplicate ACK 1st/2nd after retransmission check if timeout expired and if so retransmits

# Congestion window decrease

---

- Reno: possible to decrease the congestion window more than once during one RTT
- Vegas: in case of multiple segment loss and more than one fast retransmission, the congestion window is reduced only for the first fast retransmission

# Congestion detection

---

- Reno: reactive: loss signals congestion
- Vegas: proactive: tries to detect incipient congestion by comparing the measured  $X_{put}$  to its notion of expected  $X_{put}$

# Algorithm

---

- $\text{baseRTT} = \text{minimum RTT}$
- $\text{windowSize} = \text{size of window in bytes}$
- $\text{Expected } X_{\text{put}} = \text{windowSize} / \text{baseRTT}$
  
- $\text{rttLen} = \# \text{ of bytes during last RTT}$
- $\text{RTT} = \text{average RTT of segments acknowledge during last RTT}$
- $\text{Actual } X_{\text{put}} = \text{rttLen} / \text{RTT}$
  
- $\text{Diff} = \text{expected} - \text{actual}$
- Two thresholds:  $\alpha, \beta$
- If  $\text{diff} < \alpha$ , increase linearly; if  $\text{diff} > \beta$ , decrease linearly
- Otherwise do nothing

# Slow-start modifications

---

- Reno: double congestion window size every RTT
- Vegas: double congestion window size every other RTT
  - In-between RTTs are used for computing actual  $X_{put}$
  - When actual rate falls below expected rate, exit slow-start, enter linear-increase / linear-decrease

# Problems

---

# Problems

---

- How would Reno and Vegas compete at the bottleneck router?
- Re-routing
- Old flows versus new flows



# TCP Nice

---

# TCP for Spare Network Capacity

---

# TCP for Spare Network Capacity

---

- Rate-limiting
- Off-peak hours
- TCP connection manager
  
- How do these compare with TCP Nice?

# High-Level Question

---

# High-Level Question

---

- New TCP idea comes along?
  - How are we going to evaluate it?
- Simulation?
- Emulation?
- Large-scale Internet measurement?