A Comparison of Equation-Based and AIMD Congestion Control

Sally Floyd, Mark Handley, and Jitendra Padhye

Sept. 4-6, 2000
Workshop on the Modeling of Congestion Control Algorithms
Paris

Why look at non-TCP congestion control mechanisms?

• The congestion control mechanisms in TCP are well-suited for applications such as file transfer:

  – with the goal of transfering a file in the shortest possible time;

  – given the restrictions of fairness and the avoidance of congestion collapse.

• Some applications would prefer to avoid TCP's characteristic halving of the congestion window:

  – and would be willing to pay the price of a longer transfer time.

• Not considered:

Applications exempt from end-to-end congestion control.

Candidates for TCP-compatible but slowly-responding congestion control:

- Equation-based congestion control
  - For example, TFRC (TCP-Friendly Rate Control).

- TCP based on AIMD (Additive-Increase, Multiplicative-Decrease)
  - With increase and decrease parameters a and b.

- Binomial congestion control algorithms that are extensions of AIMD.

- And a number of others...

Equation-based congestion control: TFRC

• Use the equation for TCP's acceptable sending rate, as a function of the loss rate and RTT.

• The receiver estimates the loss event rate over the most recent eight loss intervals.

• We believe that TFRC is a viable congestion control scheme for many streaming media applications that currently don't use end-to-end congestion control.

Experiments with TFRC

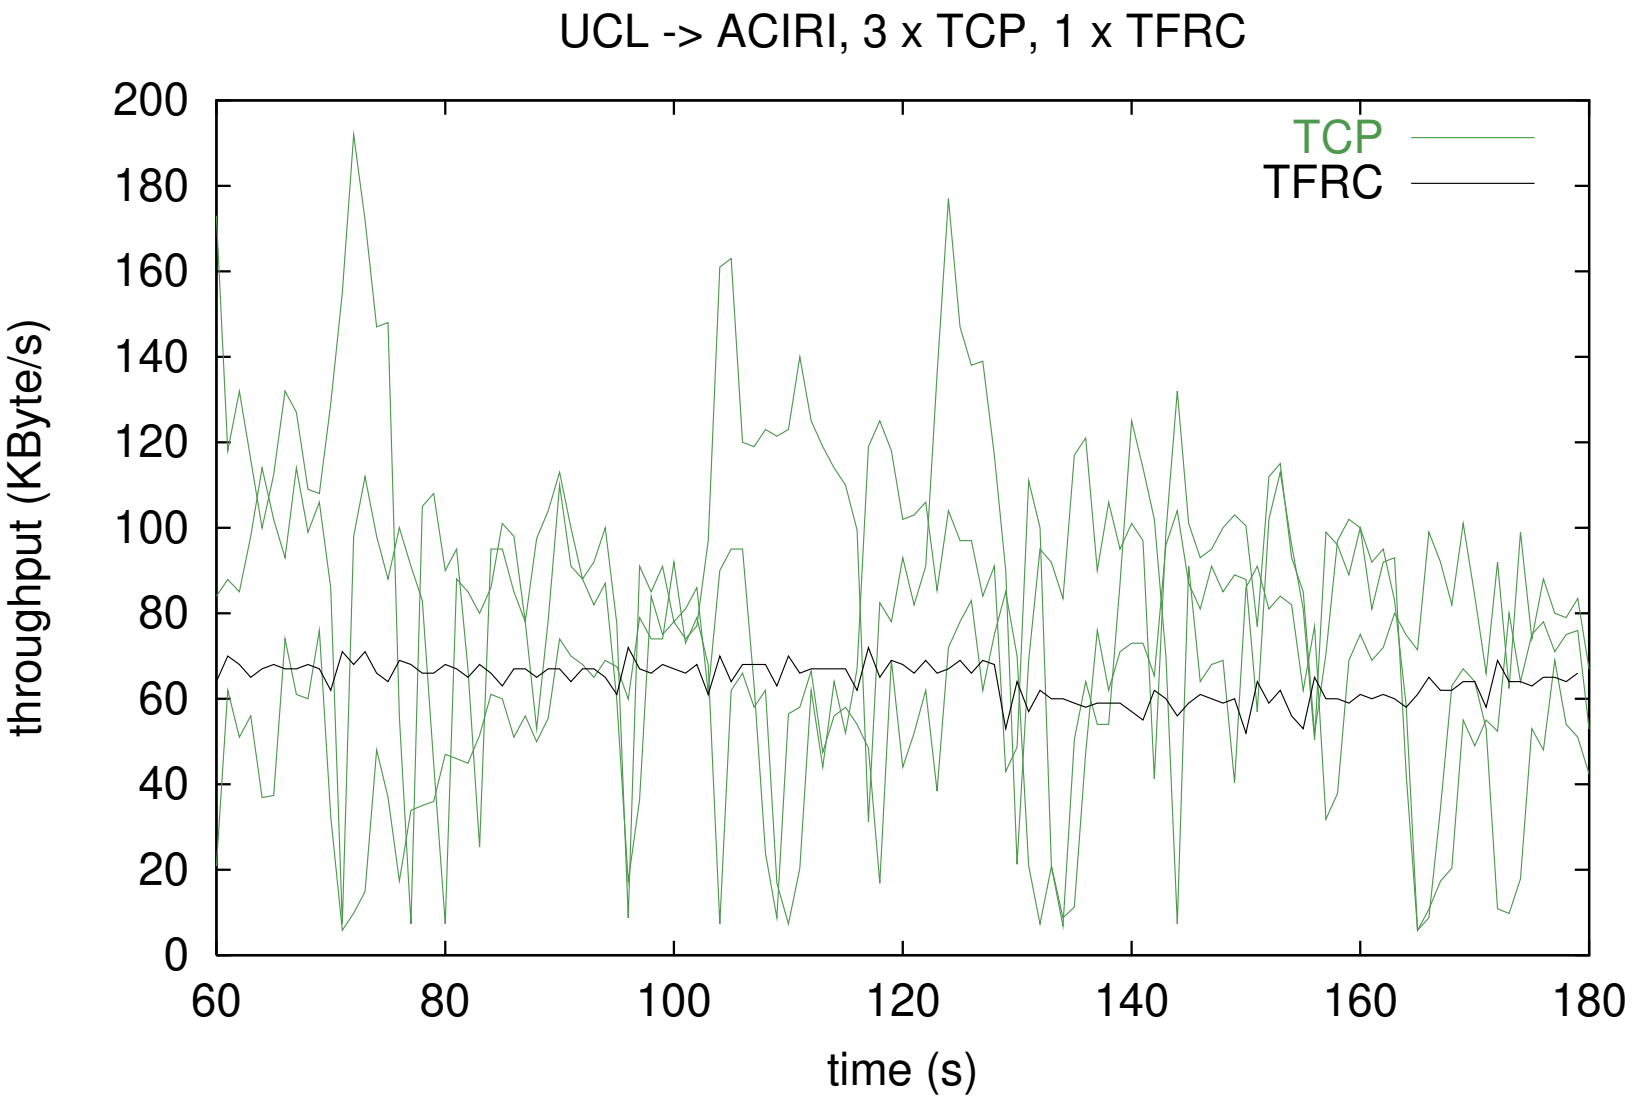

AIMD(a, b):

- The steady-state model:

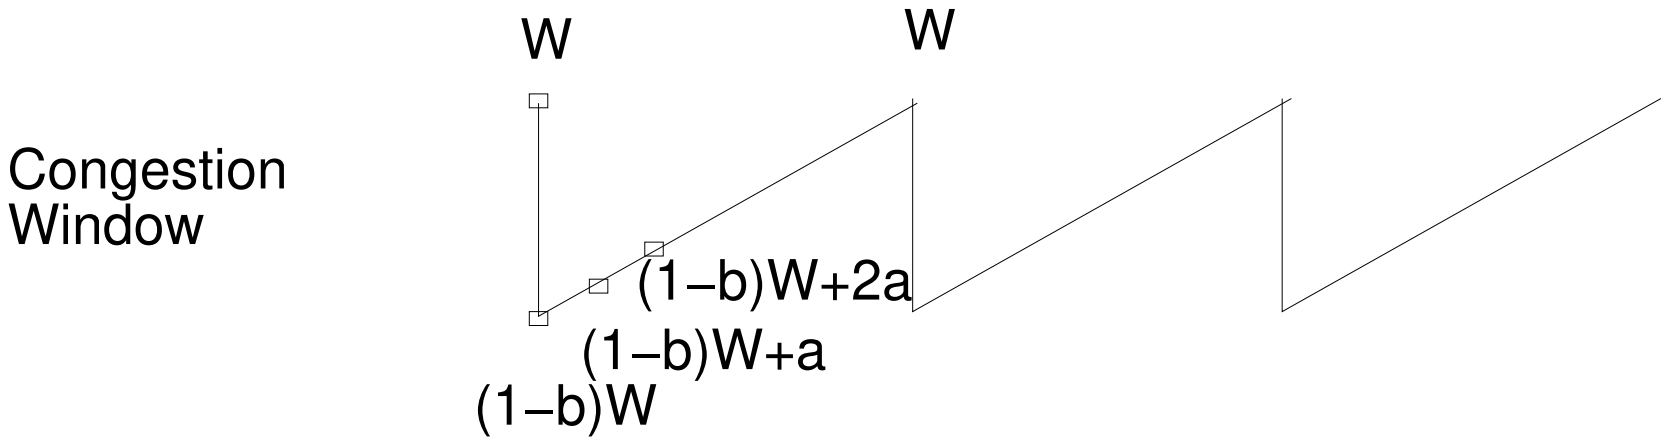

AIMD congestion window in steady-state.

- Decrease window from W to $(1-b)W$.

- Increase window from W to $W+a$ packets each roundtrip time.

AIMD(a, b), continued:

- The average sending rate S is $(1 - b/2)W$ packets per RTT.
- Each cycle has one drop in about $\frac{b(2-b)}{2a}W^2$ packets.
- Therefore, $W \approx \sqrt{\frac{2a}{b(2-b)p}}$, and $S \approx \sqrt{\frac{(2-b)a}{2bp}}$.
- For TCP, $S \approx \sqrt{1.5/p}$.
  - So AIMD(a, b) should compete fairly with AIMD(1, 1/2), in the steady-state model, if $a = \frac{3b}{(2-b)}$.

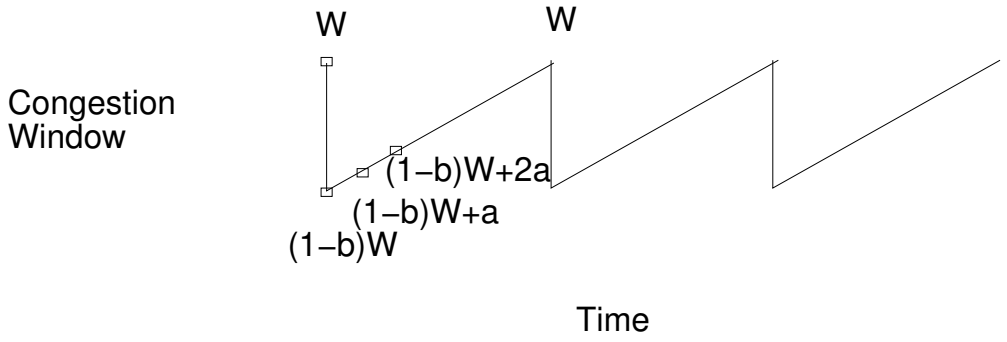

TCP(1/5, 1/8), or TCP(2/5, 1/8)?

• The previous slide suggests that TCP(1/5, 1/8) should be TCP-compatible, but actually, TCP(2/5, 1/8) looks better in simulations:

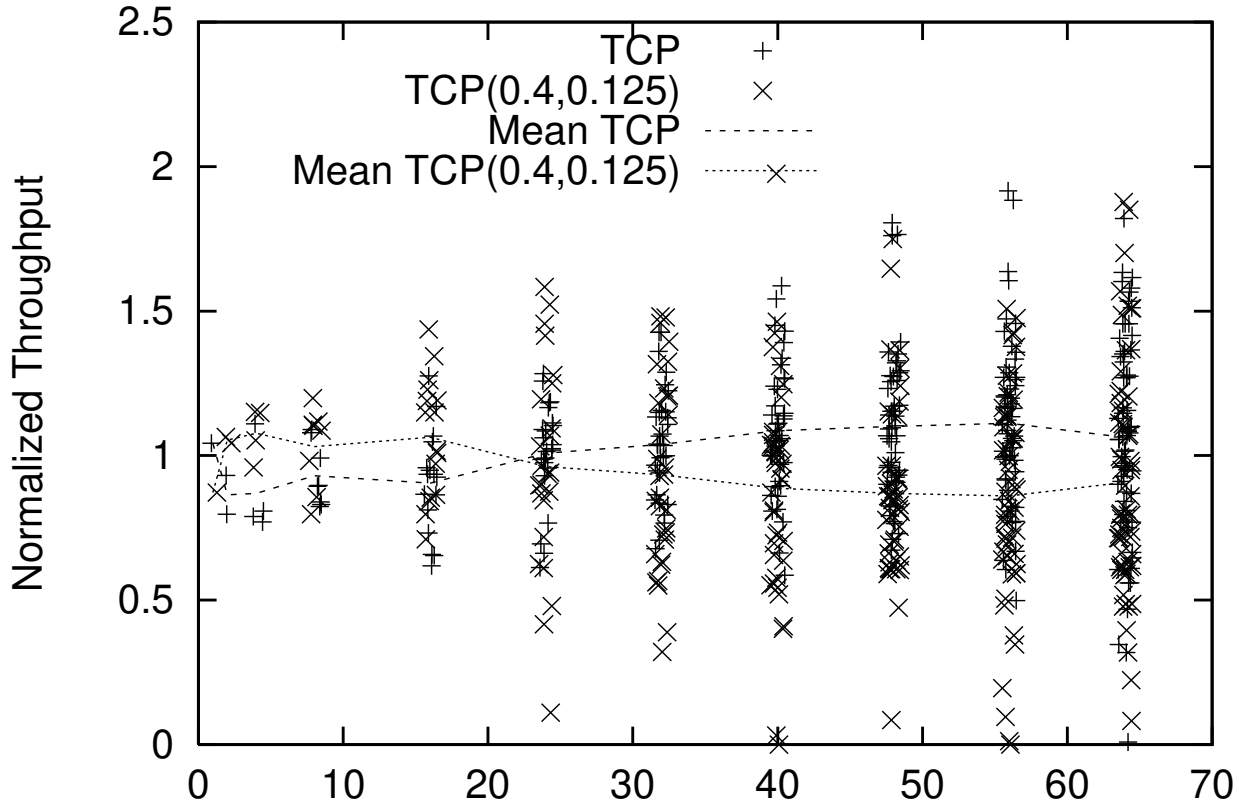

Comparing Aggressiveness and Responsiveness:

- Aggressiveness:

  Max increase in sending rate (in packets/RTT) in one RTT.

  – AIMD(a, b) or TCP(a, b): a
  – TFRC: 0.14$-$0.28 packets

- Responsiveness:

  RTTs of sustained congestion for halving the sending rate.

  – AIMD(a, b): $\log_{1-b} 0.5$
  – TCP(a, b): in heavy congestion, ack-clocking and retransmit timeouts give quick reductions
  – TFRC: 5 RTTs

Comparing Aggressiveness and Responsiveness, cont.:

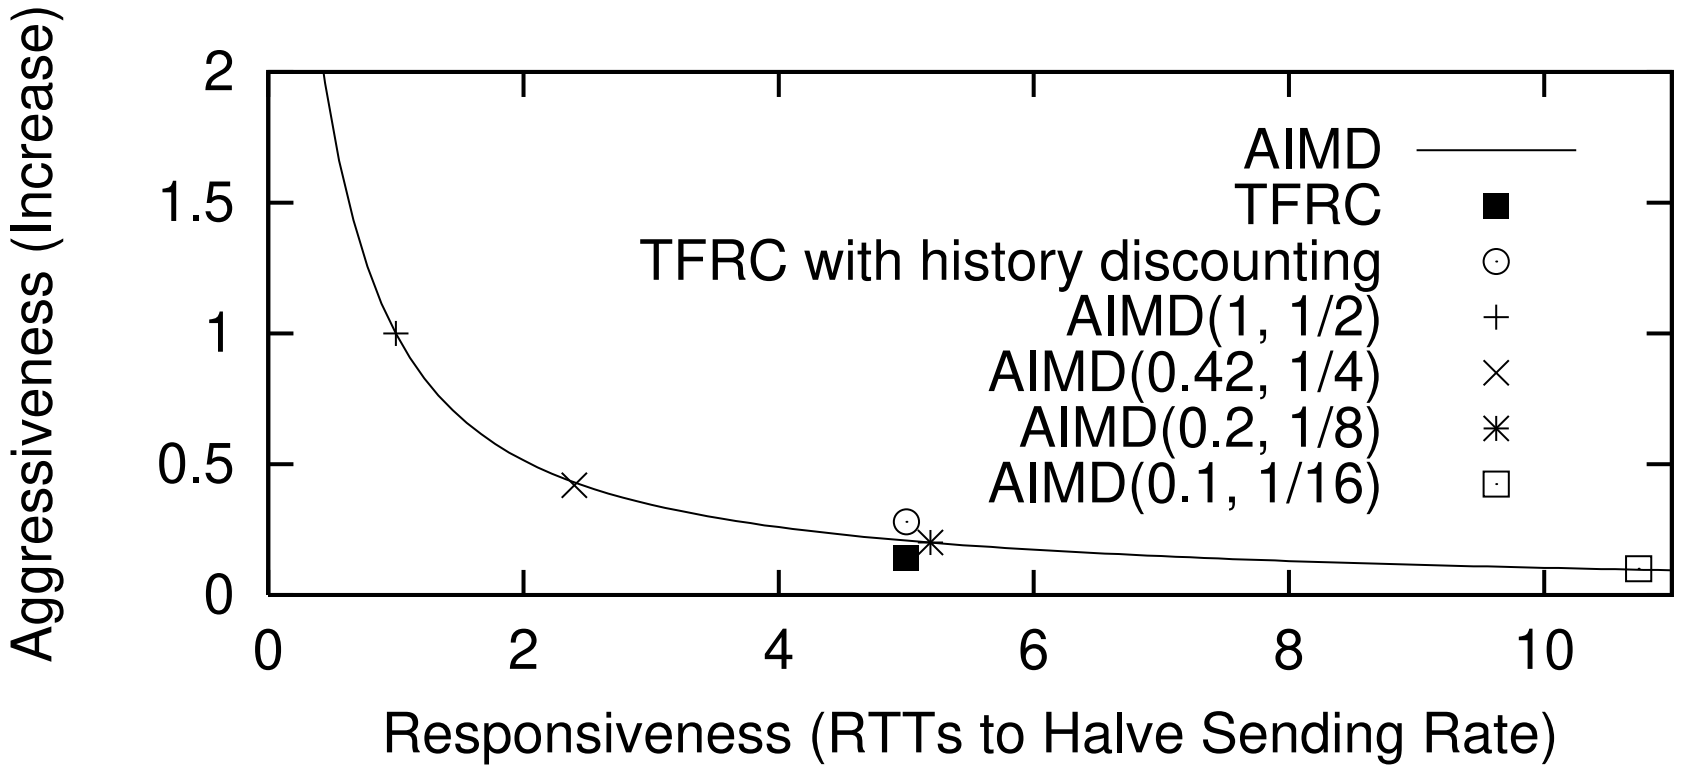

Comparing Smoothness and Responsiveness:

- Smoothness:

Largest reduction of sending rate in one RTT, in the steady-state model.

  – TCP(a, b): $1 - b$
  – TFRC: 1

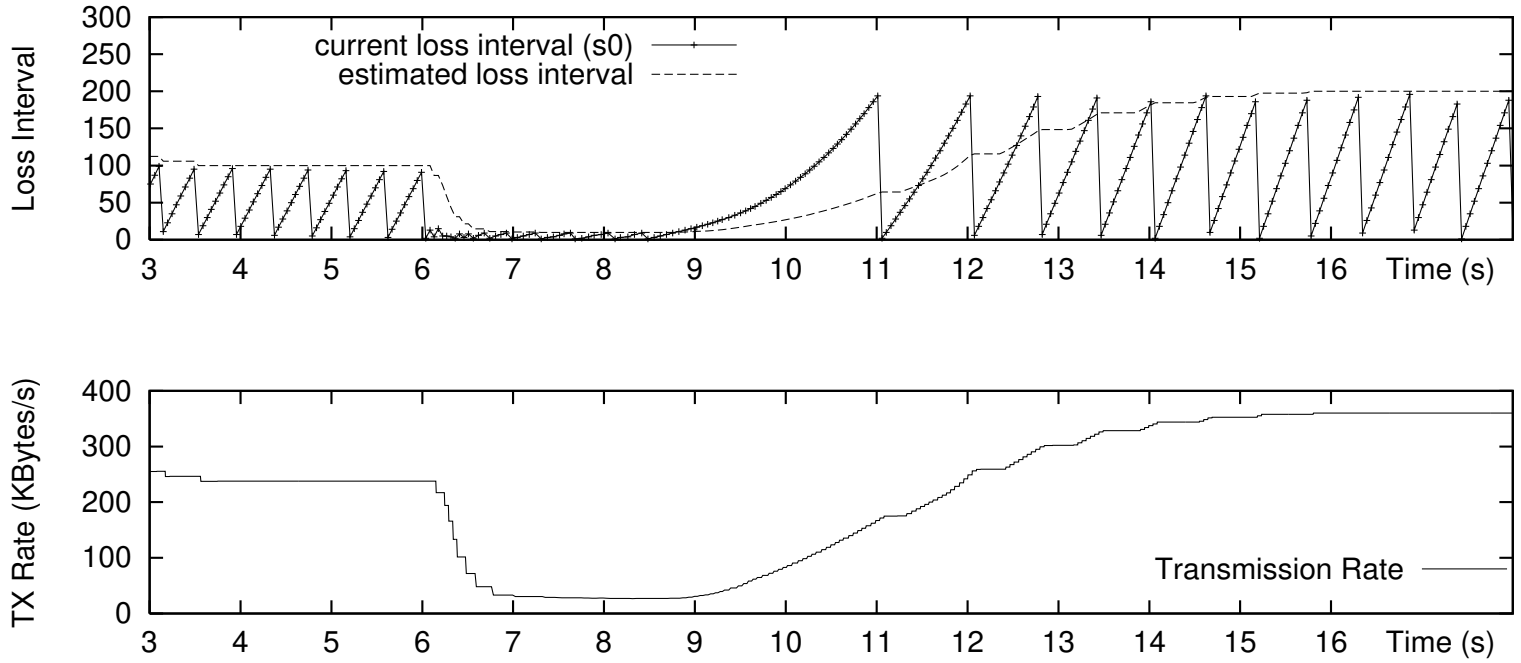

Comparing Smoothness and Responsiveness, etc.:

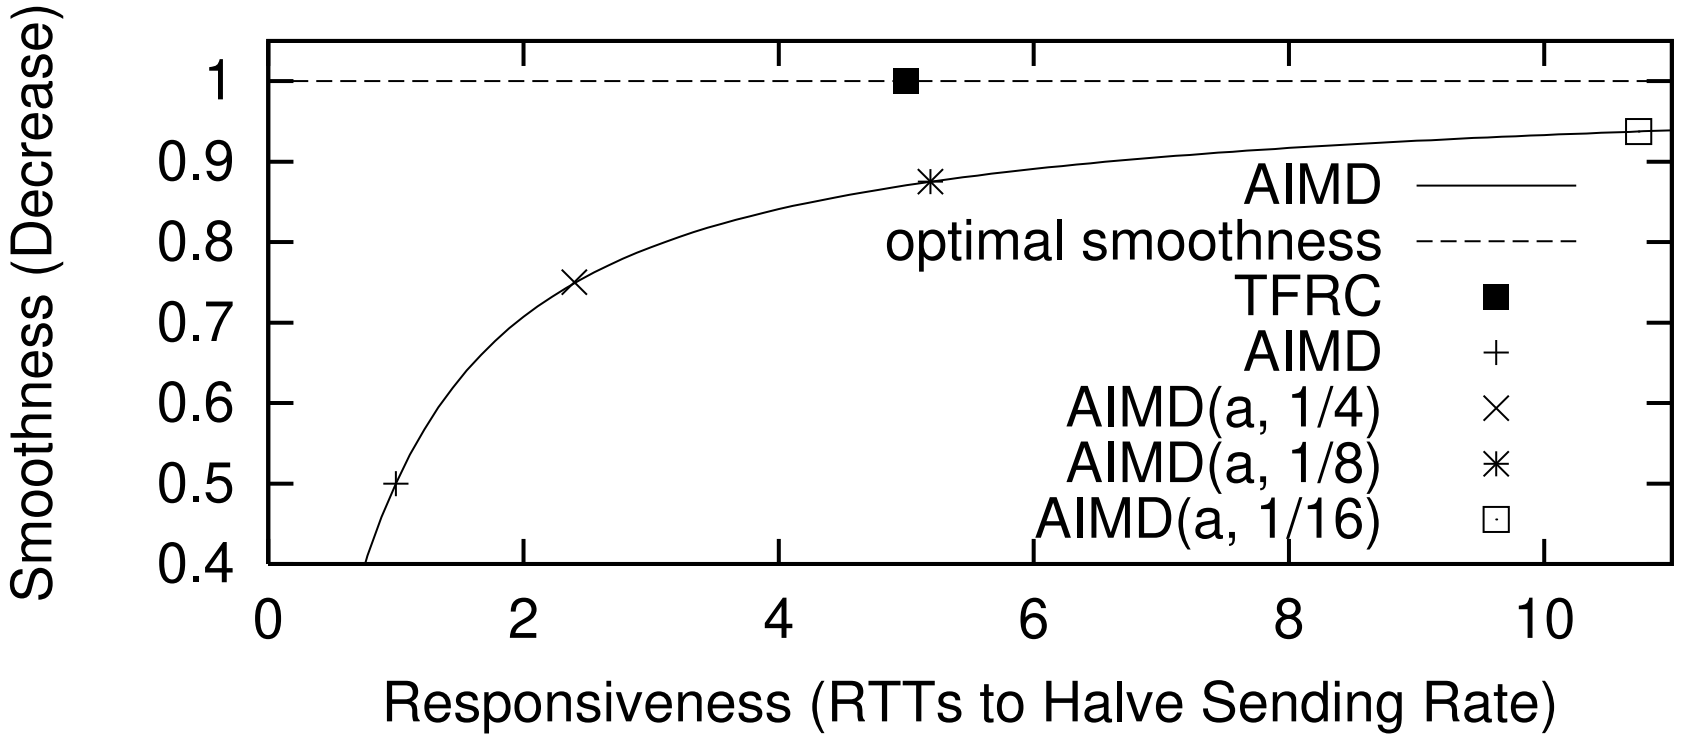

Simulation with TCP[2/5, 1/8] (top) and TCP (bottom)

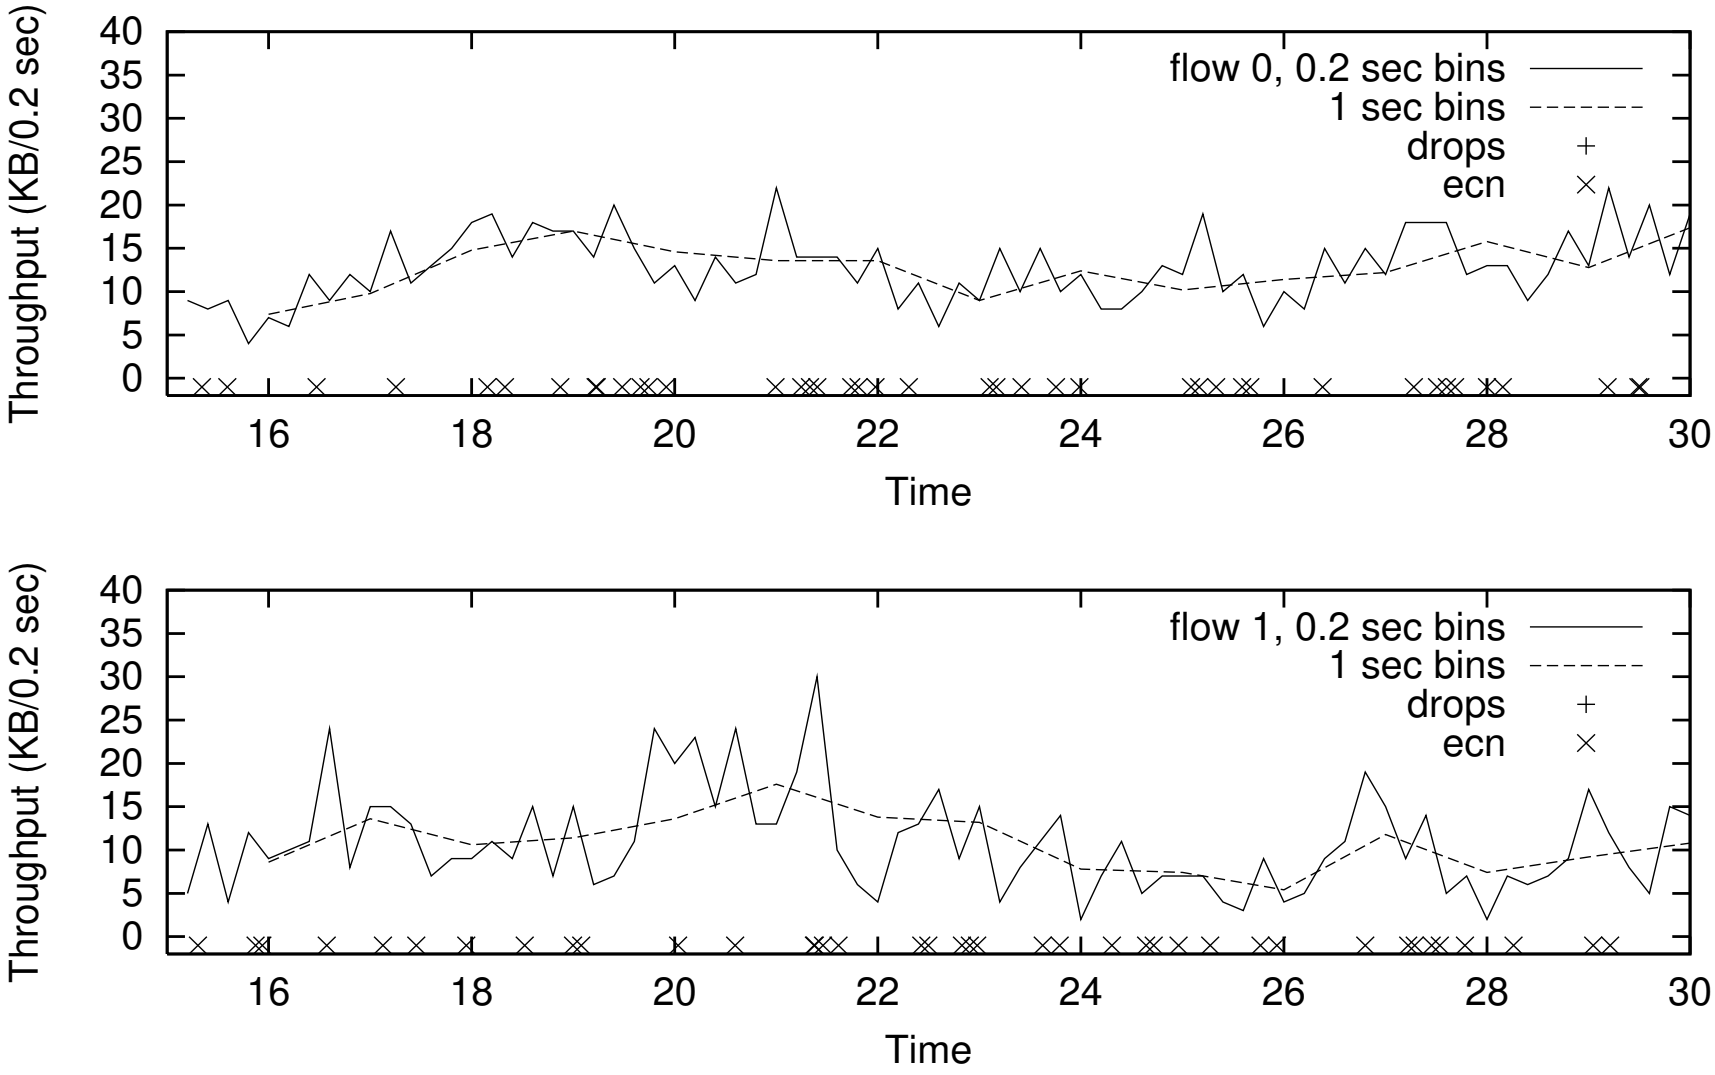

Simulation with TFRC (top) and TCP (bottom)

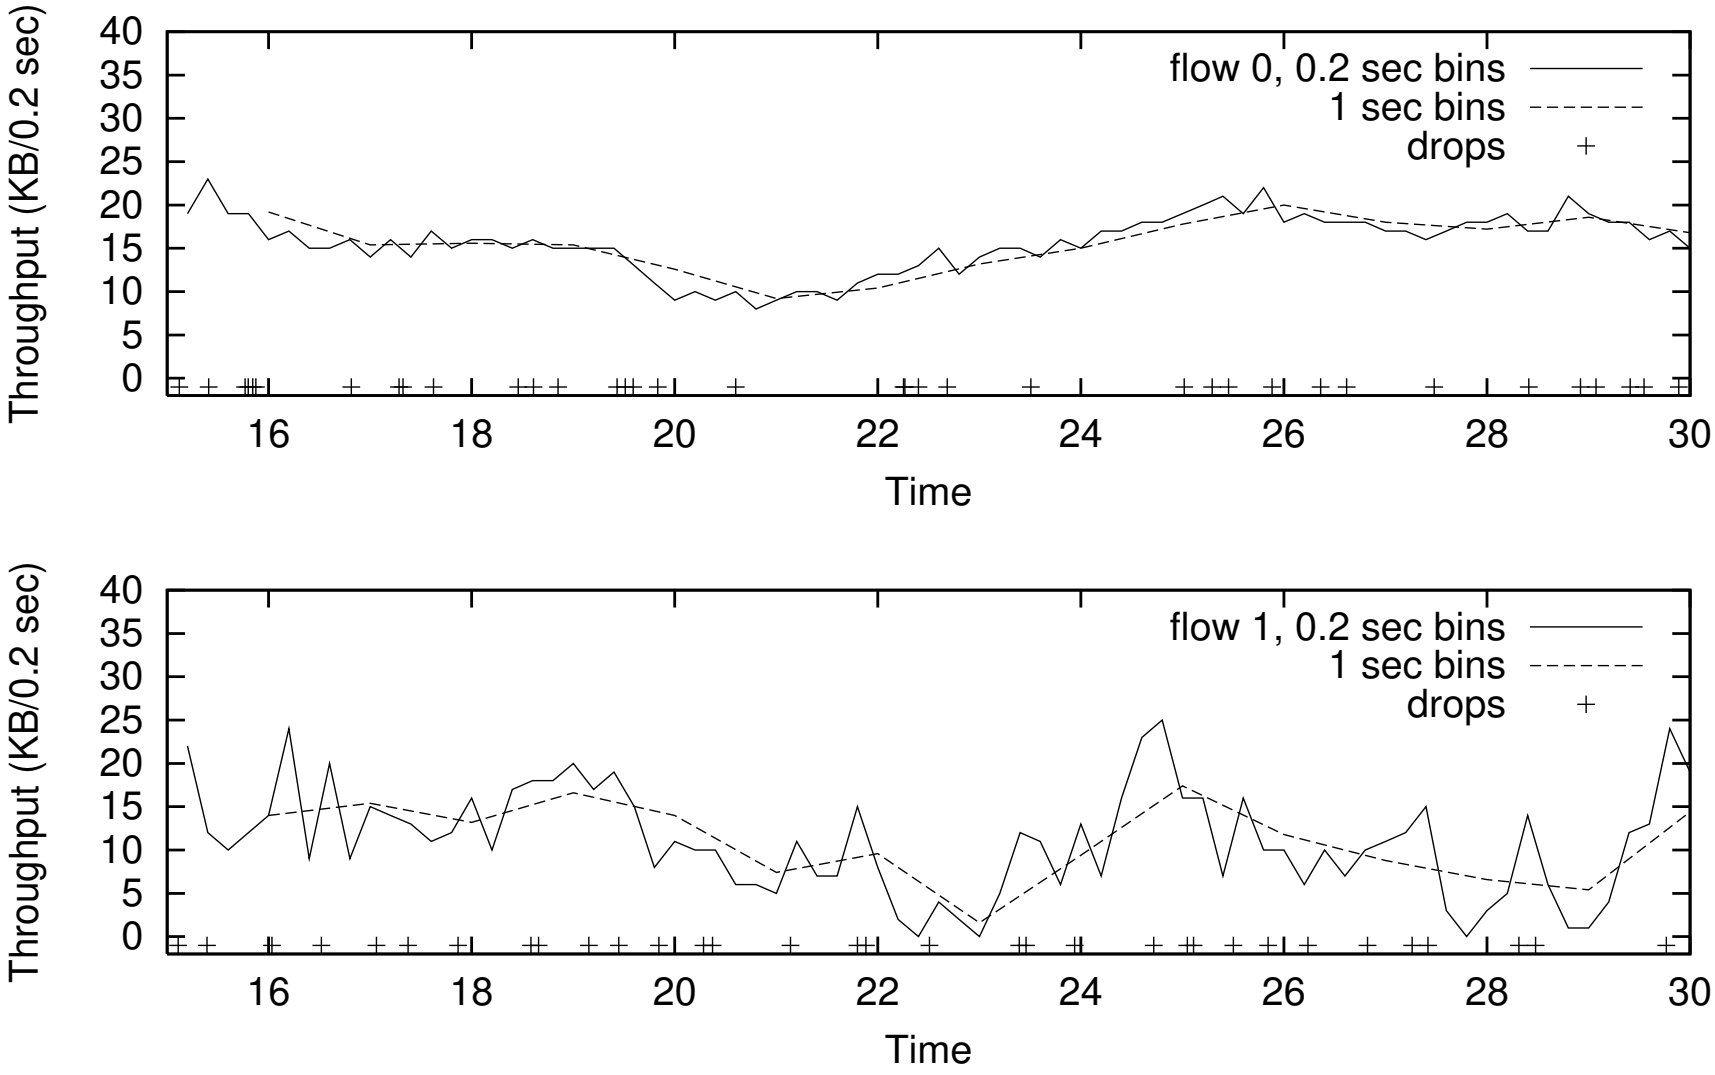

Measuring the throughput ratio

- T_i: the sending rate for a flow over the i-th time interval.

- Throughput ratio for the i-th interval: $\frac{T_i}{T_{i-1}}$

- For a flow, we look at the distribution of throughput ratios, for 0.2-second, 1-second, and 10-second time intervals.

The cumulative distribution of throughput ratios

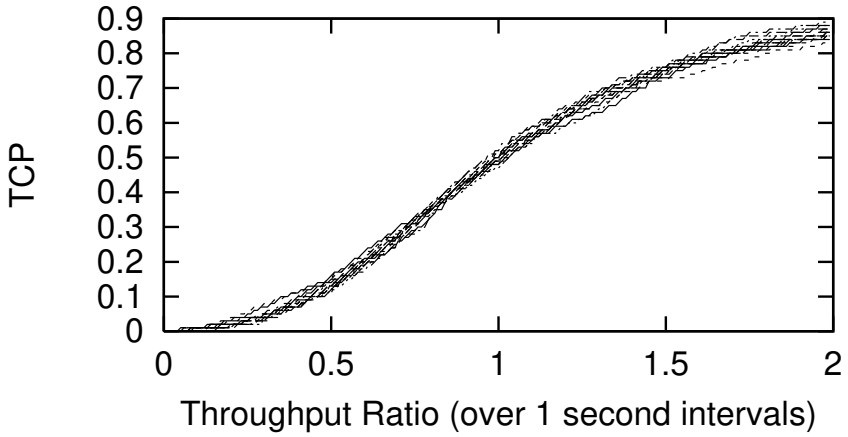

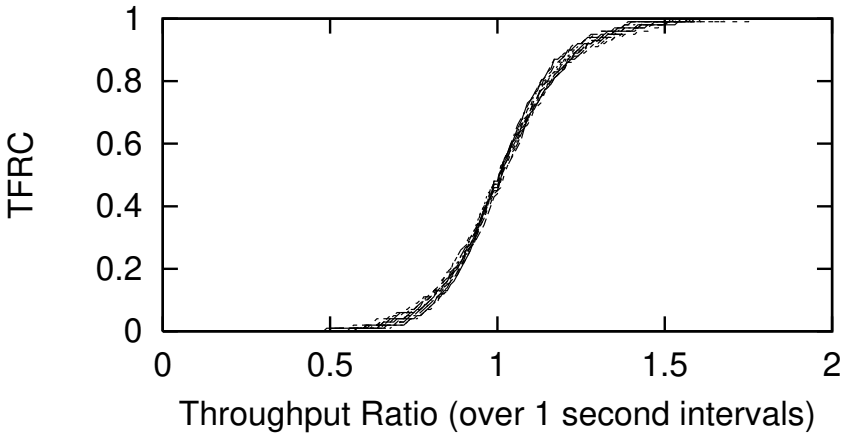

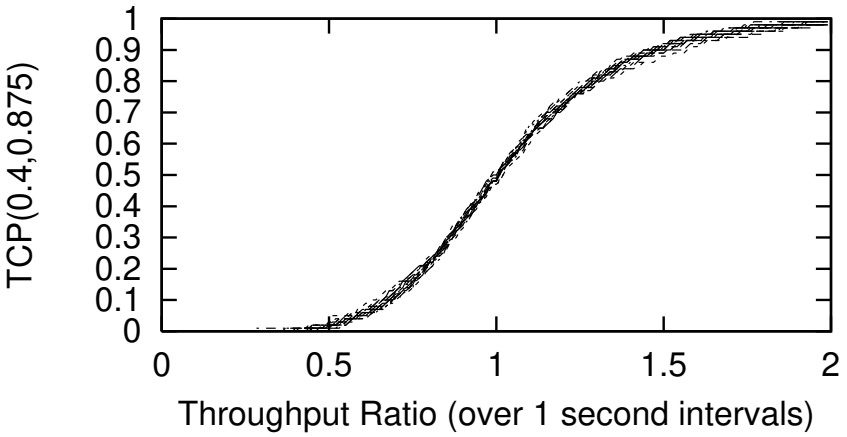

Summary:
Why are we bothering with TFRC, instead of using TCP(2/5, 1/8)?

- TFRC gives a smoother sending rate than TCP(2/5, 1/8).

- TFRC uses less traffic on the return path than TCP:
  - TFRC uses a feedback message once per RTT.
  - TCP uses ACK packets.

- TFRC is a promising building block for multicast congestion control.
  - Receiver-based.
  - Sender slowly adjusts its sending rate based on feedback from receivers.

Differences between memoryless and memory-based congestion control

- Memory-based congestion control:
  – Equation-based congestion control, such as TFRC, with its memory of past loss intervals.

- Memoryless congestion control:
  – TCP, which responds to the presence or absence of congestion in the most recent RTT.
  – TCP's slow-start threshold ssthresh retains some minimal memory during a slow-start.

Simulations of TFRC (top) and TCP(2/5, 1/8) (bottom), with packet bursts perfectly adapted to TFRC's memory

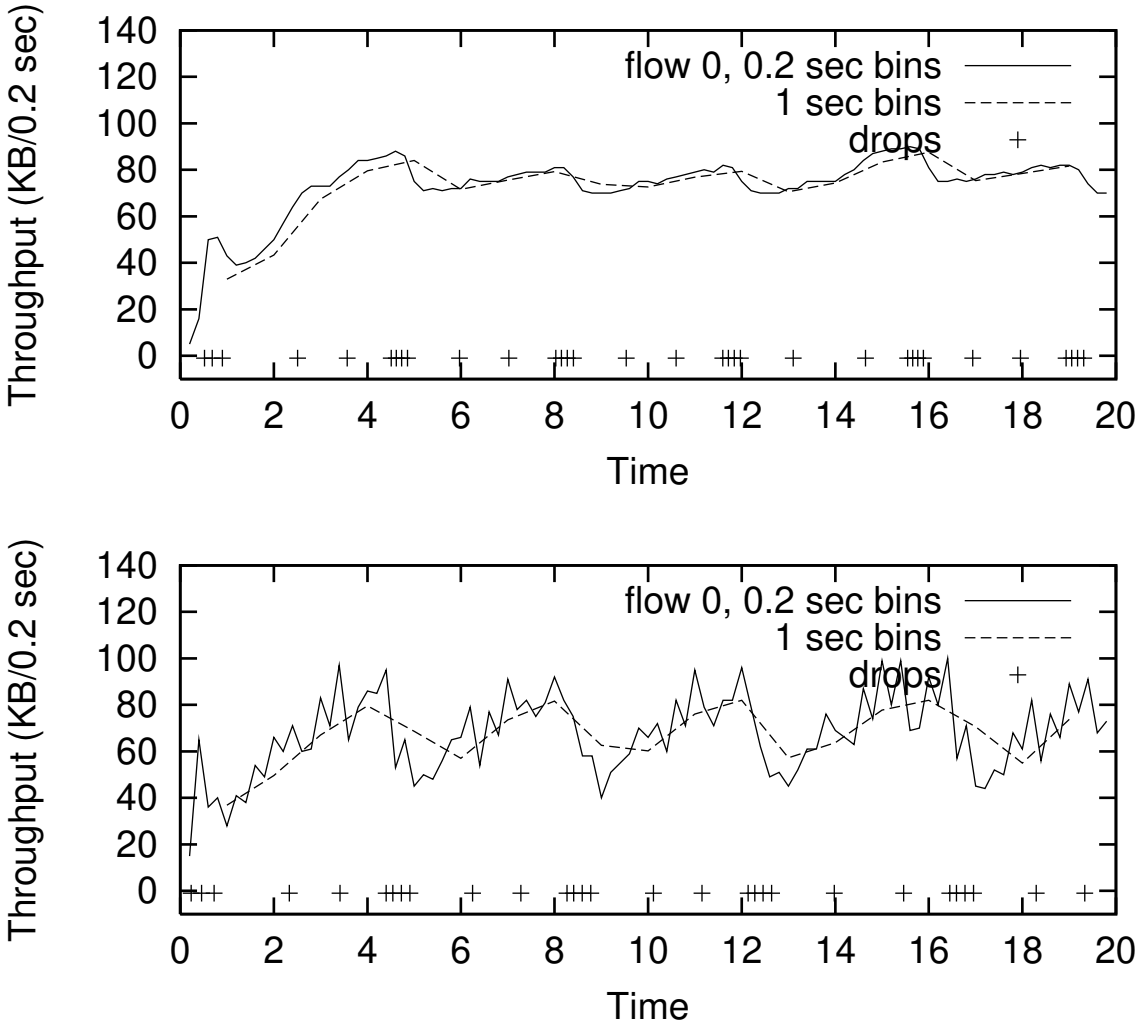

Simulations of TFRC (top) and TCP(2/5, 1/8) (bottom), with packet bursts *not* perfectly adapted to TFRC's memory

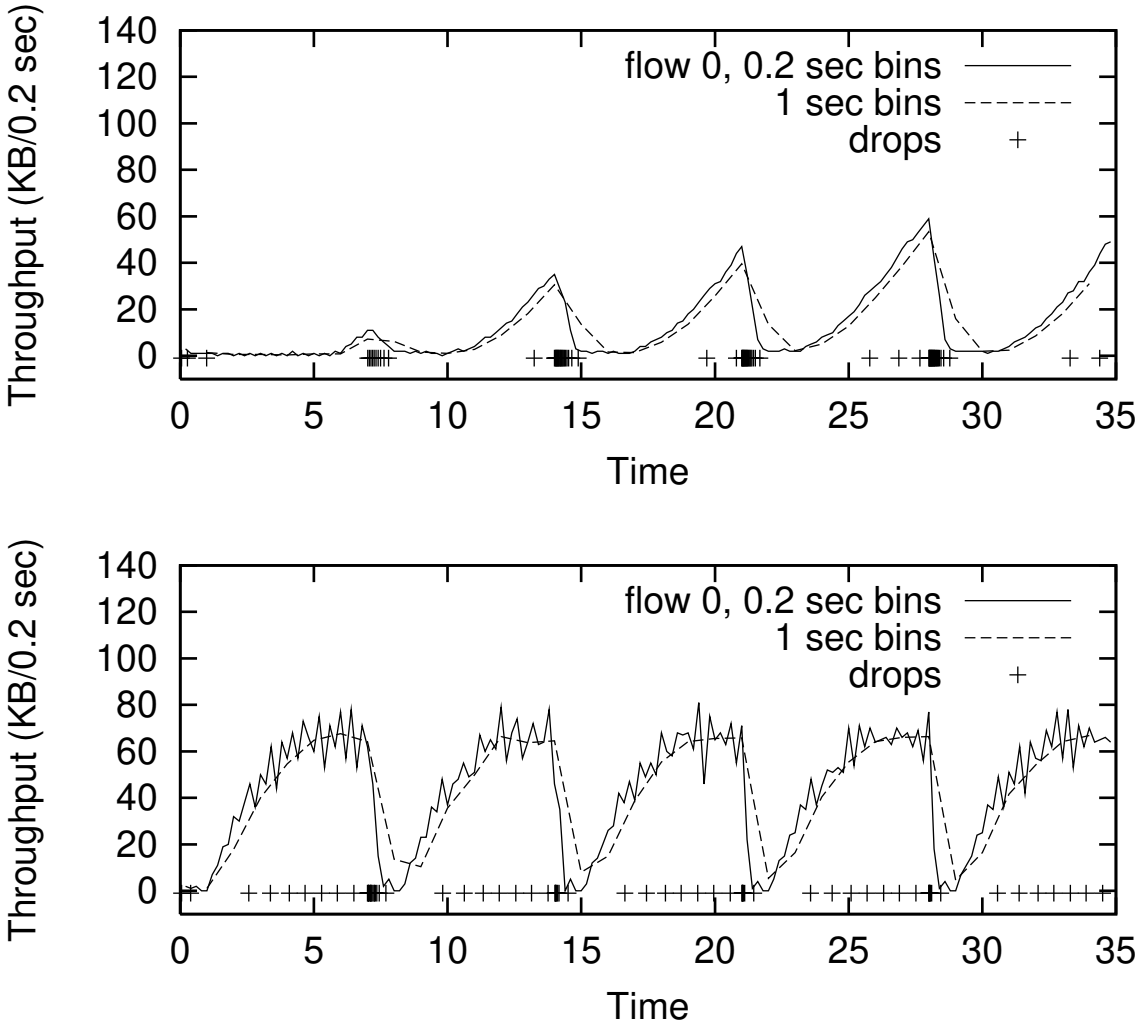

Concerns about Slow Congestion Control (SlowCC):

Joint work, in progress, with Deepak Bansal, Hari Balakrishnan, and Scott Shenker.

Concerns about Slow Congestion Control (SlowCC):

- (1) Less aggressive: SlowCC's limited ability to take advantage of an increase in the available bandwidth.

- (2) Less responsive: An extended period of transient congestion, with high packet drop rates, after a sharp decrease in the available bandwidth.

- (3) Fairness with TCP in a changing environment:
  - Long-term fairness with TCP.
  - Transient fairness with TCP.
  - Fairness to competing TCP web mice.

Aggressiveness: SlowCC's limited ability to take advantage of increases in the available bandwidth:

- Assume that the bandwidth available to a flow is doubled.

- For a TCP flow with increase parameter a, the sending rate is increased by:
  - a/R packets/sec each RTT,
  - a/R^2 packets/sec each second.

  - For a TCP with a sending rate of λ packets/sec, it takes $\lambda R^2/a$ seconds to double its sending rate, in congestion avoidance.

- This "cost" of SlowCC is independent of the scheduling or queue management at the routers.

Responsiveness: SlowCC's slow decrease in response to congestion:

• In mild but persistent congestion (e.g., with ECN), TCP(a, b) decreases its sending rate by a multiplicative factor $(1 - b)$ each RTT.

• In severe congestion, TCP(a, b)'s decrease is quick, dominated by ack-clocking and retransmit timeouts.

• For equation-based congestion control, the decrease rate is determined by the method for estimating the loss event rate at the receiver.

• For TFRC, which averages over a number of loss intervals, the decrease rate is determined by the number of loss intervals in the average.

– For heavy congestion, the "ack-clocking-emulation" is the dominant mechanism.

Summary of fairness results:

- Long-term fairness with TCP in a changing environment.
- Transient fairness with TCP in a changing environment..
- Fairness to competing TCP web mice in a changing environment.

Summary: We could find no fairness reasons for not deploying SlowCC (that is, TCP(1/8) or TFRC) in the current Internet.

SlowCC's long-term fairness with TCP in changing environments:

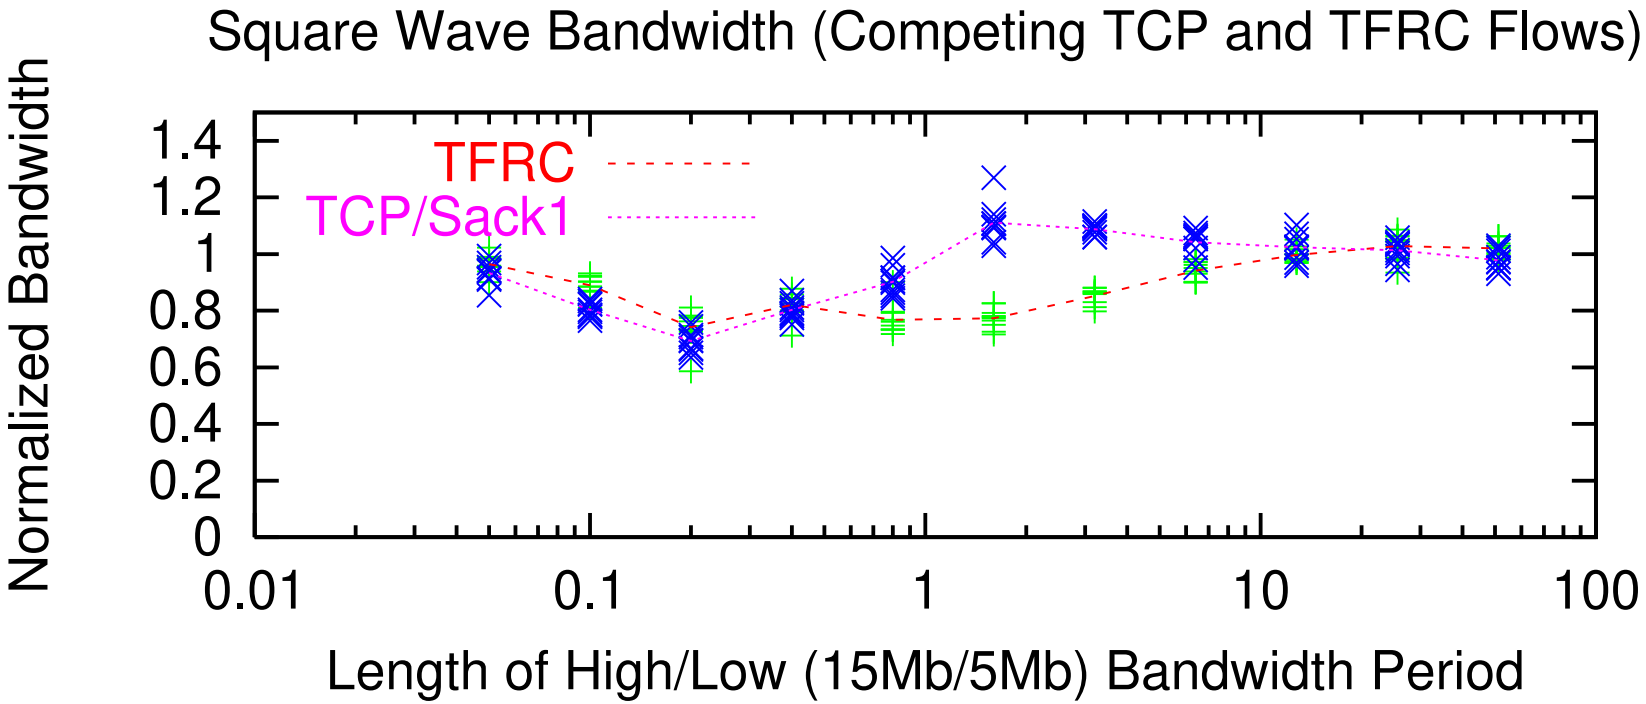

Simulations of TCP competing with TFRC.

SlowCC's long-term fairness with TCP in changing environments:

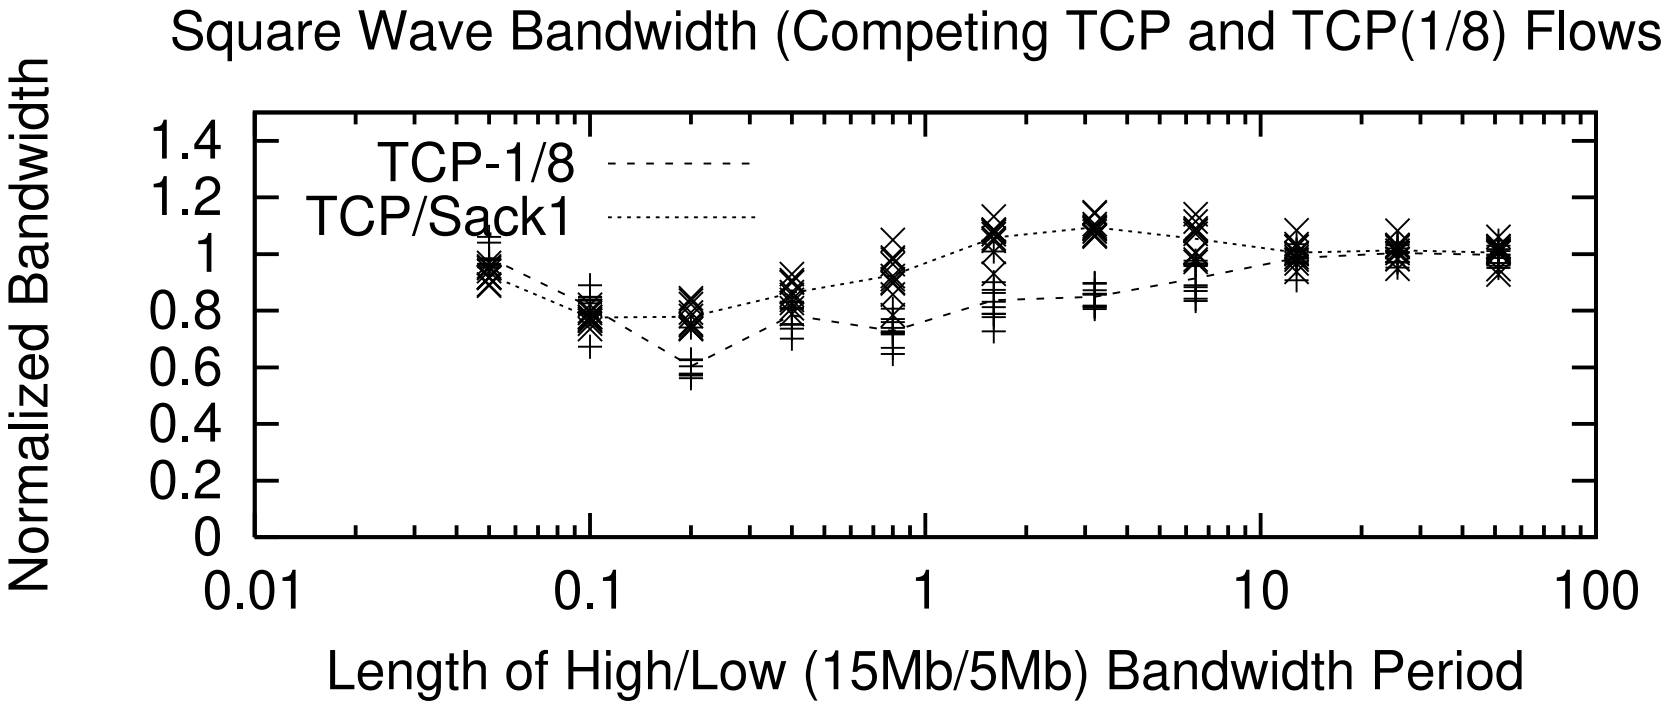

Simulations of TCP competing with TCP(1/8).