NADA Algorithm Update and Test Case Evaluations

draft-zhu-rmcat-nada

Xiaoqing Zhu, Rong Pan, Michael Ramalho, Sergio Mena,
Charles Ganzhorn, Paul Jones, and Stefano D'Aronco

IETF 91, Honolulu, Hawaii, USA

2014-11-09

Outline

- Revised algorithm
  - TCP co-existence
  - Accelerated ramp-up
- Simulation-based test case evaluation
- Testbed results

Issues with Currently Documented Algorithm (as in draft-zhu-rmcat-nada-04)

- Slow ramp-up at start up
- Ignored loss/marking information
- Lack of TCP coexistence

Revised Receiver Behavior

Non-linear warping

Queuing delay d_n

Warped delay $\tilde{d}_n$

Packet marking ratio p_M

Packet loss ratio p_L

Receiving rate r_R

Calculate **aggregate** congestion signal and its derivative:

Delay weight for marking

Delay weight for loss

$$x_n = \tilde{d}_n + p_M d_M + p_L d_L$$

$$x'_n = \frac{x_n - x_{n-k}}{\delta}$$

ACK interval

Non-linear Warping of Delay

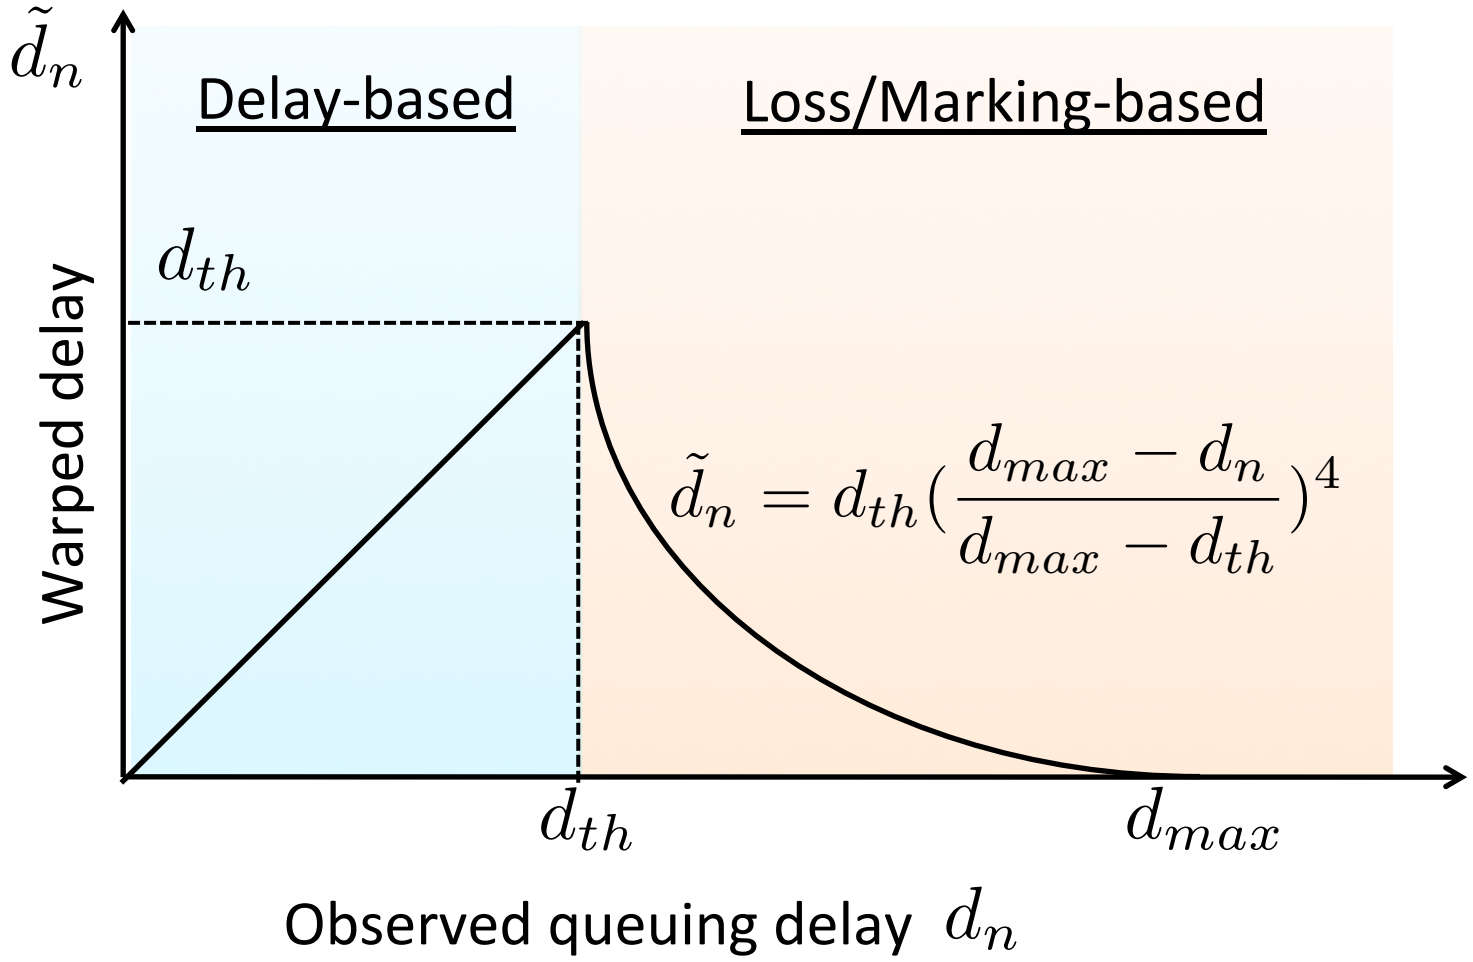

[Budzisz et al. ToN 2011]

Revised Sender Behavior

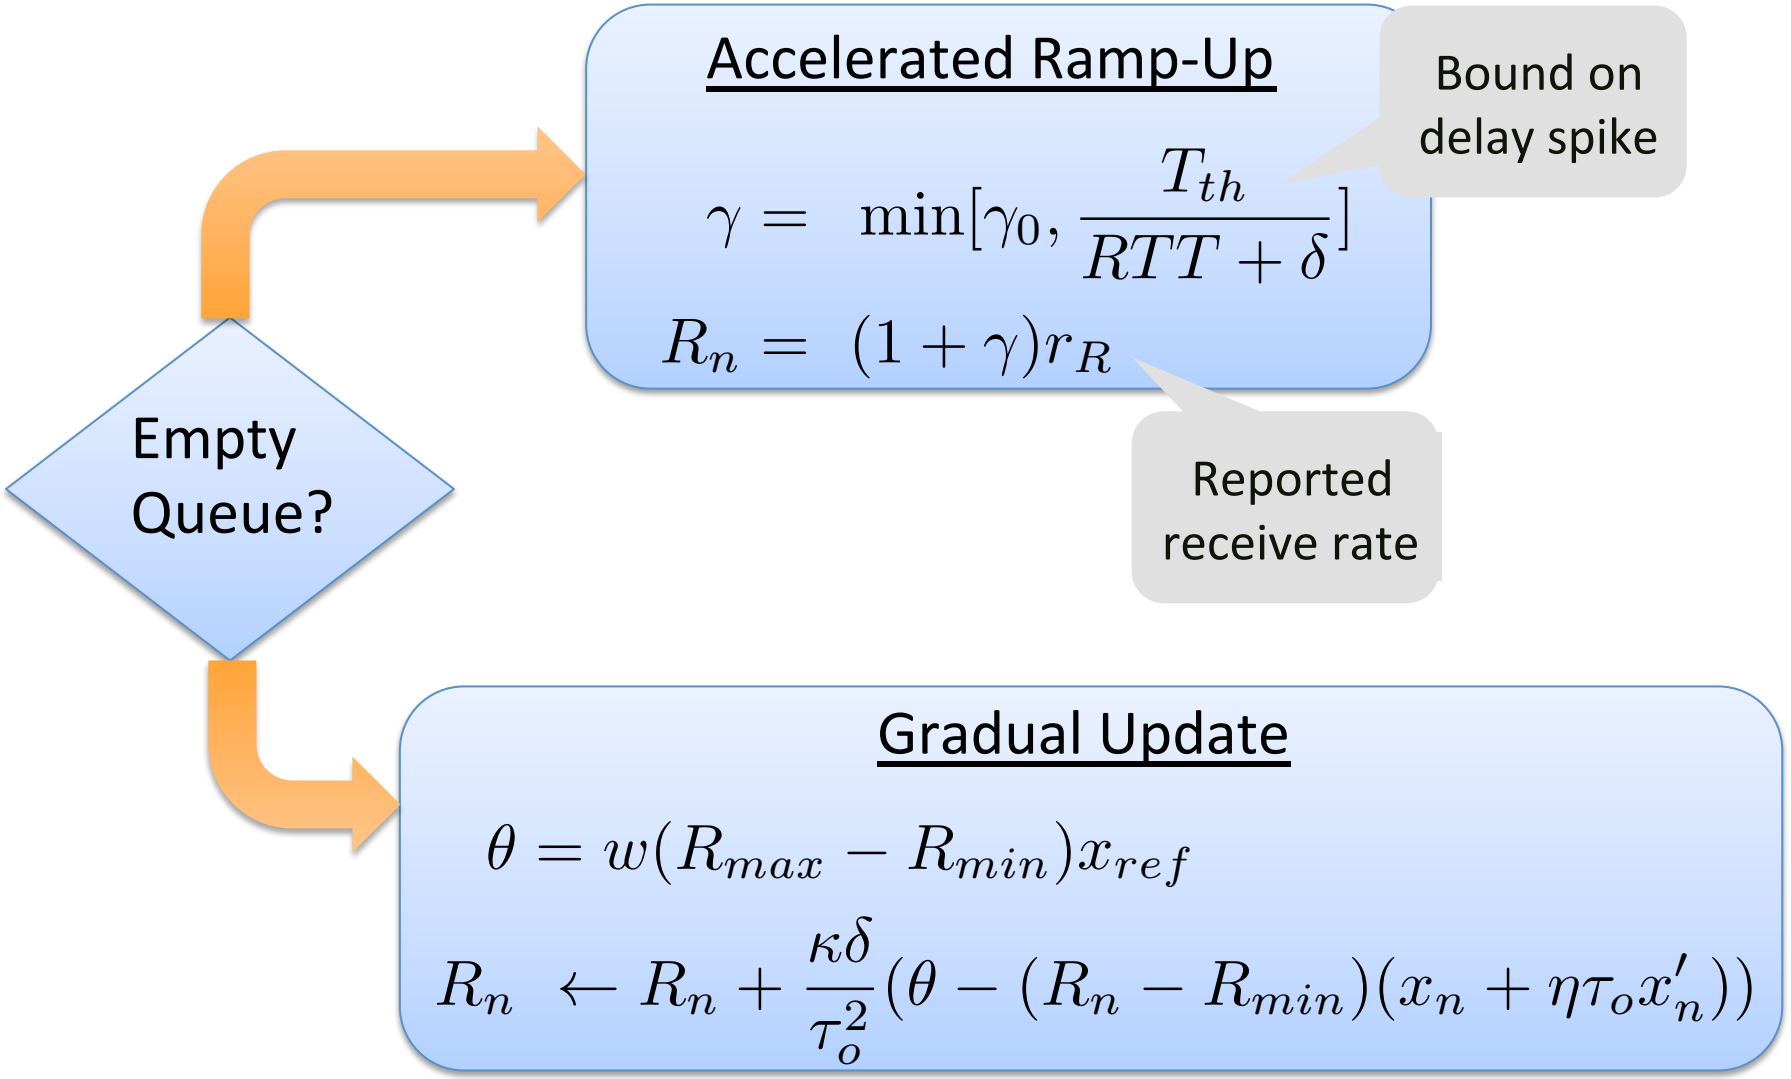

RTT-Aware Accelerated Ramp-Up

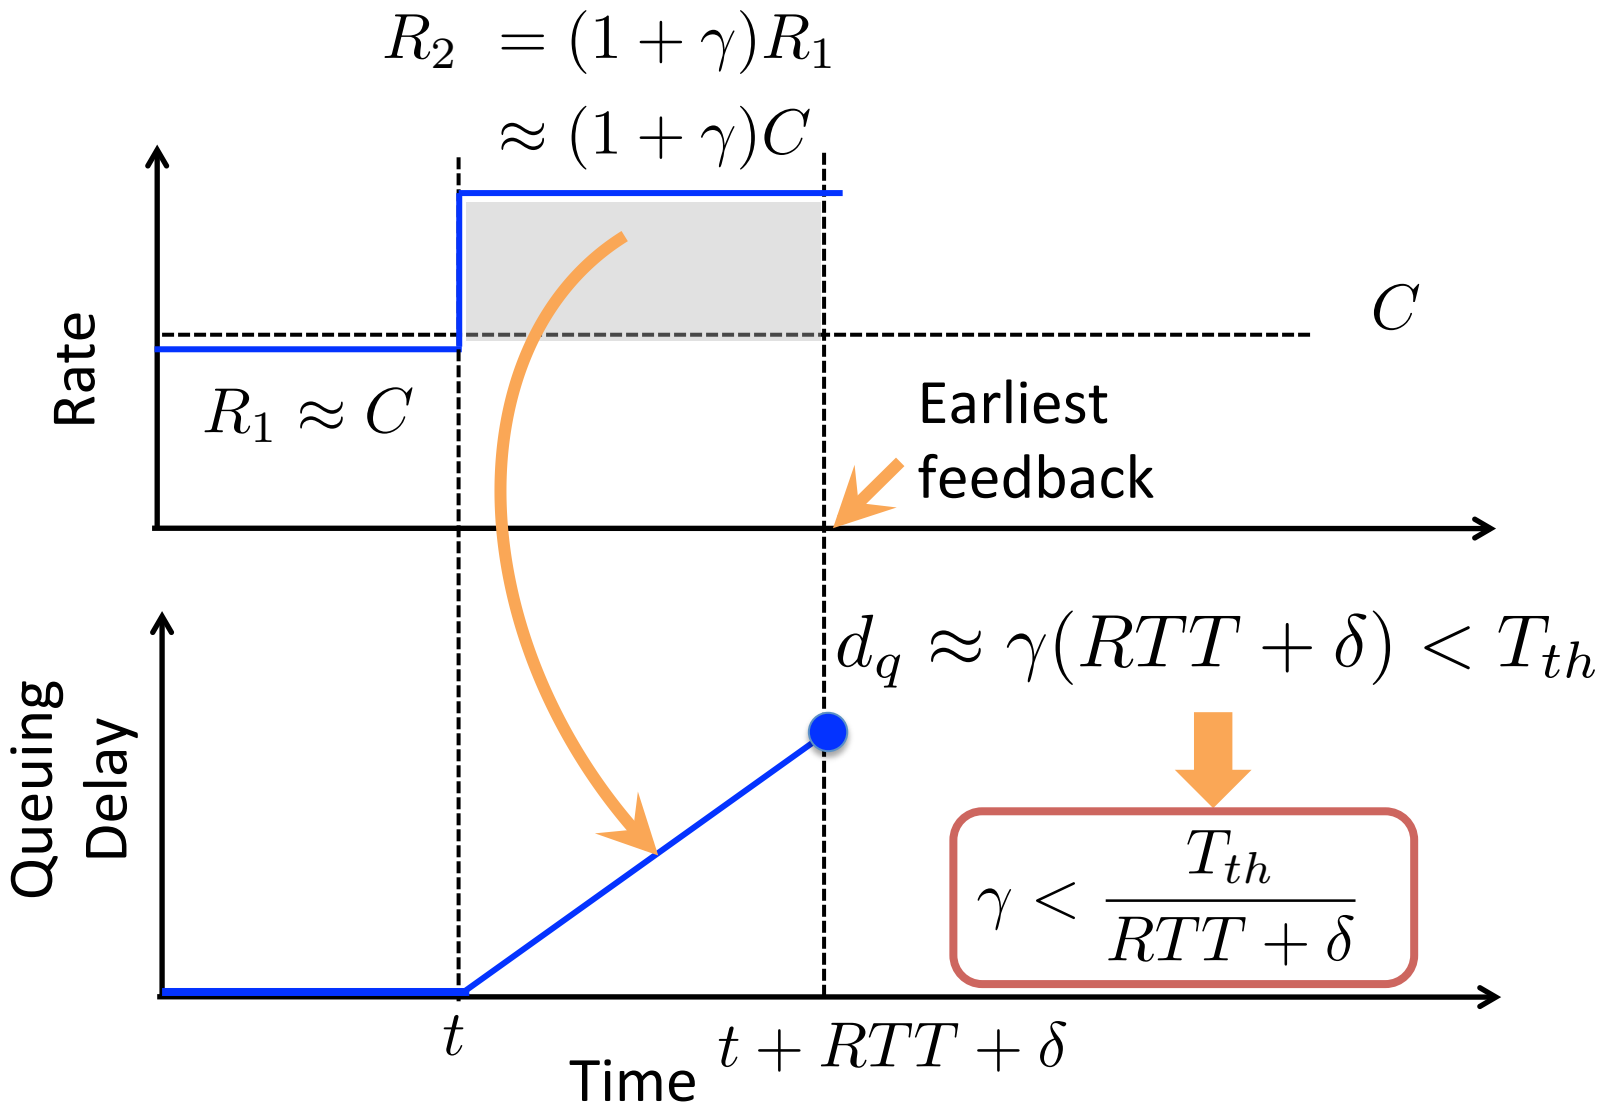

List of Test Cases

- Variable Available Capacity w/ Single RMCAT Flow
- Variable Available Capacity w/ Multiple RMCAT Flows
- Congested Feedback Link w/ Bidirectional Flows
- Competing Flows with same RMCAT Algorithm
- Round Trip Time Fairness
- RMCAT Flow Competing with a Long-Live TCP Flow
- RMCAT Flow Competing with Short TCP Flows
- Media Pause and Resume

Default Settings

Path propagation delay: 50ms
Bottleneck queue depth: ~ 300ms

Default Algorithm Parameters

- ACK interval: $\delta_0 = 100$ ms
- MTU size: 1000 bytes
- Rate range: 150 Kbps ~ 1.5 Mbps
- Scaling parameters: $\kappa = 1.0, \eta = 2.0$
- Upper bound on RTT: $\tau_o = 500$ ms
- Reference delay: $x_{ref} = 10$ ms
- Delay warping: $d_{th} = 50$ ms, $d_{max} = 400$ ms
- Delay weights: $d_M = 200$ ms, $d_P = 1.0$ s
- Accelerated ramp-up: $T_{th} = 50$ ms, $\gamma_0 = 0.5$

Variable Available Capacity w/ Single RMCAT Flow

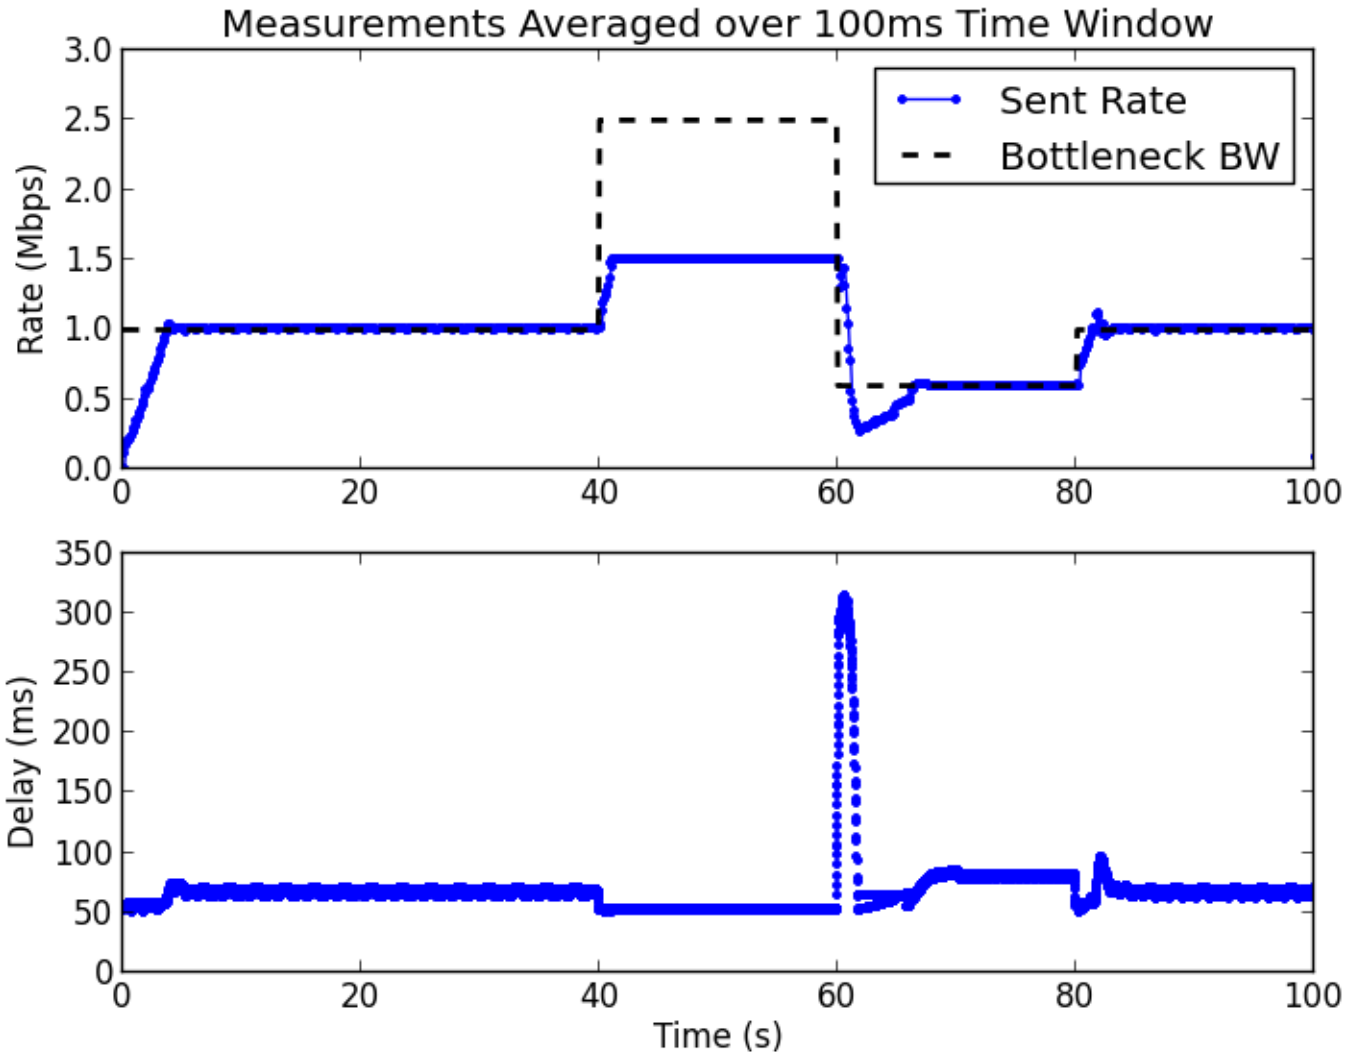

Variable Available Capacity w/ Single RMCAT Flow (Zoom-In View)

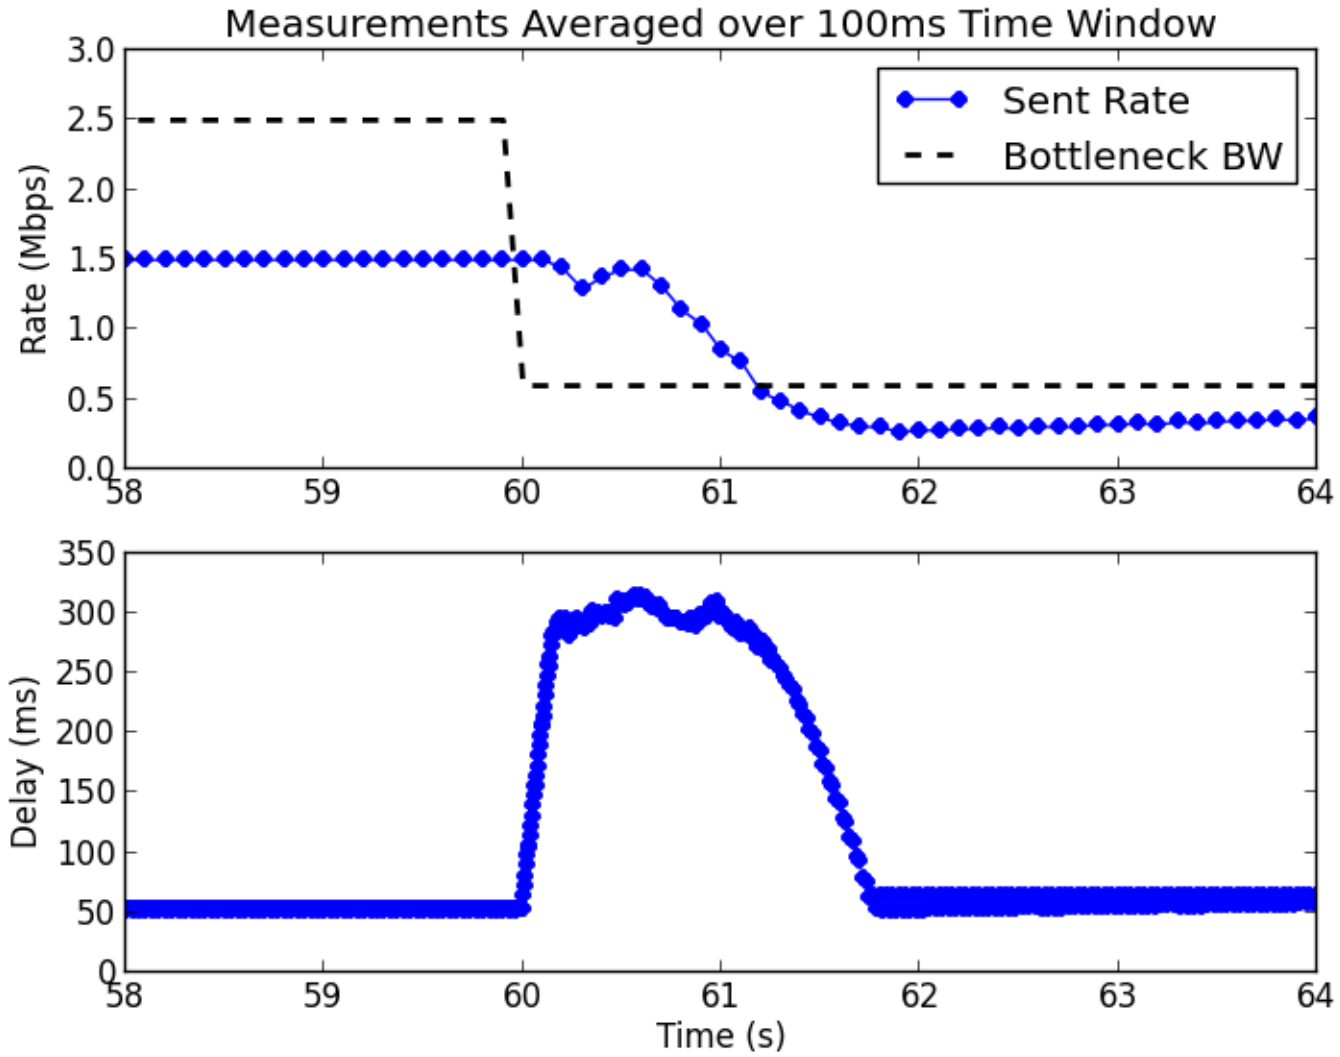

Variable Available Capacity w/ Multiple RMCAT Flows

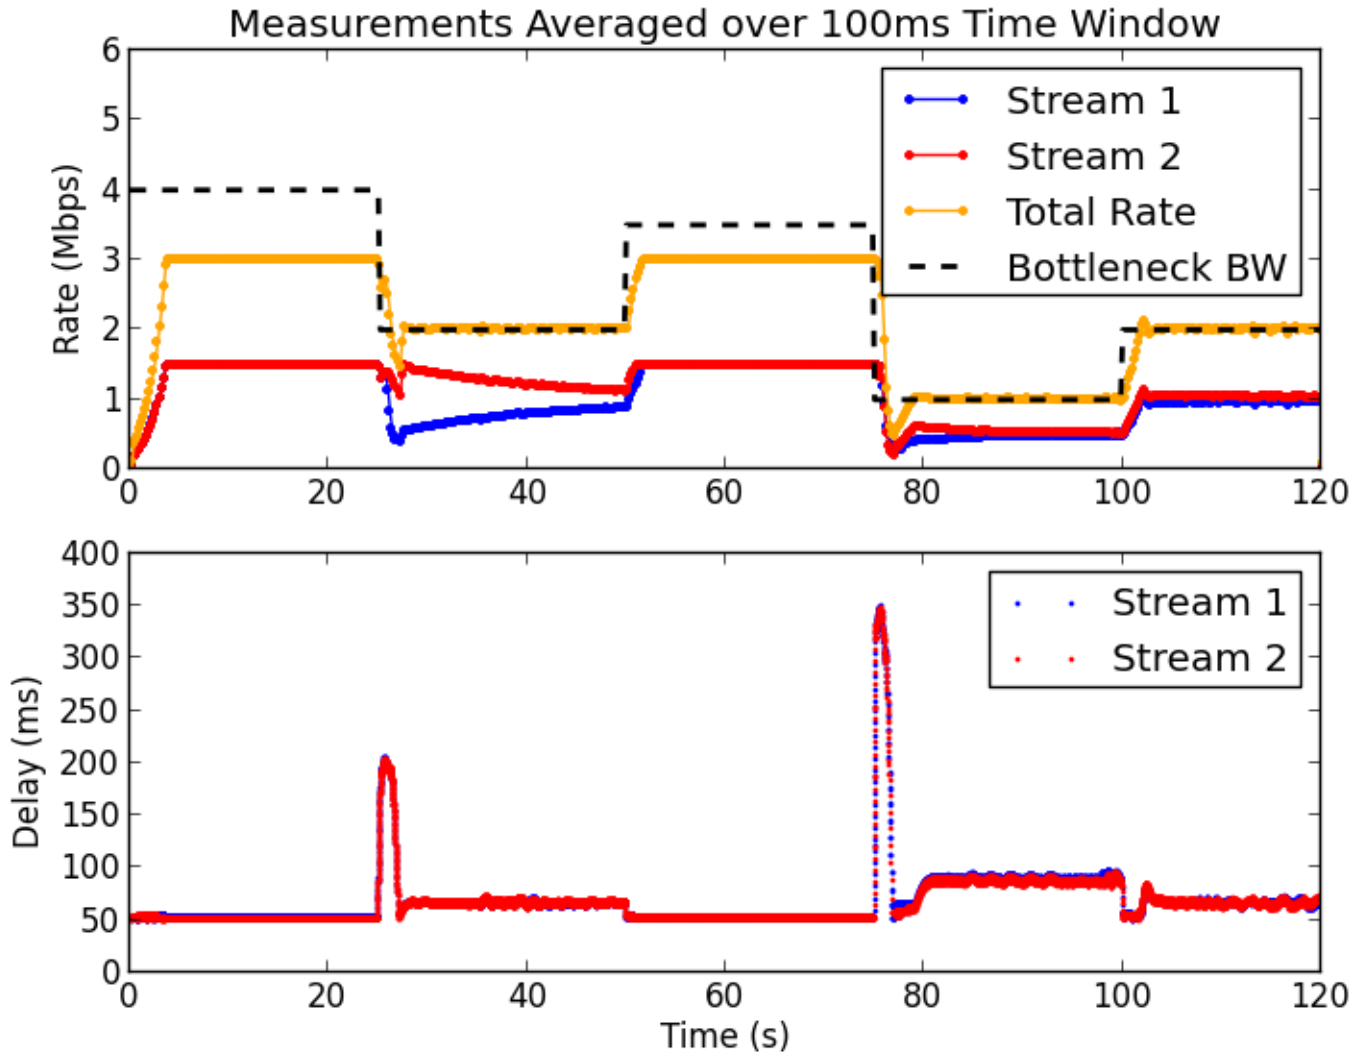

Variable Available Capacity w/ Multiple RMCAT Flows (Zoom-In View)

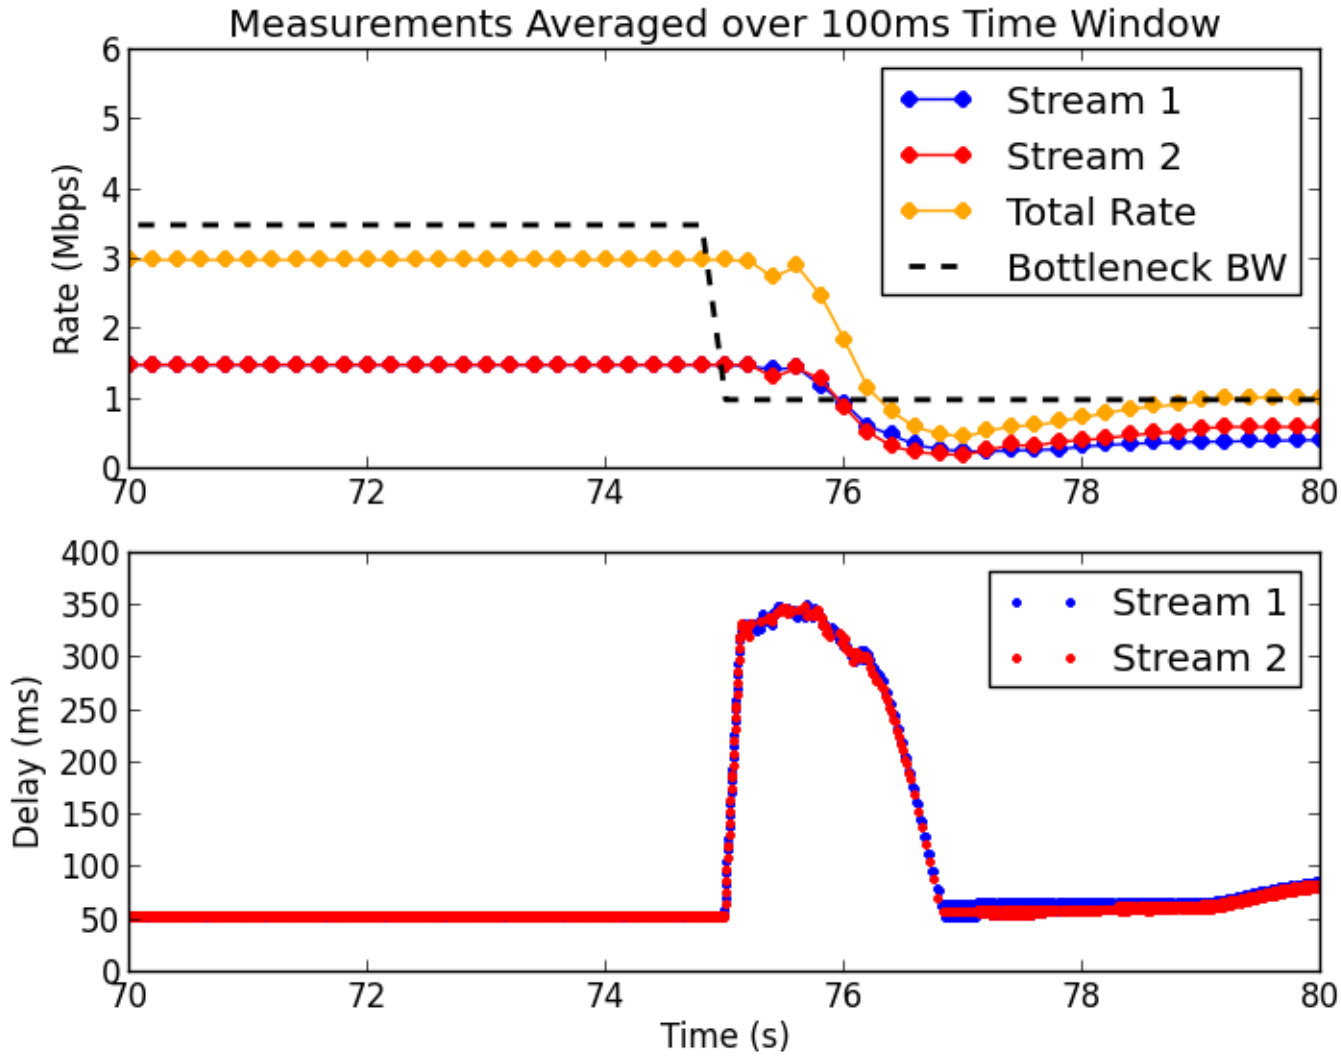

Congested Feedback w/ Bi-directional RMCAT Flows: Forward Direction

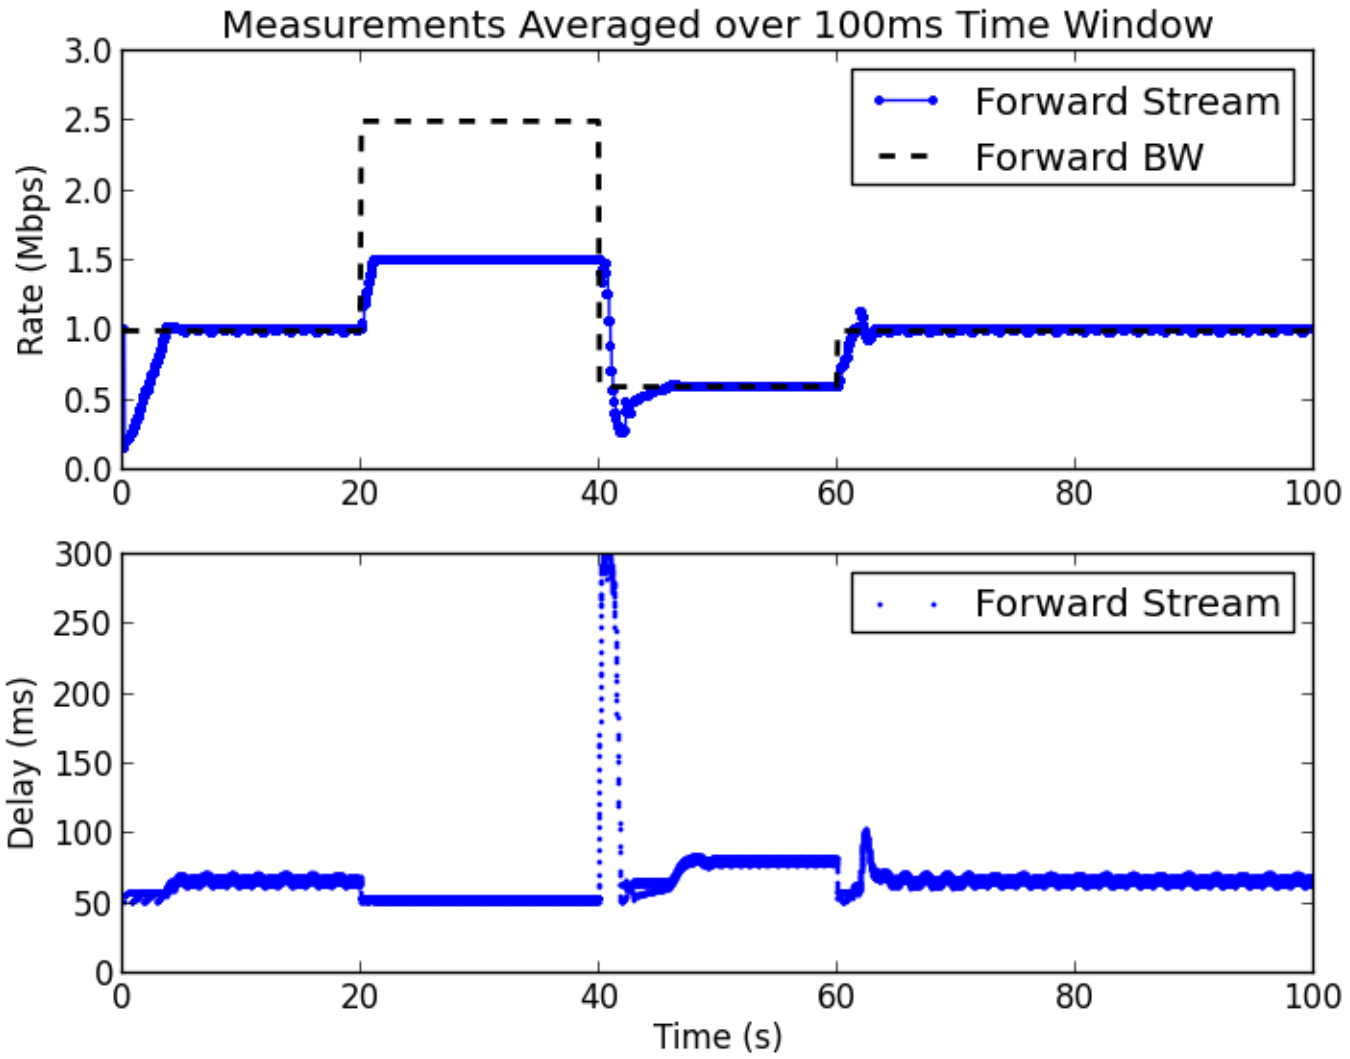

Congested Feedback w/ Bi-directional RMCAT Flows: Backward Direction

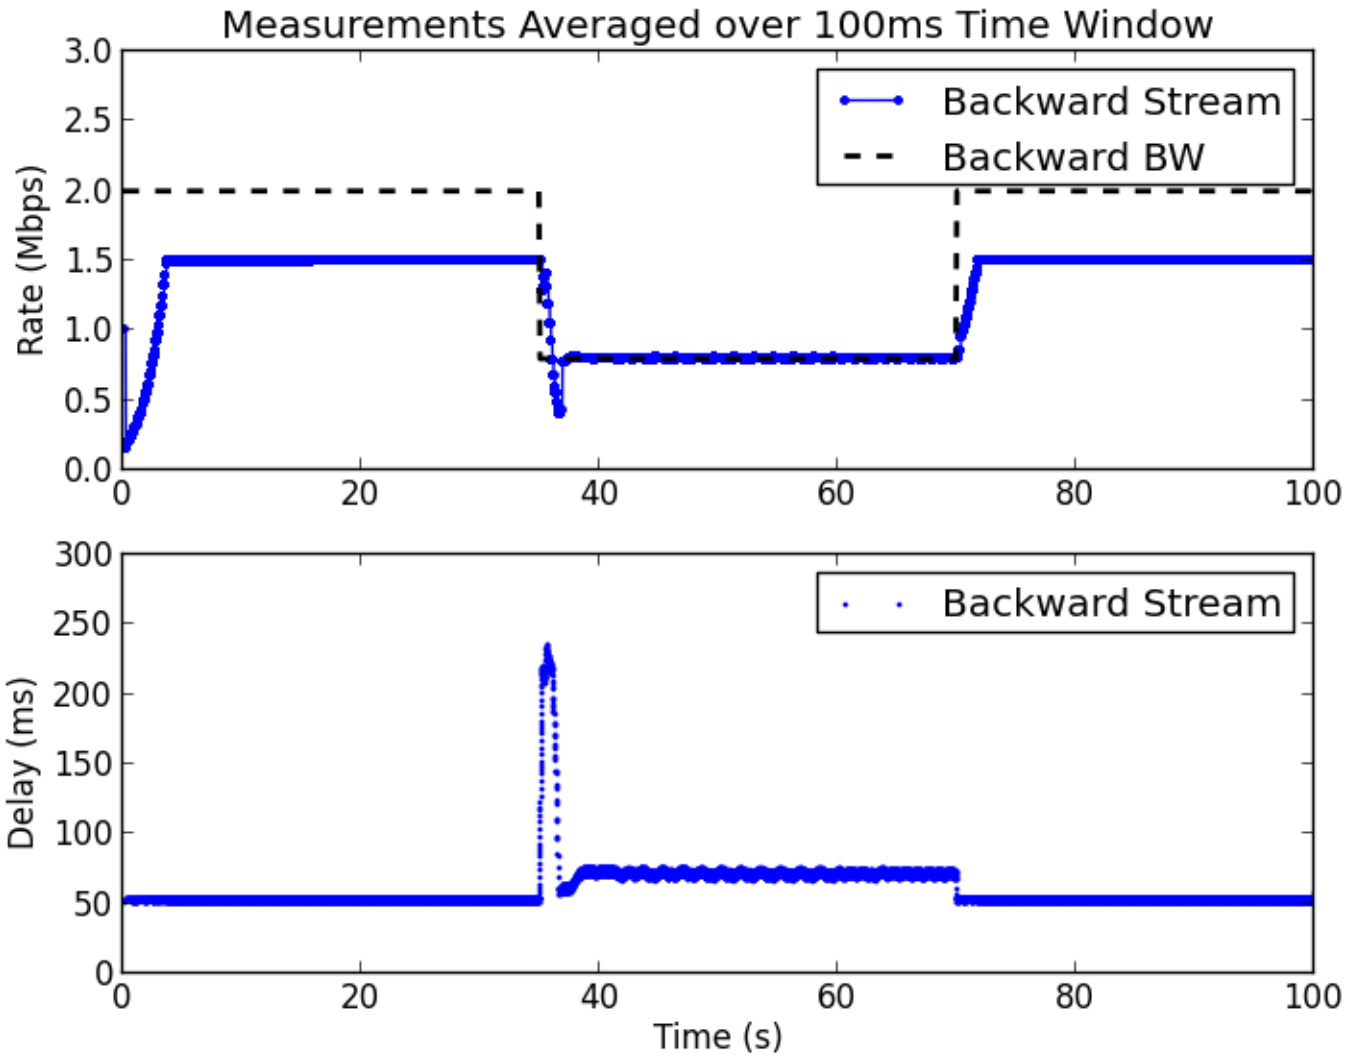

Congested Feedback w/ TCP Backward Traffic

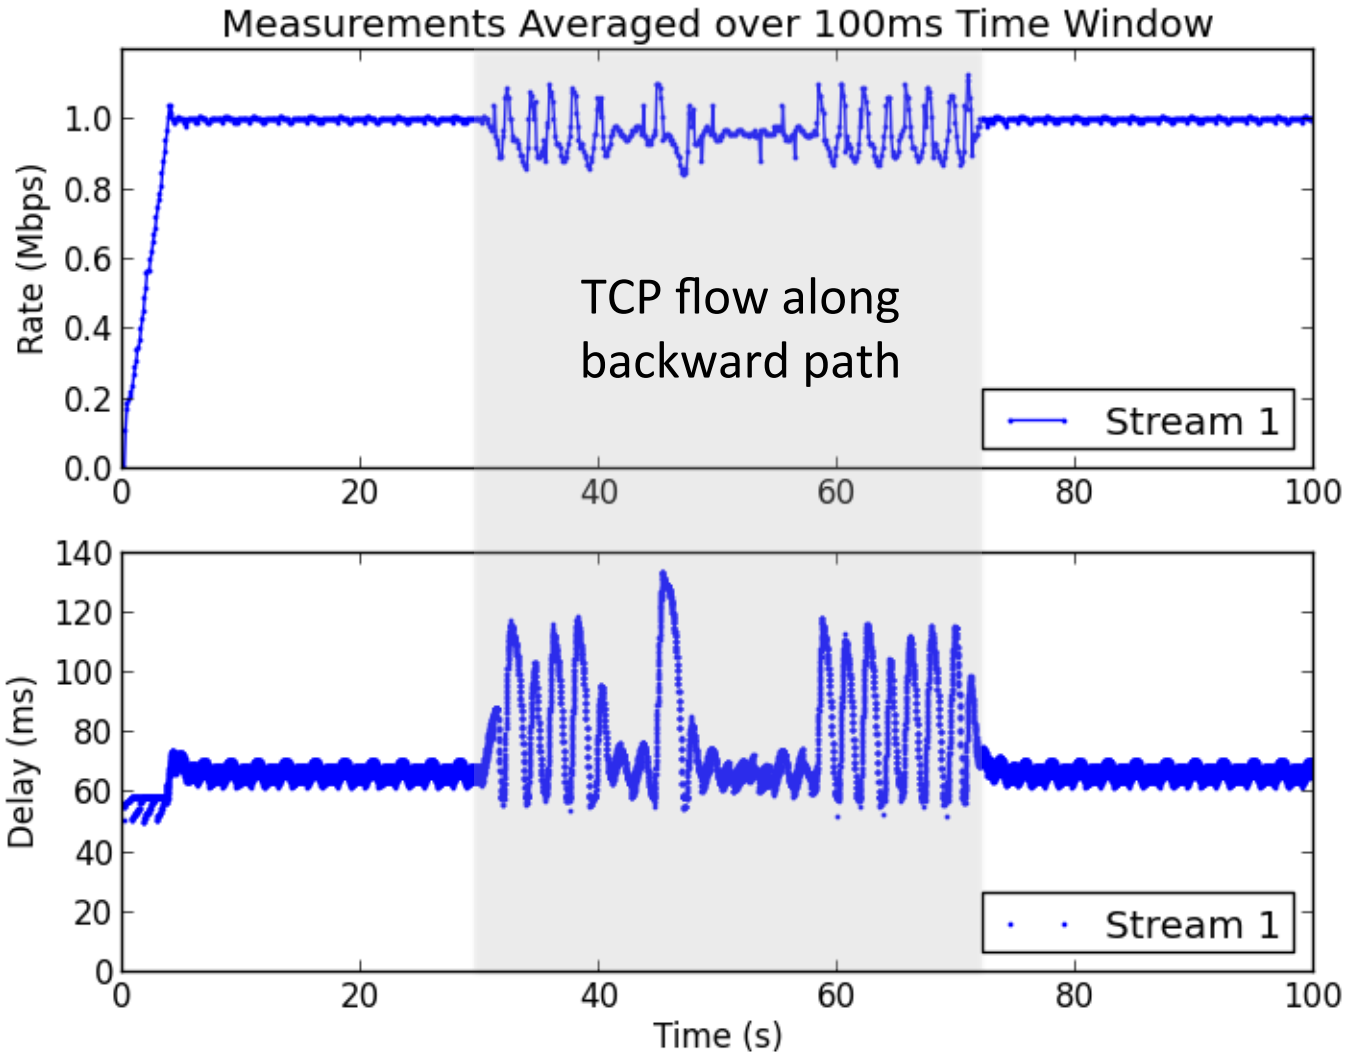

Competing Flows w/ Same RMCAT Algorithm

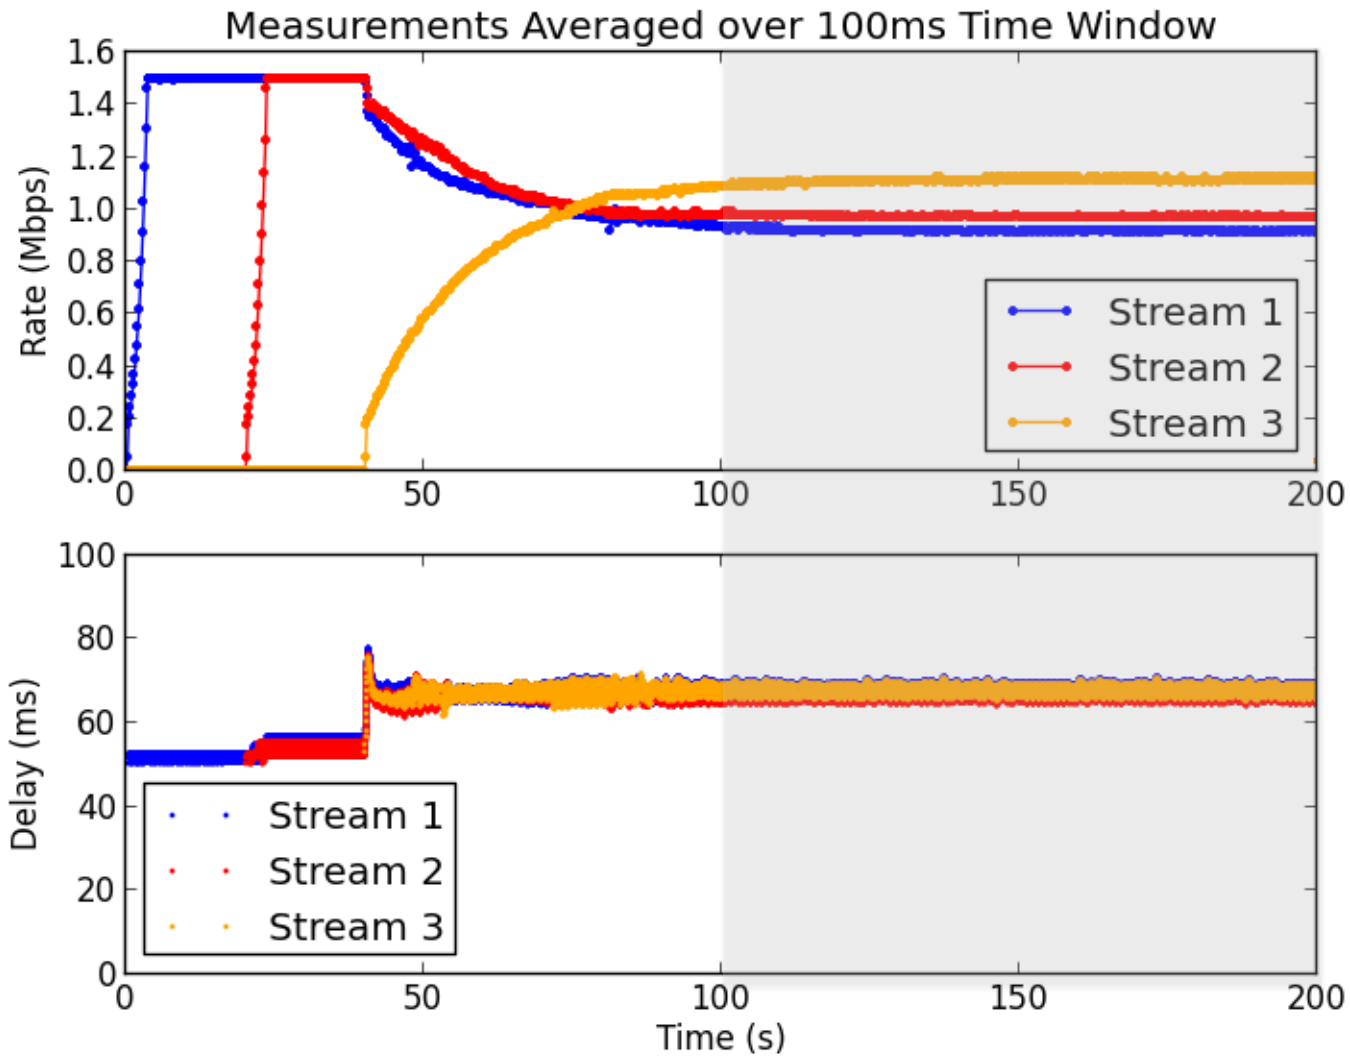

Competing Flows w/ Same RMCAT Algorithm Per-Flow Rate and Delay Statistics

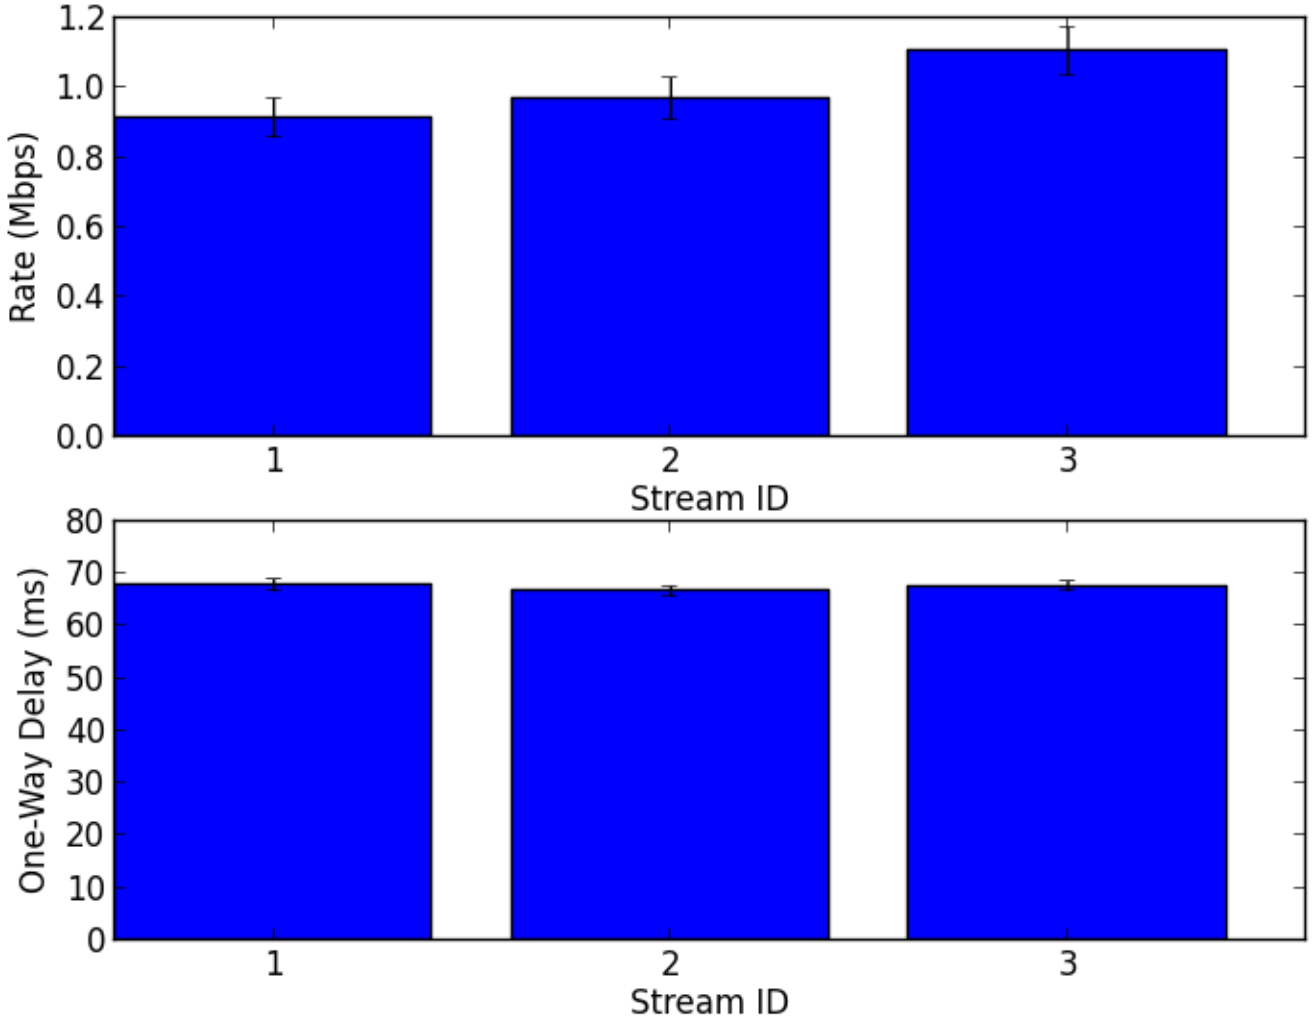

Competing Flows w/ Same RMCAT Algorithm Varying Path Propagation Delay

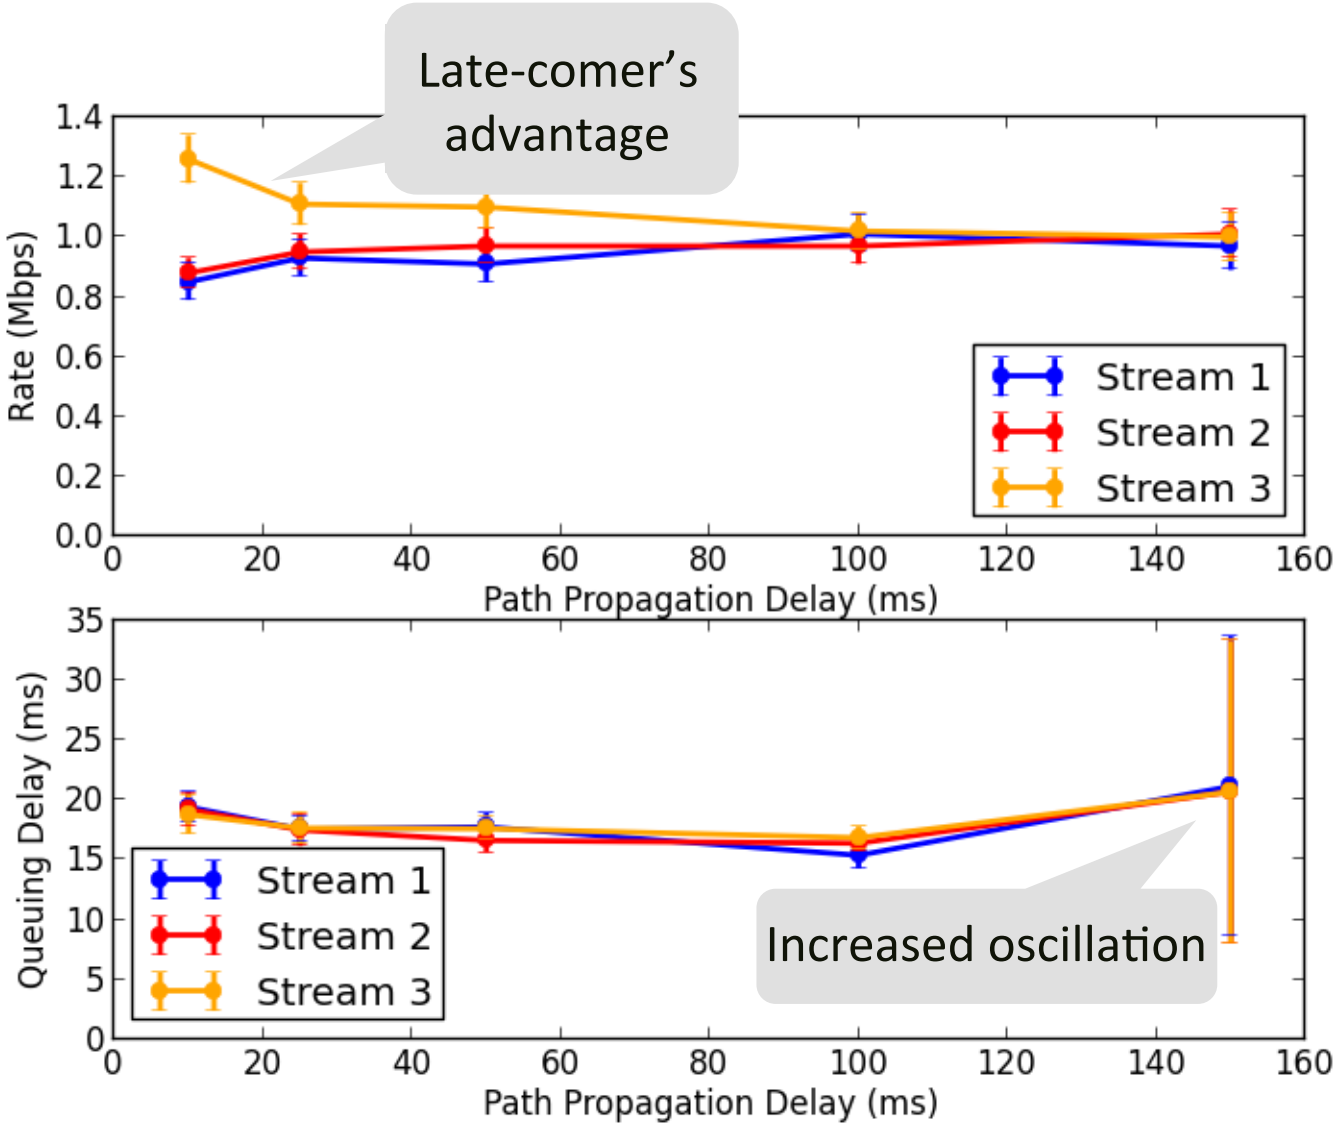

Round Trip Time Fairness

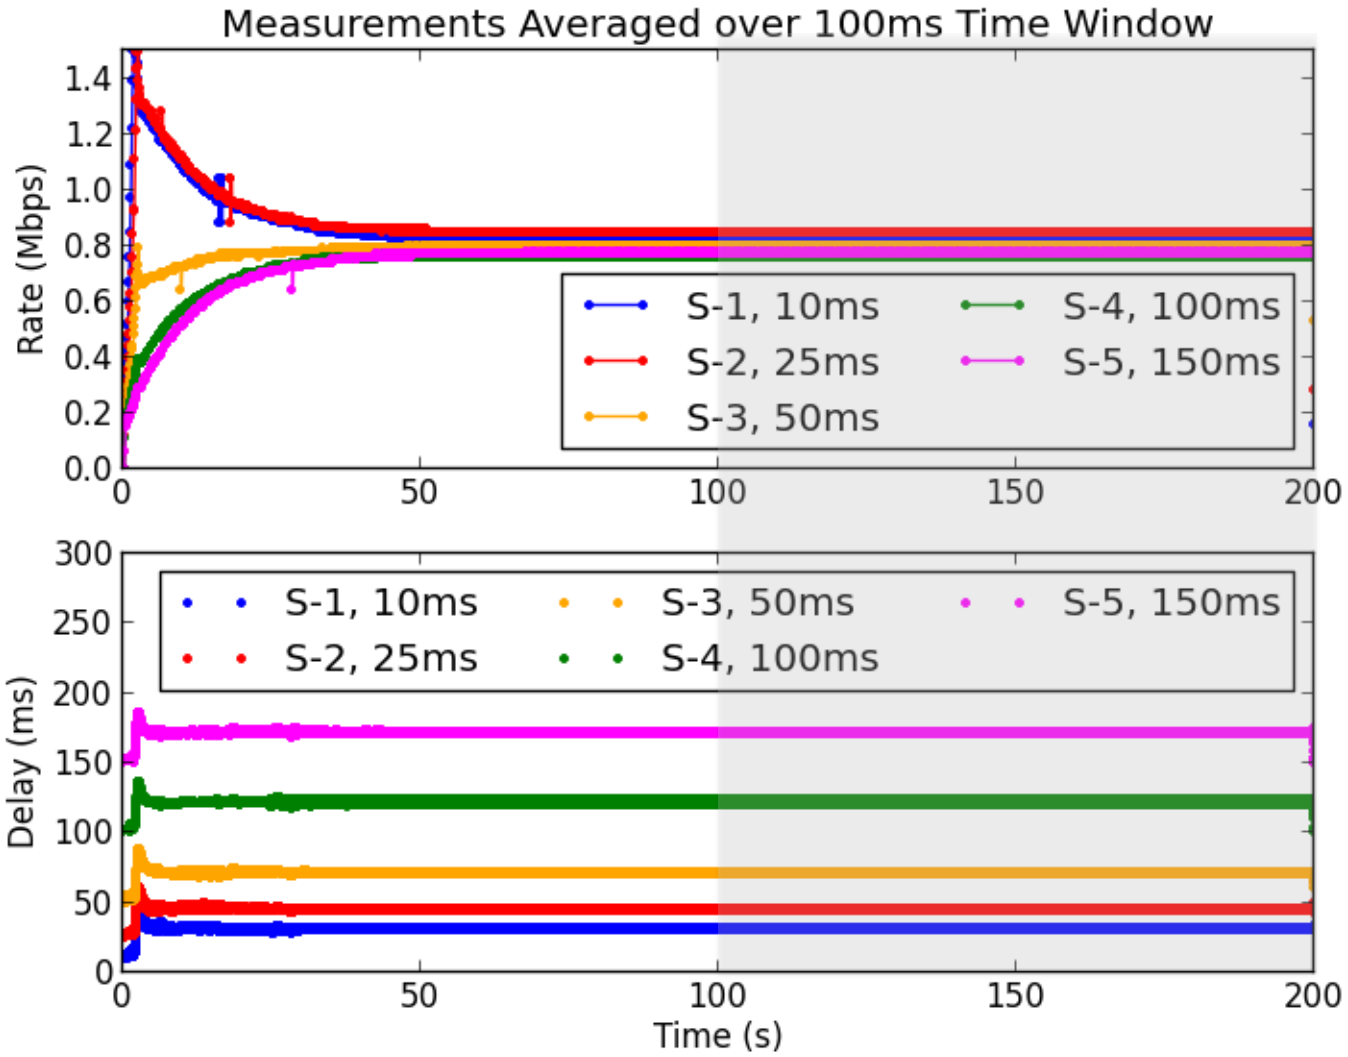

Round Trip Time Fairness
Per-Flow Rate and Delay Statistics

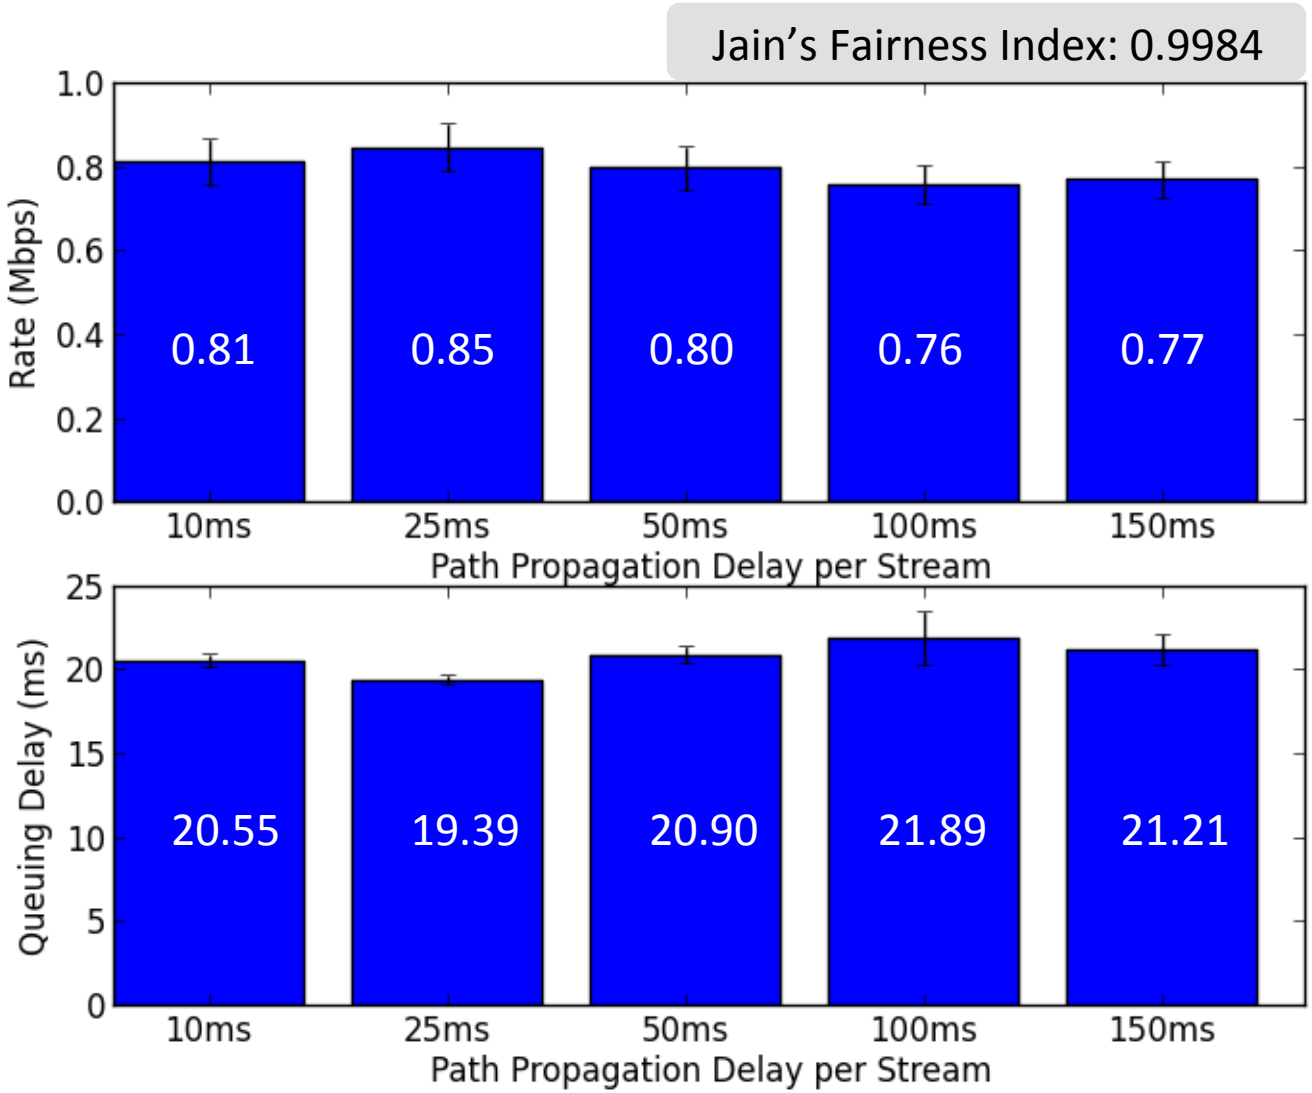

RMCAT Flow competing w/ a Long TCP Flow

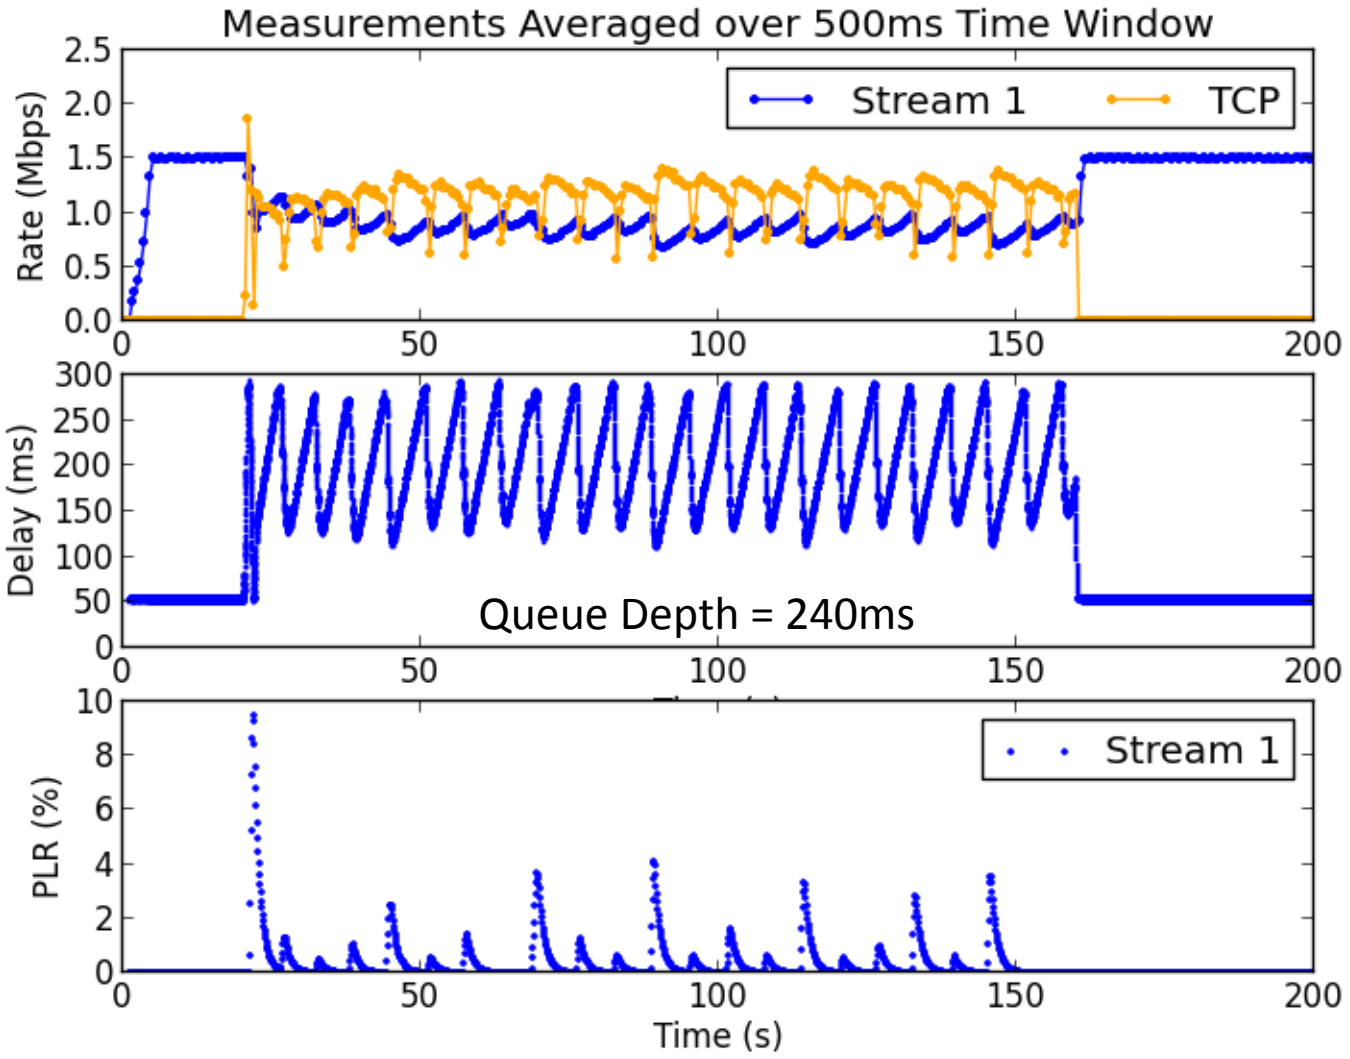

RMCAT Flow competing w/ a Long TCP Flow: Varying Bottleneck Queue Depth

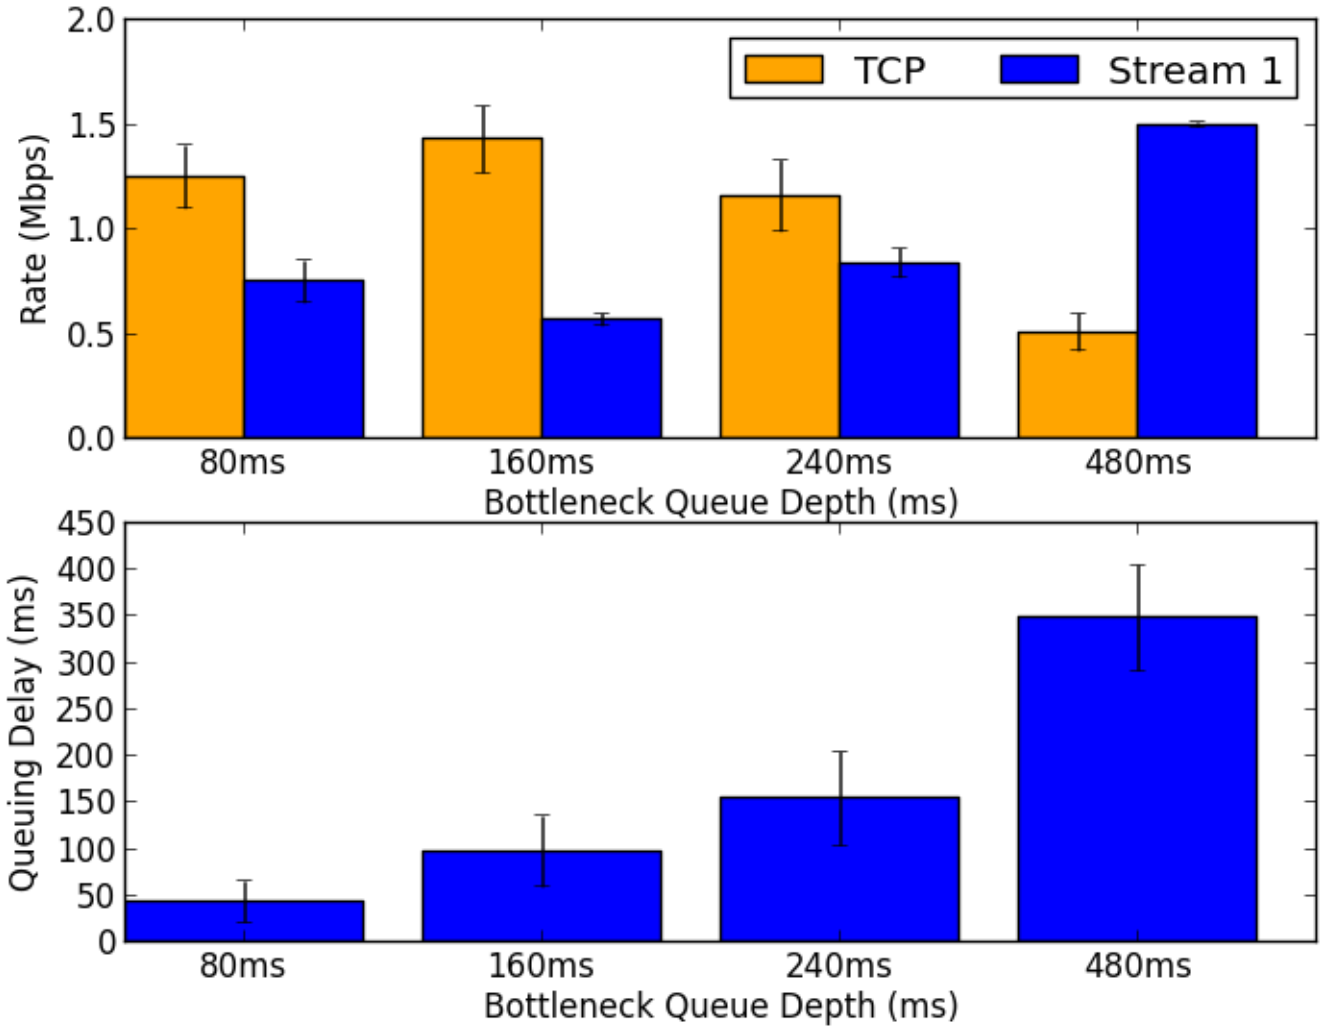

RMCAT Flow competing w/ a Long TCP Flow: Varying Path Propagation Delay

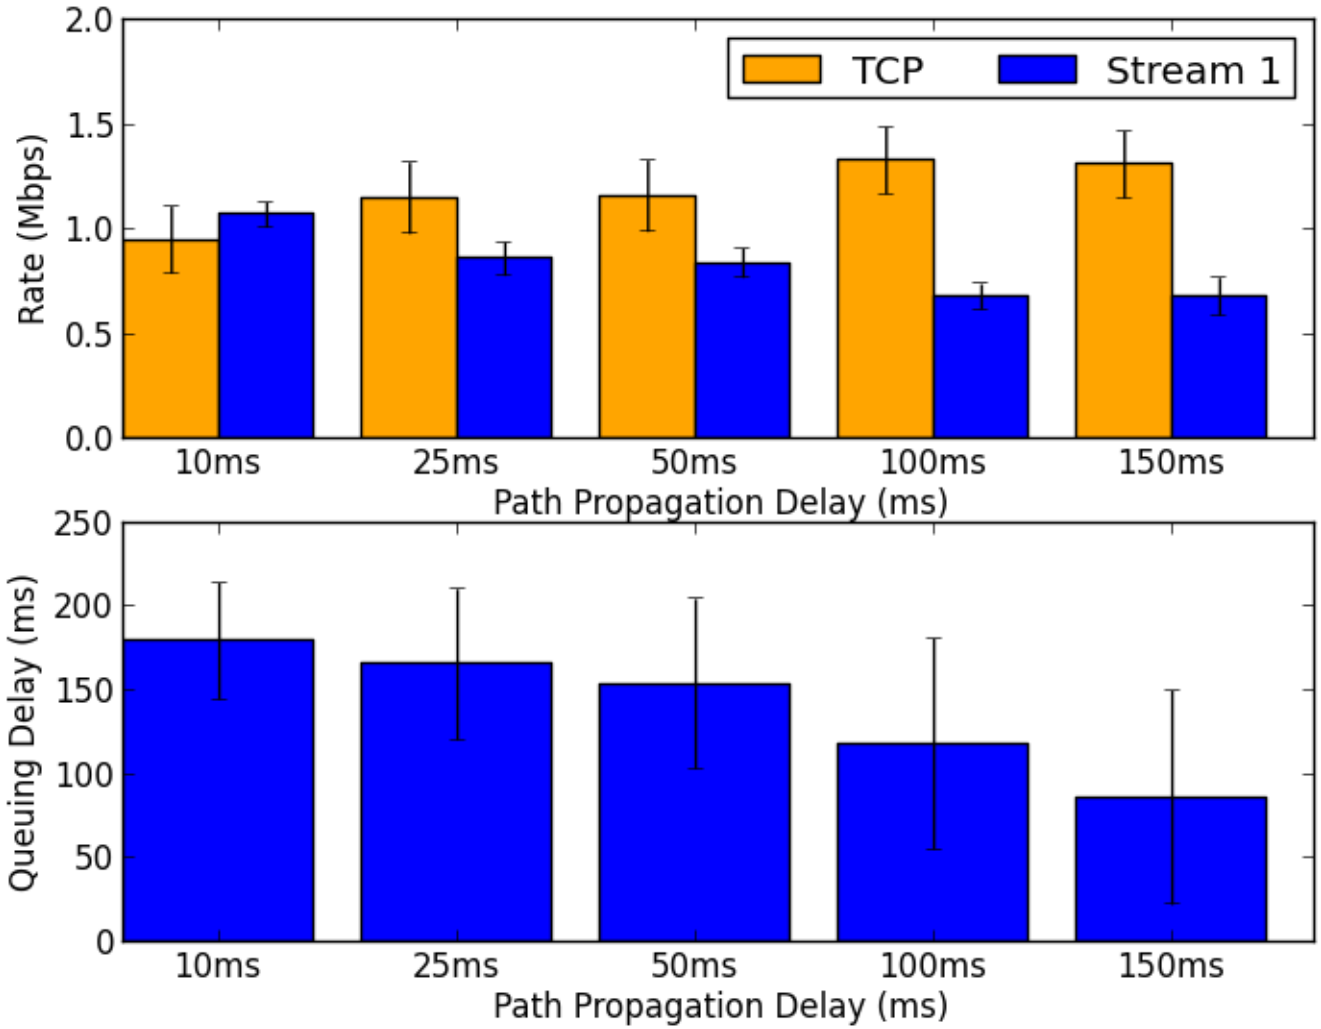

RMCAT Flow competing w/ Short TCP Flows

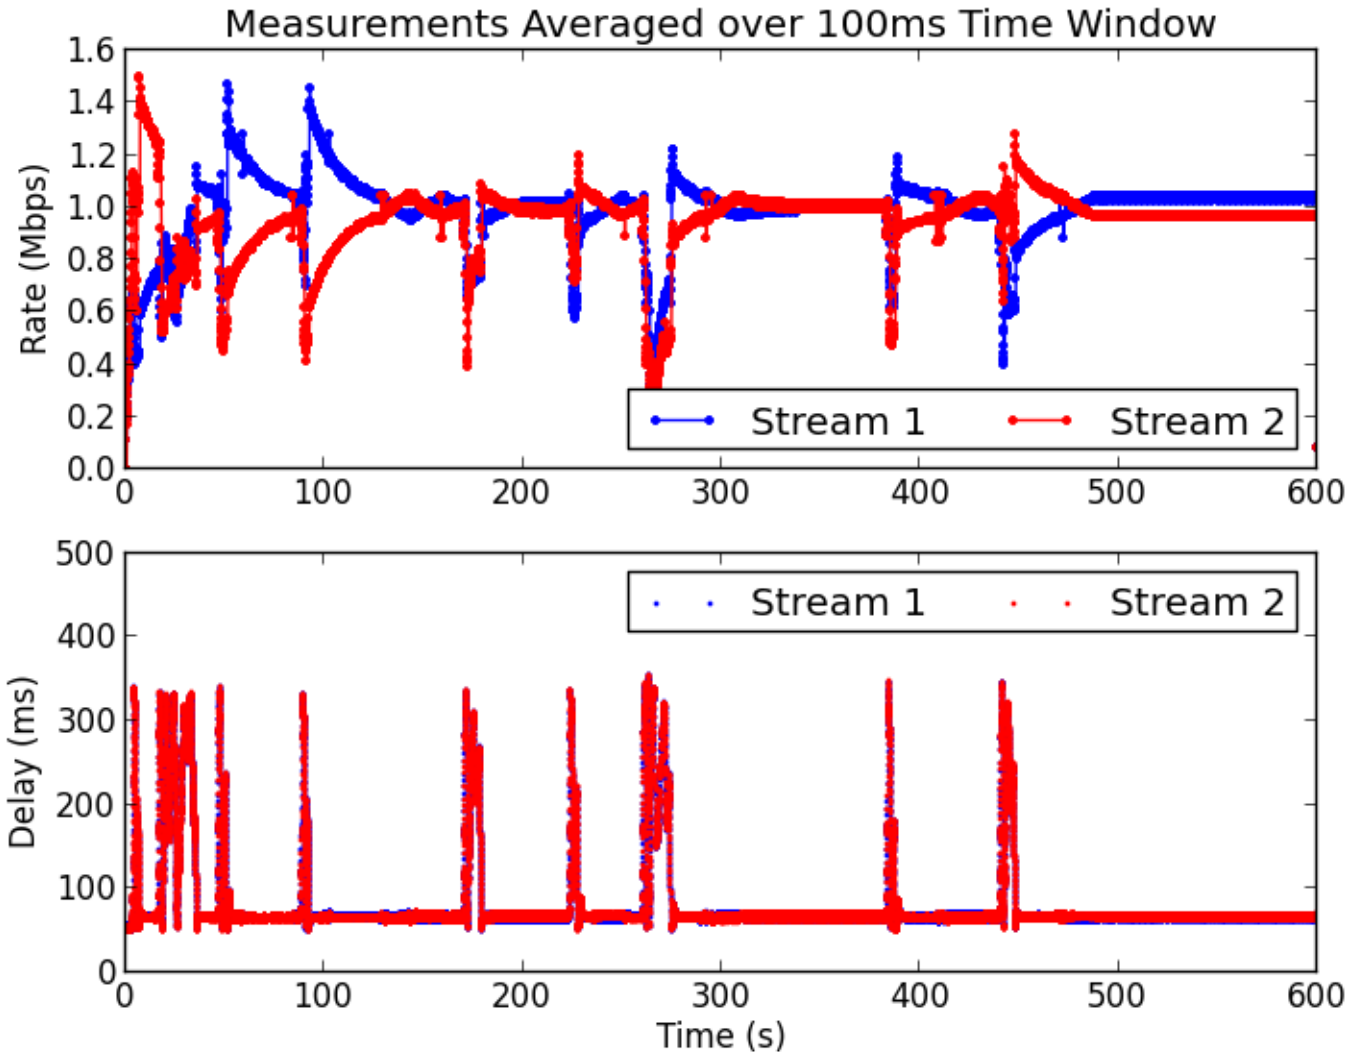

Media Pause and Resume

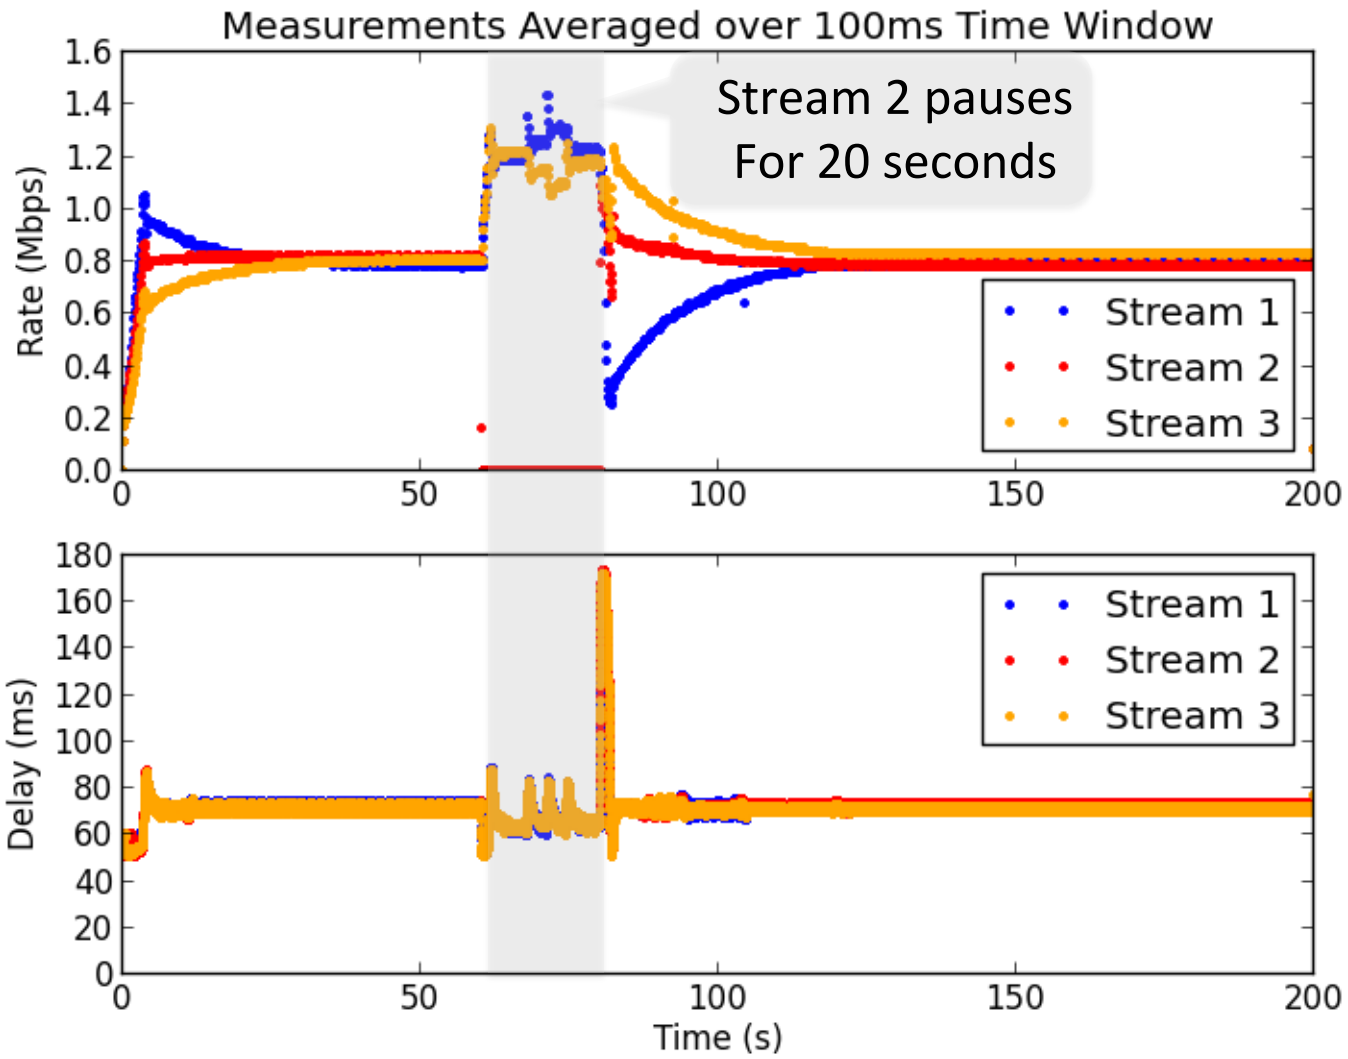

Preliminary Testbed Evaluations

- Testbed Setting:
  - Bottleneck link capacity: 1Mbps
  - Path propagation delay: 0ms
  - Bottleneck queue depth: 500ms
  - VM-based sender and receiver

- Windows-based implementation of NADA:
  - Synthetic traffic source
  - Rate range: 120Kbps – 3.6Mbps
  - Feedback interval target: 100ms
  - *Limitation: Sender/receiver timestamp granularity of 16ms*

Testbed Result: Single Flow

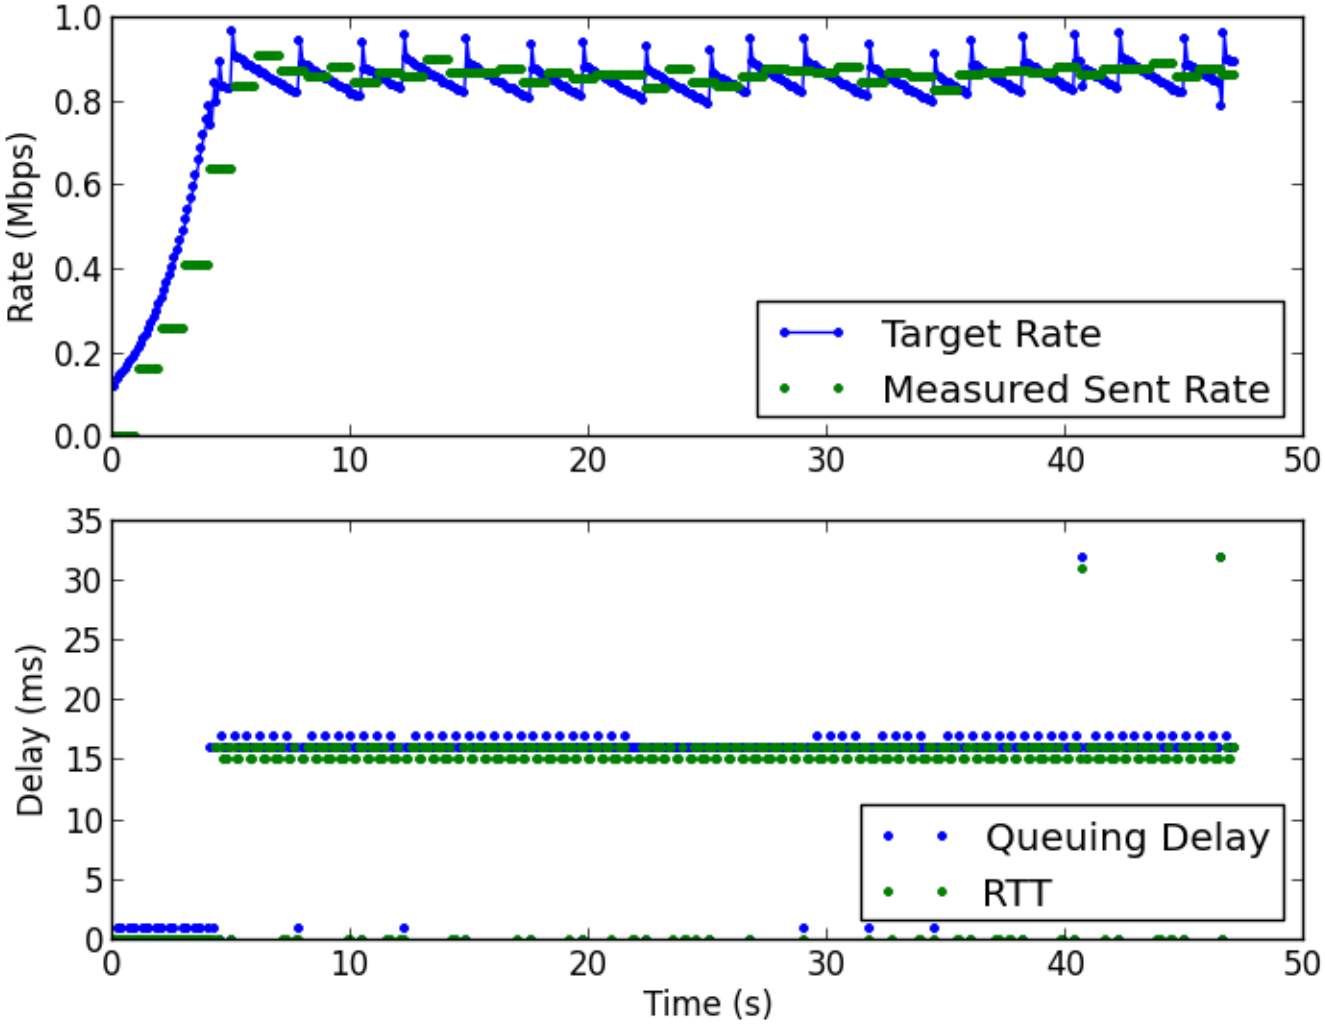

Testbed Result: Two Flows

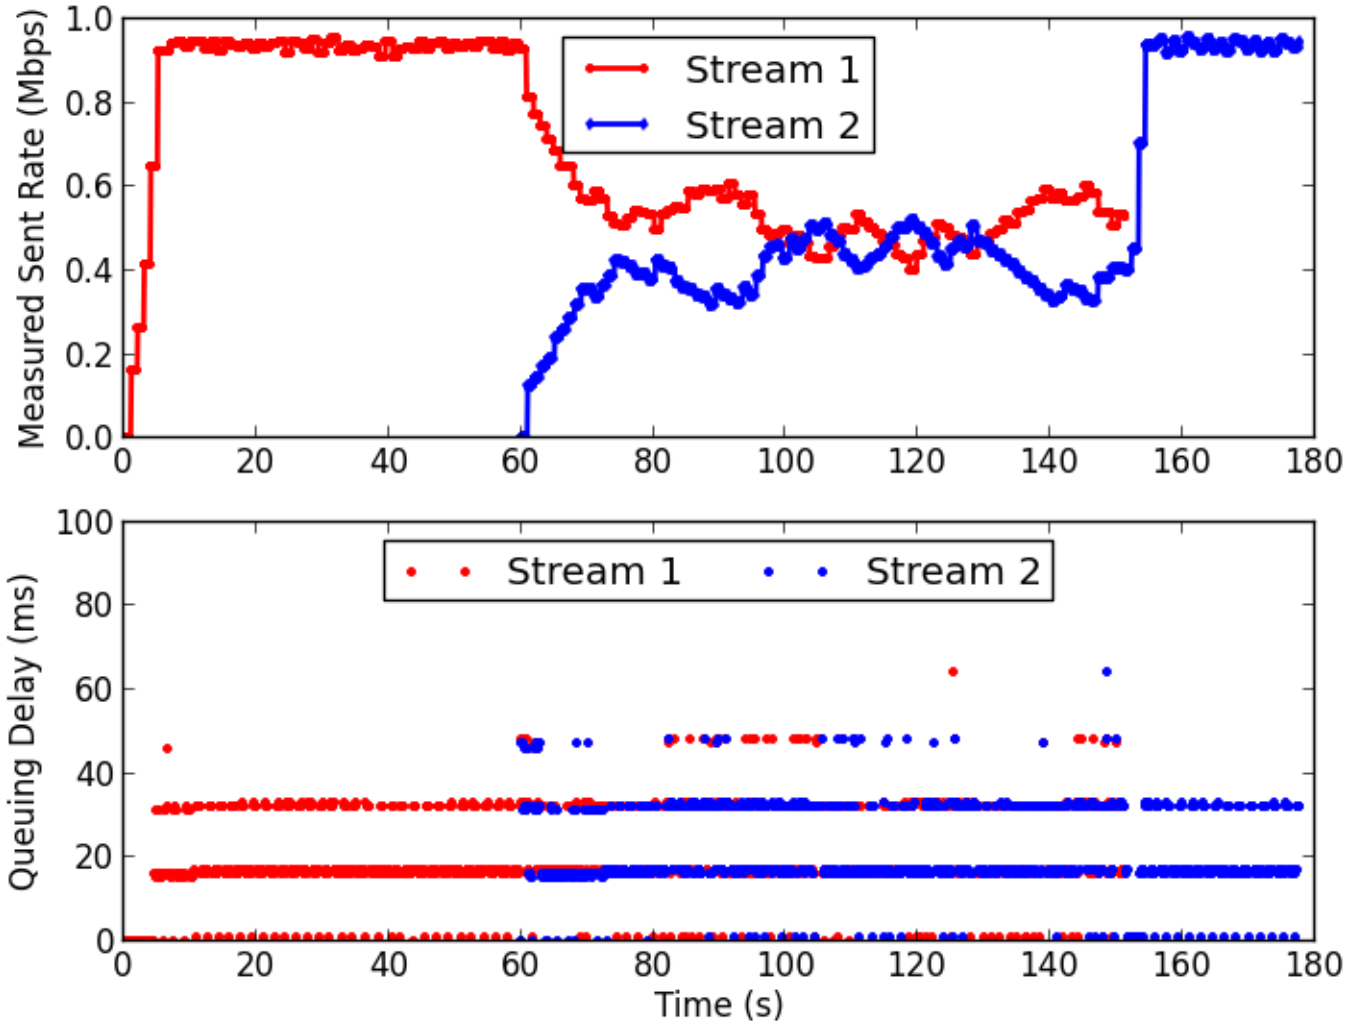

Summary and Next Steps

- Revised algorithm and updated evaluation results:
  - Improved convergence speed
  - Coexists with TCP
  - Maintains stability and RTT fairness

- Next steps:
  - Analysis and tuning of algorithm parameters
  - Calculation of performance metrics
  - More extensive testbed evaluations

Backup Slides

RMCAT Lab: Topology

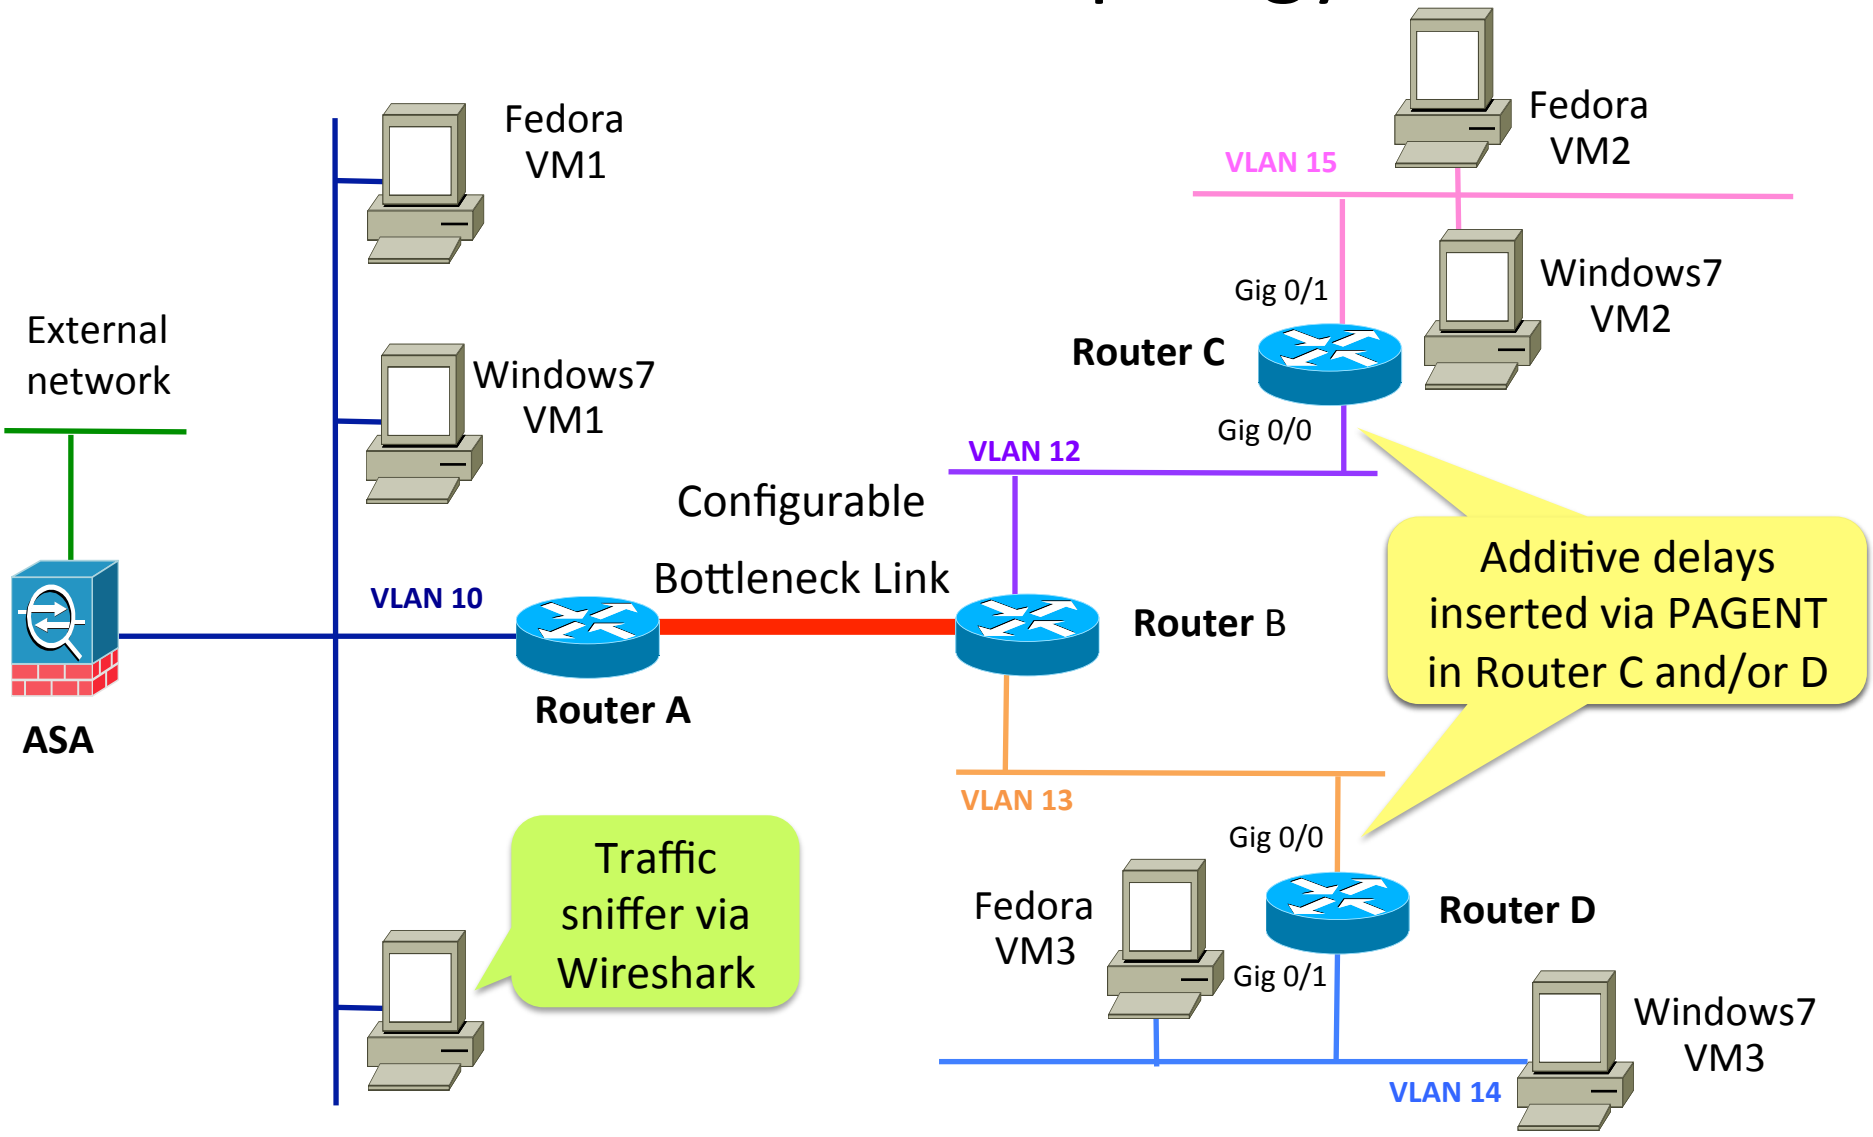

All virtual machines are hosted on a UCS server

Delay Measurement Framework

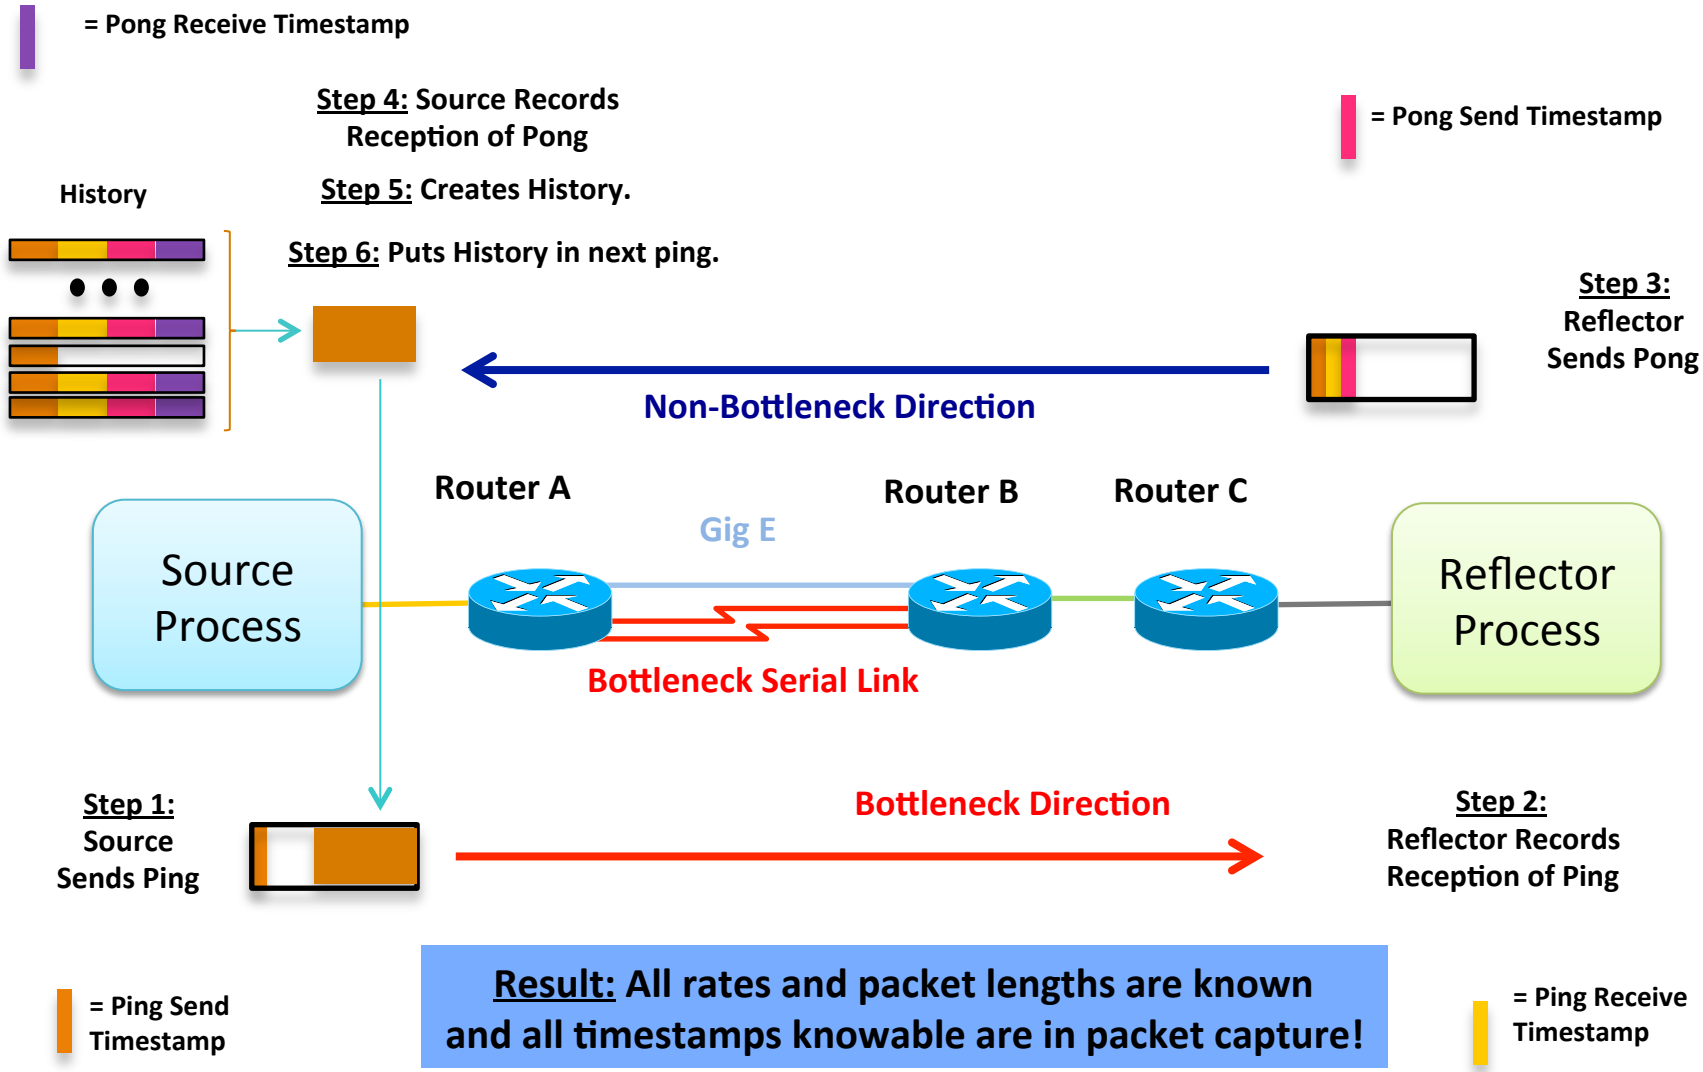

Result: All rates and packet lengths are known and all timestamps knowable are in packet capture!

NADA algorithm implemented at sender process

Feedback Control Loop in Laplace Domain

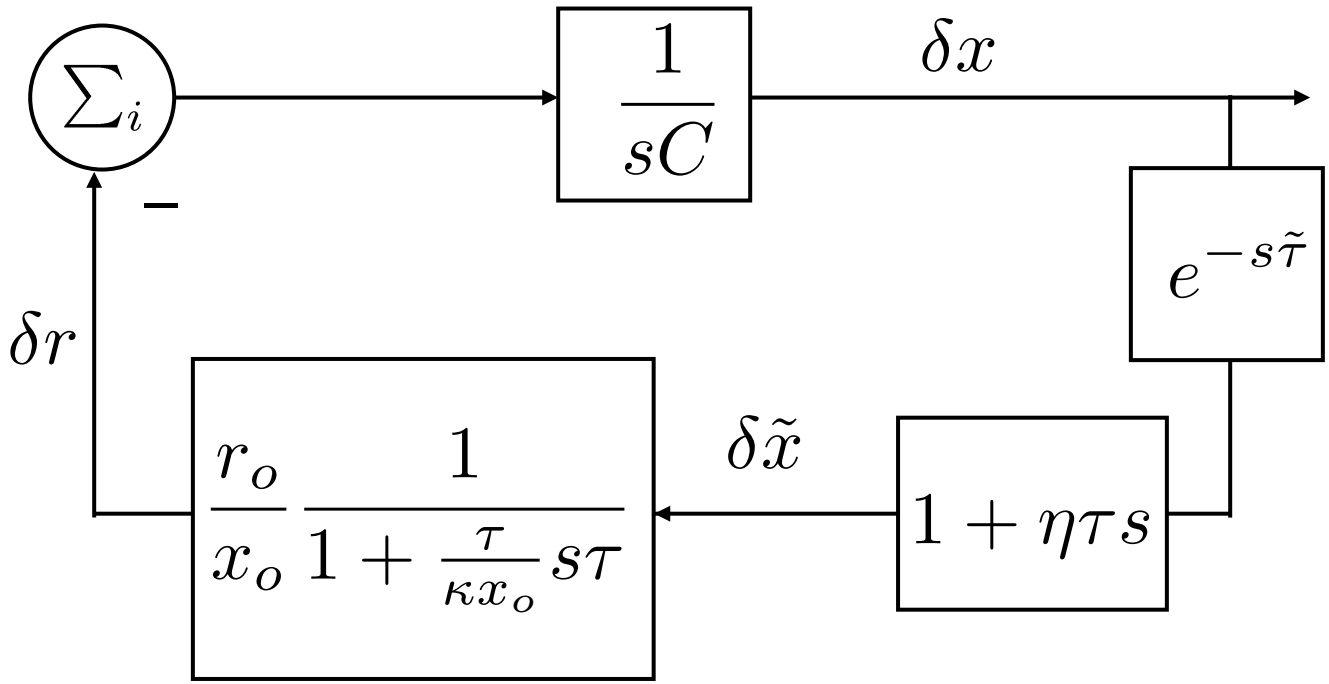

Stability Analysis

- Open-loop transfer function:

$$\mathcal{G}(s) = -\frac{r_o}{C}\frac{1+\eta s\tau}{1+\frac{\tau}{\kappa x_o}s\tau}\frac{e^{-s\tilde{\tau}}}{sx_o}$$

- Low-frequency response:

$$|\mathcal{G}(s)| \propto \frac{r_o}{x_o}\frac{\tilde{\tau}}{C}$$

Ensures weighted fair bandwidth sharing

- High-frequency response:

$$\mathcal{G}(s) \approx -\kappa\eta\frac{r_o}{C}\frac{\tilde{\tau}}{\tau}\frac{e^{-s\tilde{\tau}}}{s\tilde{\tau}}$$

Stability guaranteed by

$$\kappa\eta\frac{\tau}{\tilde{\tau}} < \frac{\pi}{2} \text{ and } \eta\tau \gg 1$$

Numerical Results on Stability Region