Experimental Evaluation of BBR Congestion Control

Mario Hock, Roland Bless, Martina Zitterbart

Published at IEEE ICNP 2017, Oct 10–13, Toronto, Canada

Institute of Telematics (TM)

Google's Congestion Control BBR

- Overall objectives:
  - Replace loss-based congestion control
  - High throughput with a small queue
- Model-based approach
- Experimental evaluation based on [1] and Linux 4.9 implementation

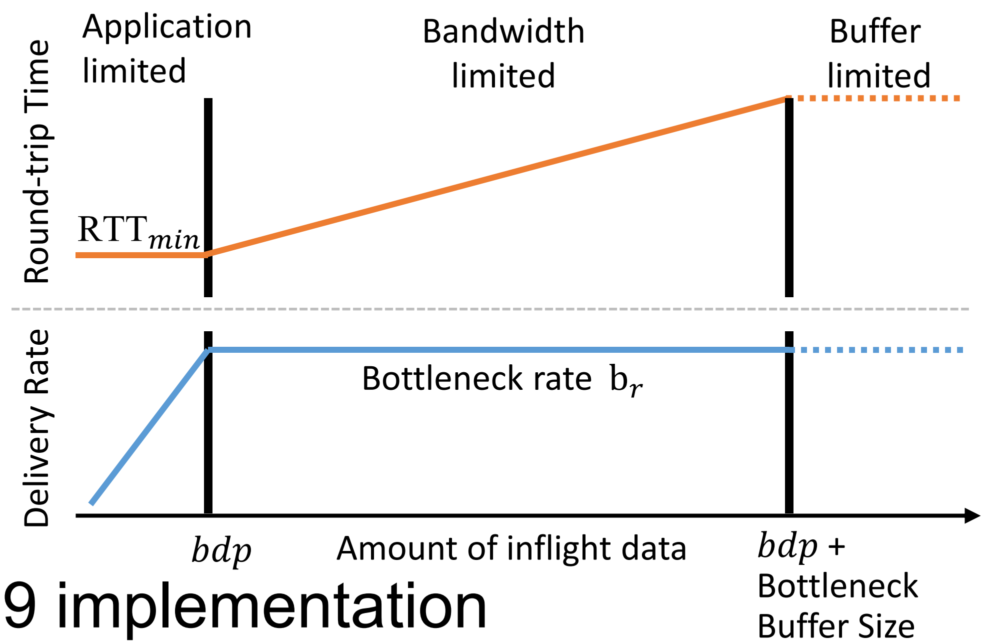

- Key findings [2]
  - Model does not work for multiple flows at the bottleneck
  - Massive packet loss in small buffers
  - Unfairness
  - Suppression of loss-based congestion control

So?

What's wrong with the model?

Network model ok for bottleneck

But used at the sender!

Model lacks dynamics of multiple senders!

Behavior in the Single Flow Case

- Simplified example with a single flow
  - Bottleneck 100Mbit/s, fully utilized

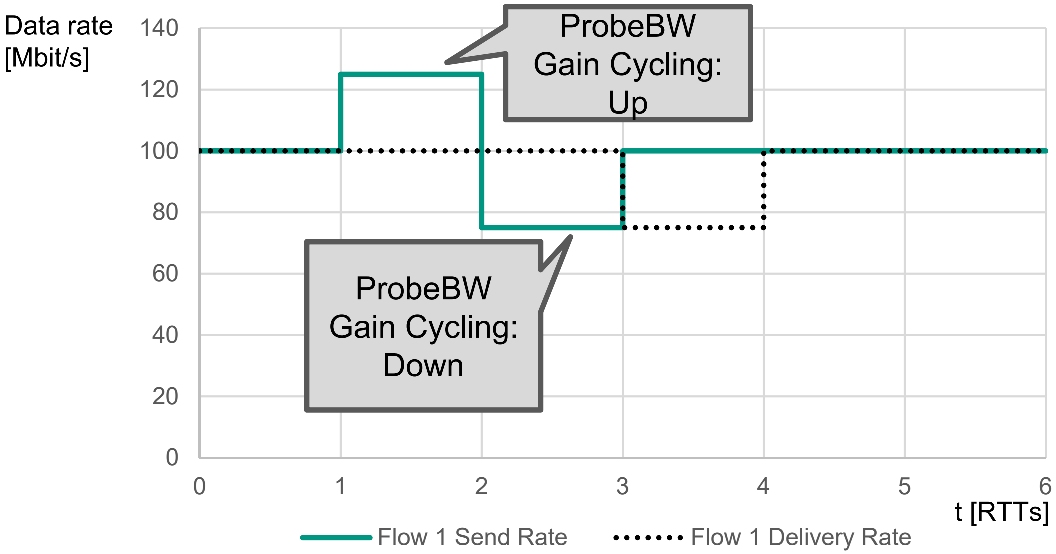

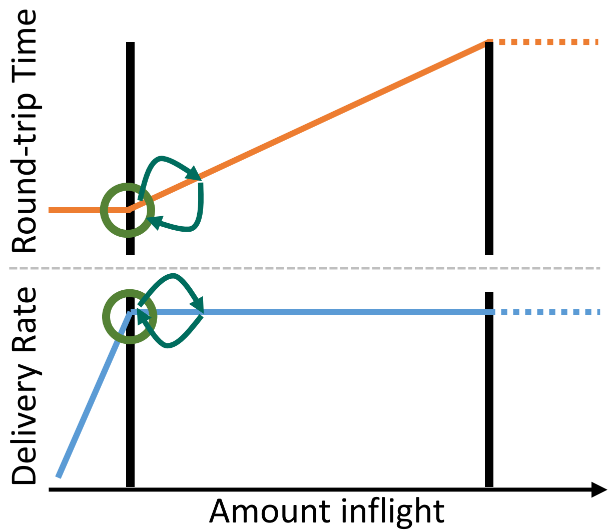

- Flow probes and cannot get higher delivery rate, since bottleneck fully utilized
- Excess data gets queued and removed afterwards

Experimental Evaluation – Single Flow

- 1 Gbit/s bottleneck

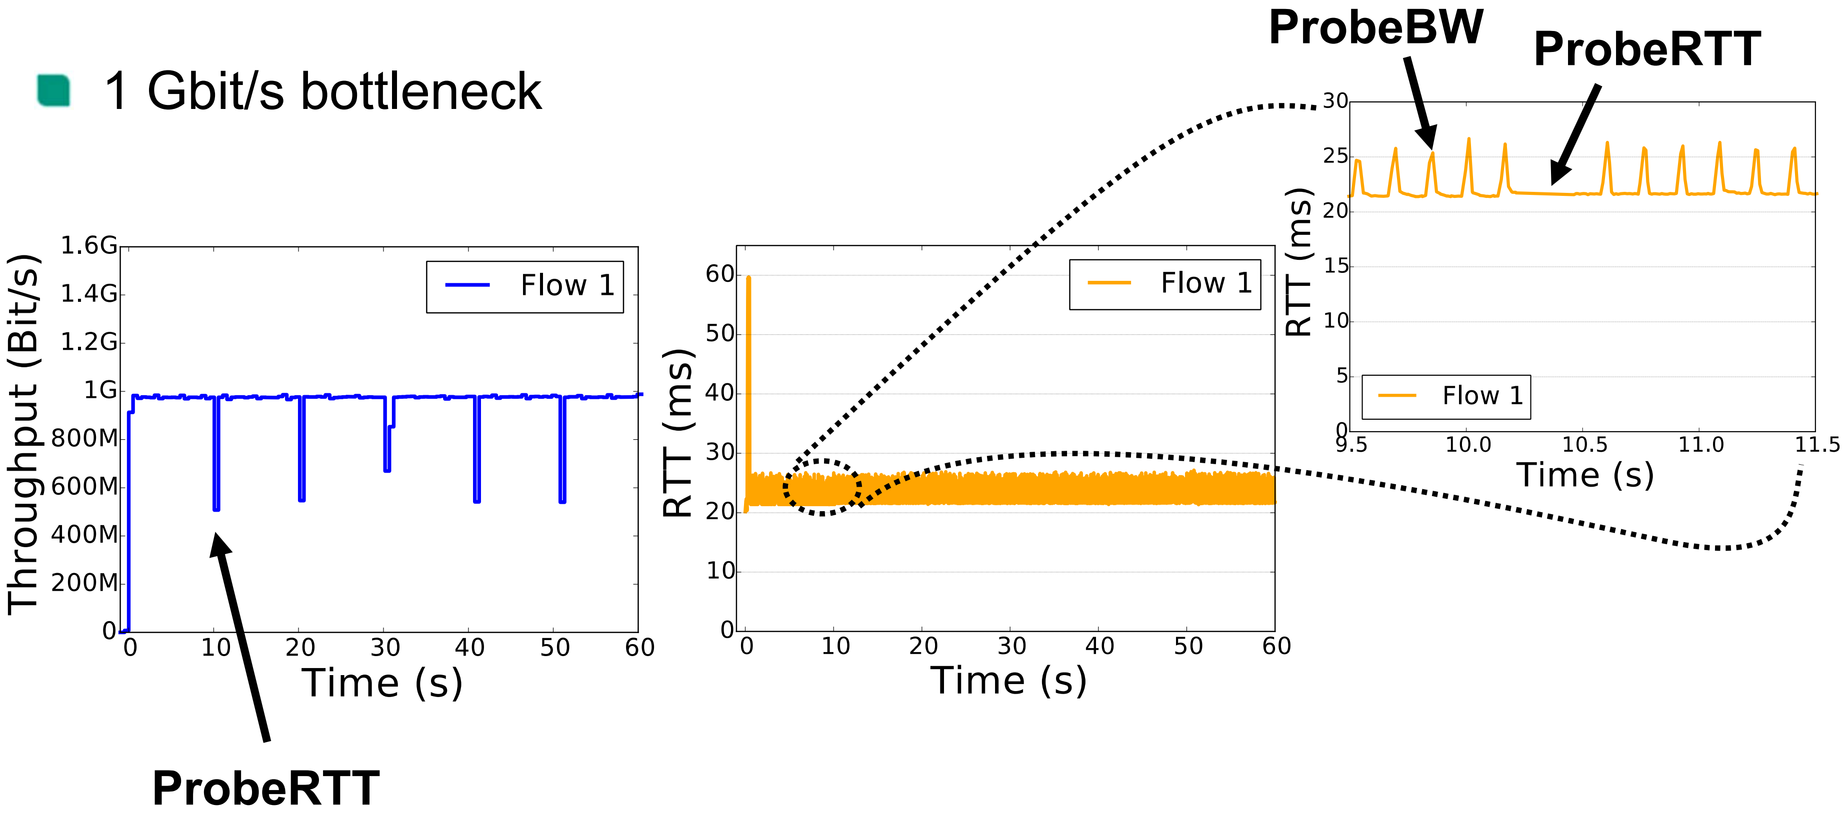

Works as expected since the model fits!

Behavior in the Multiple Flows Case

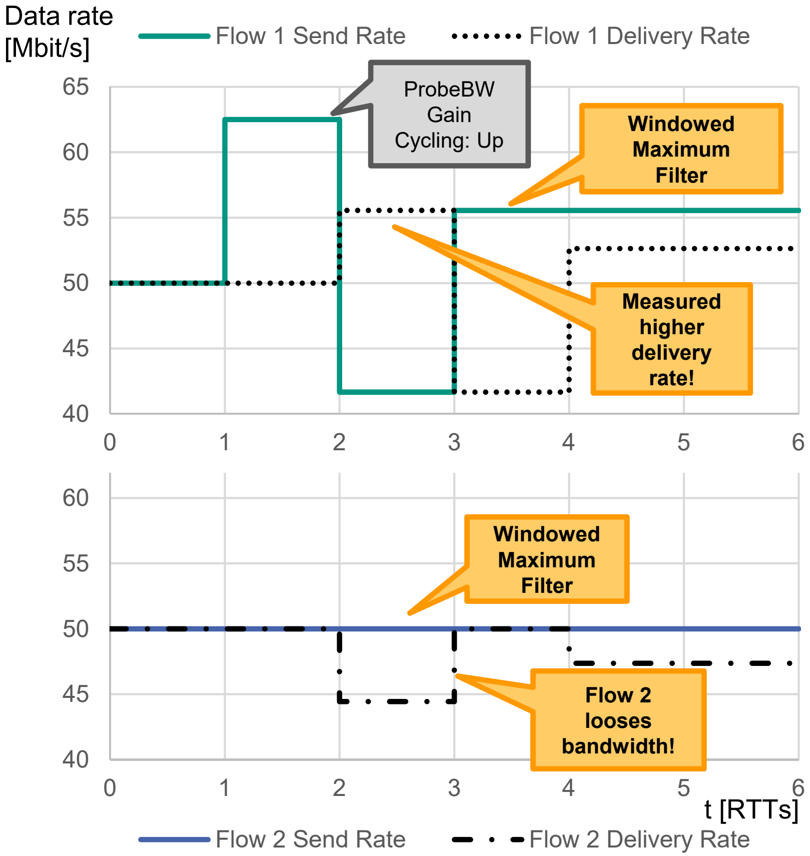

- Simplified example with two flows:
  - Bottleneck 100Mbit/s, fully utilized
  - Each flows sends with 50Mbit/s initially

- Flow probes and actually gets higher delivery rate, although bottleneck fully utilized!
- Windowed maximum filter keeps send rate too high

Behavior in the Multiple Flows Case

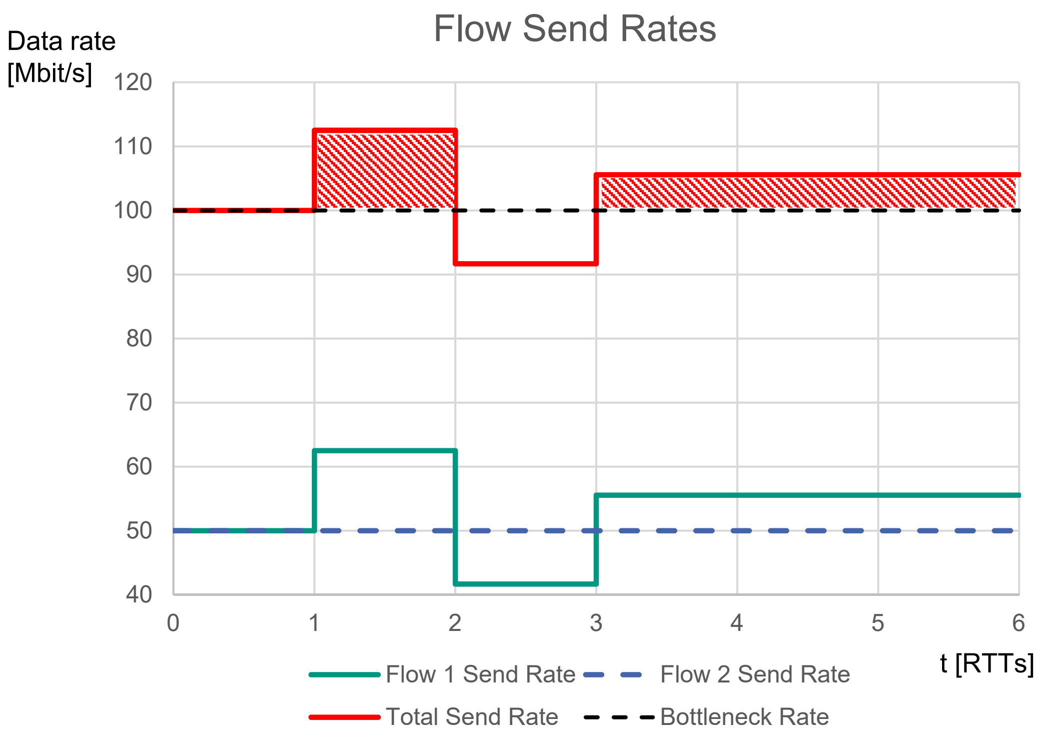

- Rate-based approach: amount of inflight data steadily increases
- Bottleneck becomes overloaded

Large buffer (>= bdp)

- BBR operates at its inflight cap (1 to 1.5bdp queued!)

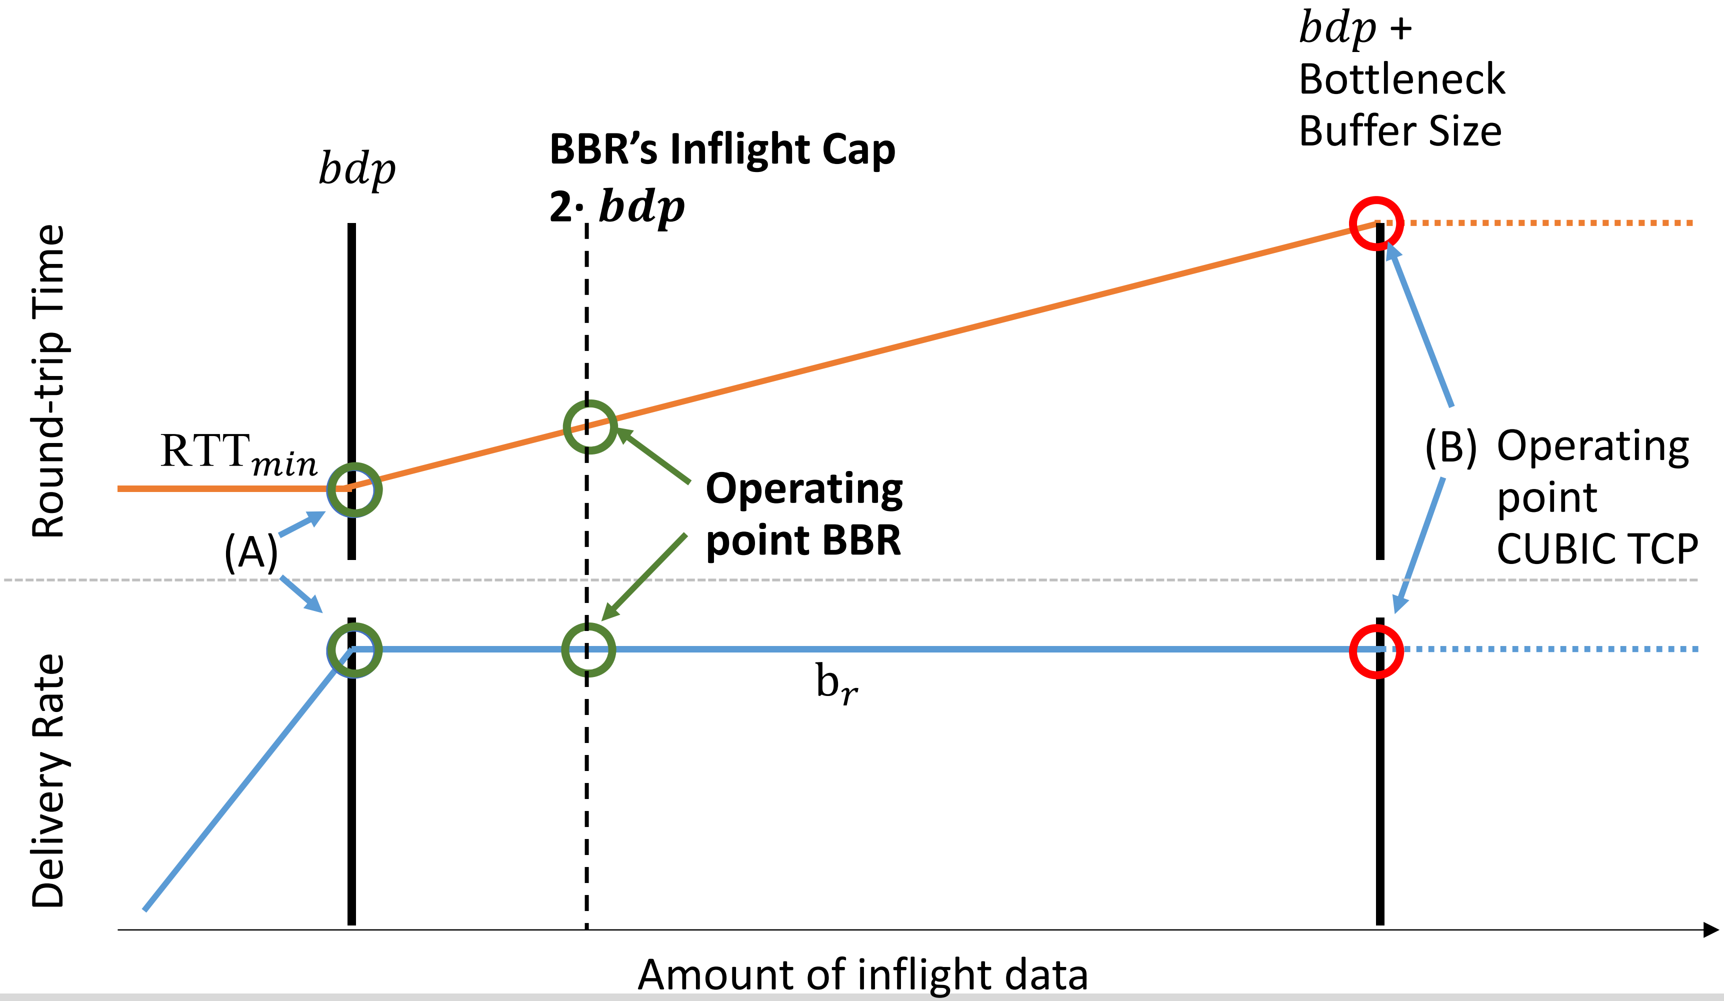

Small buffer (< bdp)

- BBR ignores packet loss as congestion signal

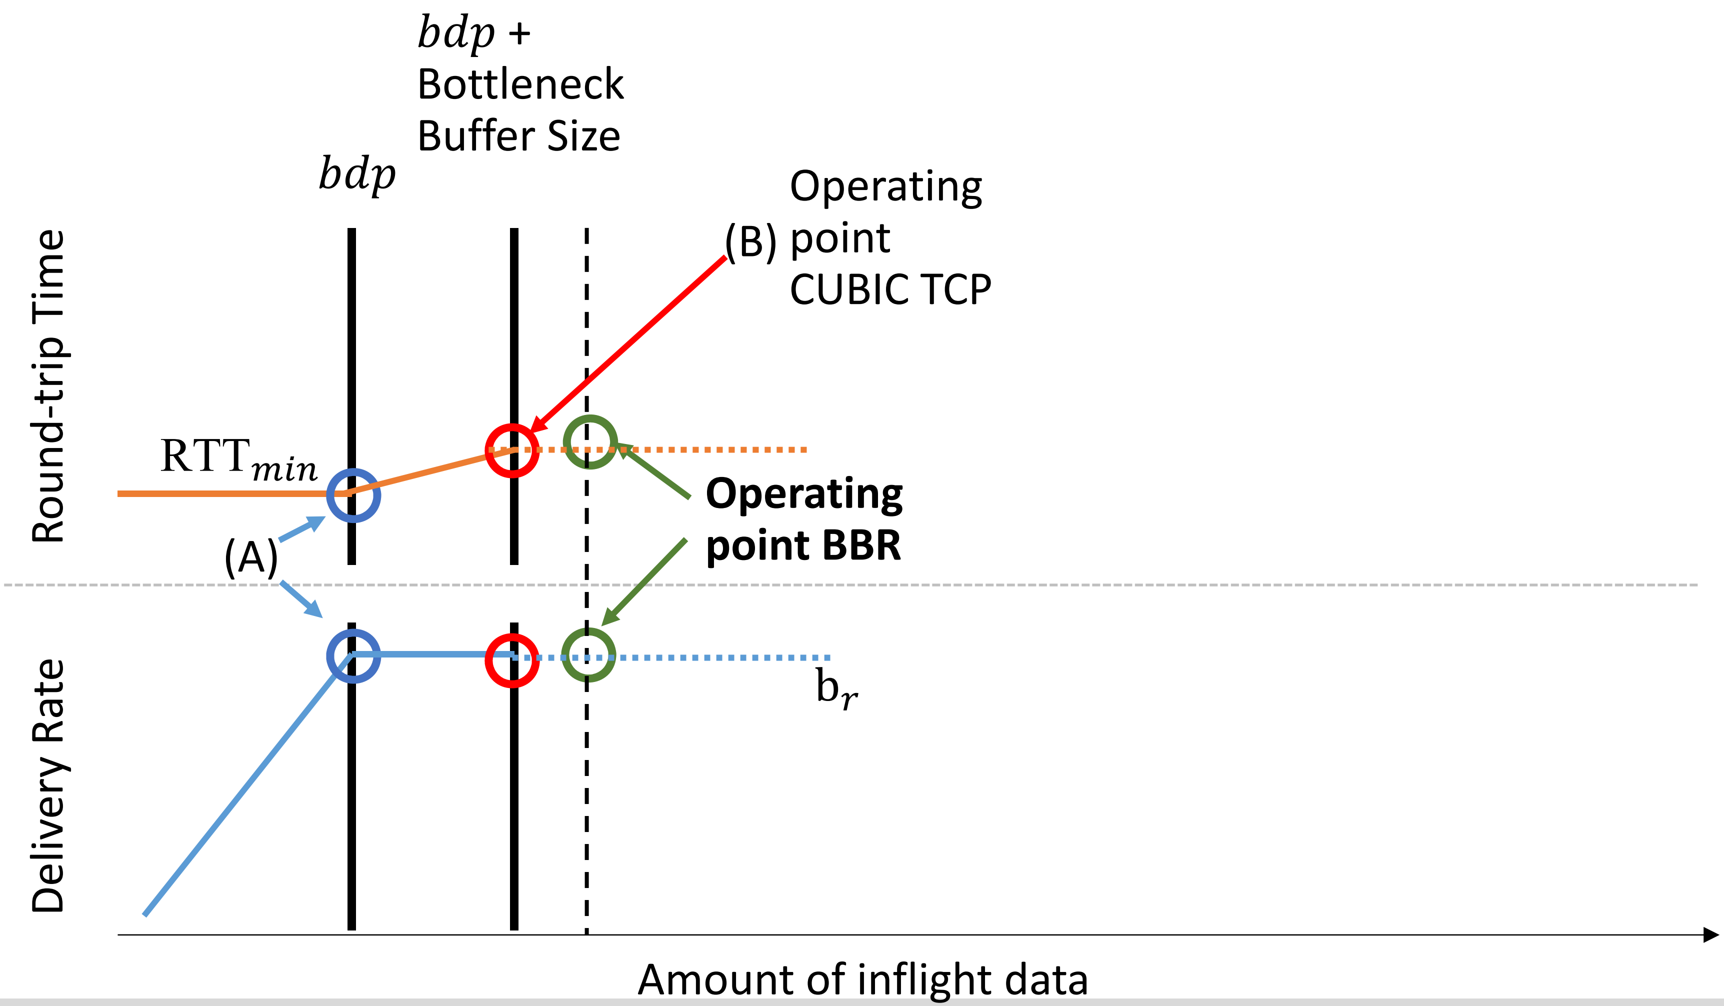

Experimental Evaluation – Setup

- Several experiments with BBR (Linux v4.9) at **1 Gbit/s** and **10 Gbit/s**

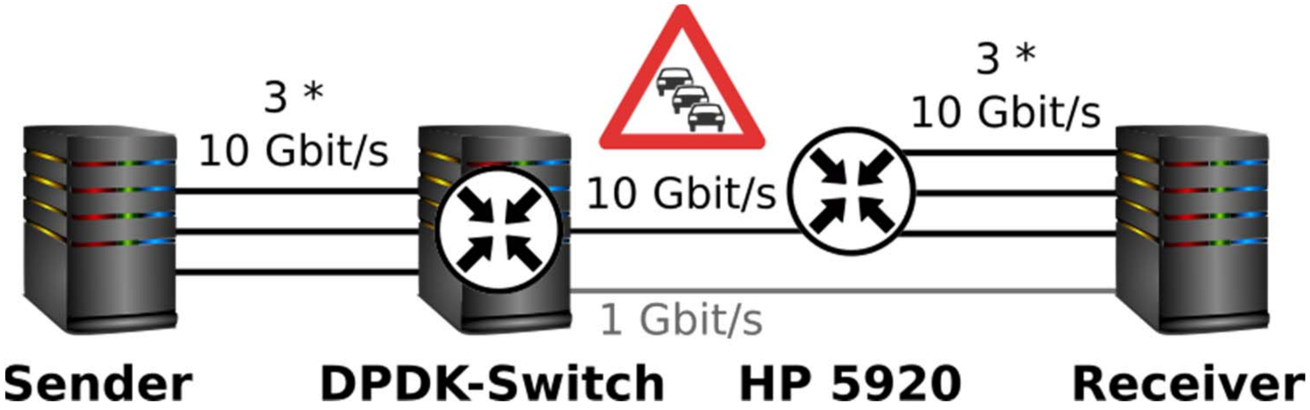

- RTT: 20ms
- Bottleneck buffers
  - Large: 160ms ($= 8\ bdp$)
  - Small: 16ms ($= 0.8\ bdp$)
- Sender is not application-limited (iperf3)
- Repeated every experiment 5 times

Different RTTs, Two Flows – Large buffer

- 2 Flows, same RTT_{min}: RTT is doubled → BBR queues $1\ bdp$

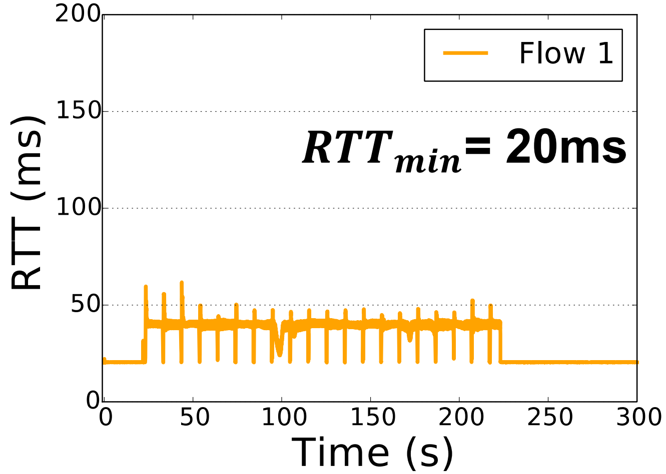

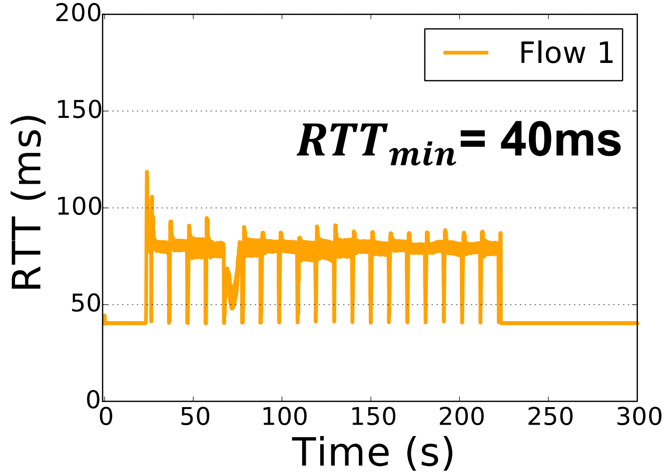

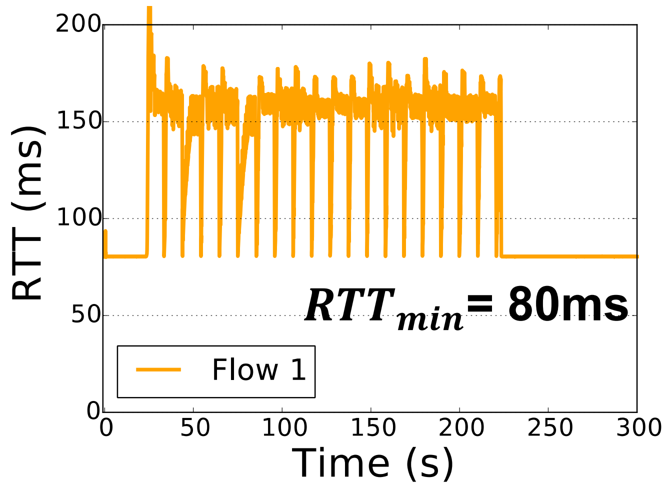

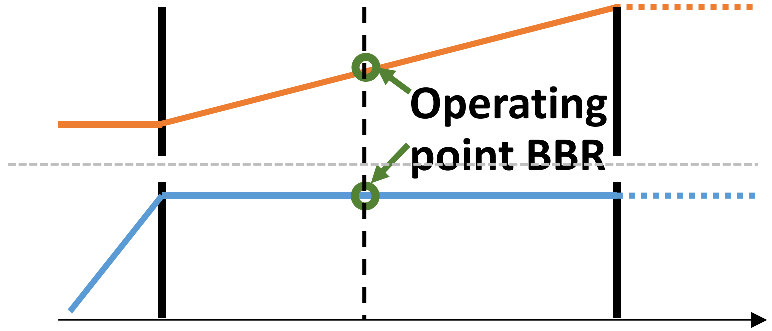

Multiple Flows and Small Buffer (0.8 BDP)

- 6 BBR flows (2 per interface)
- BBR causes massive packet loss

Sender‘s transmission rate

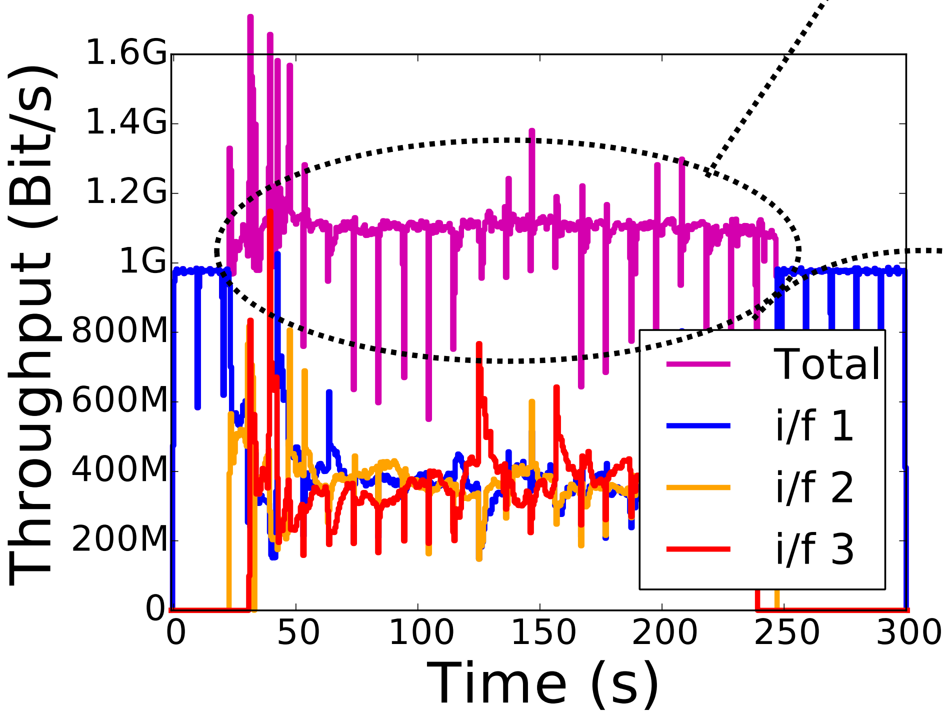

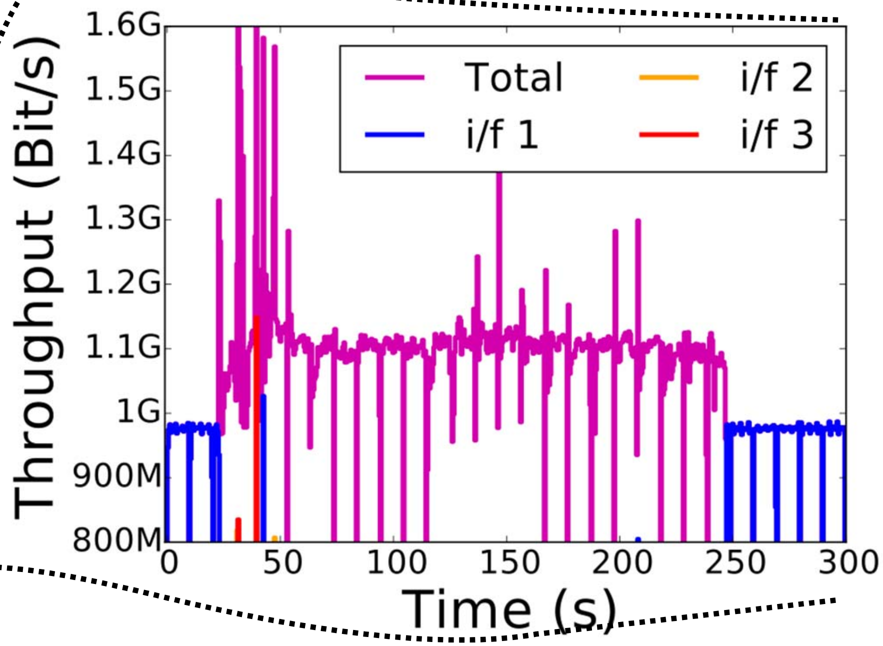

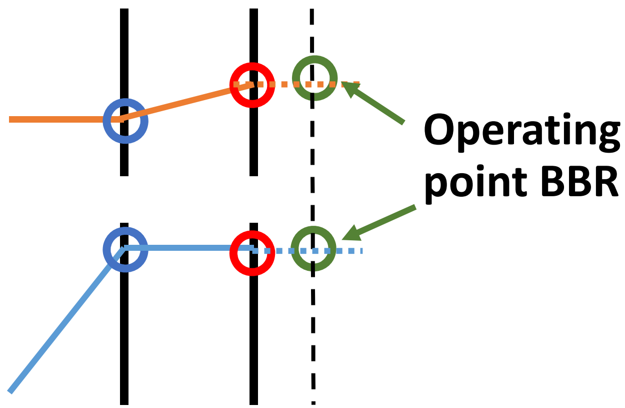

Comparison to CUBIC

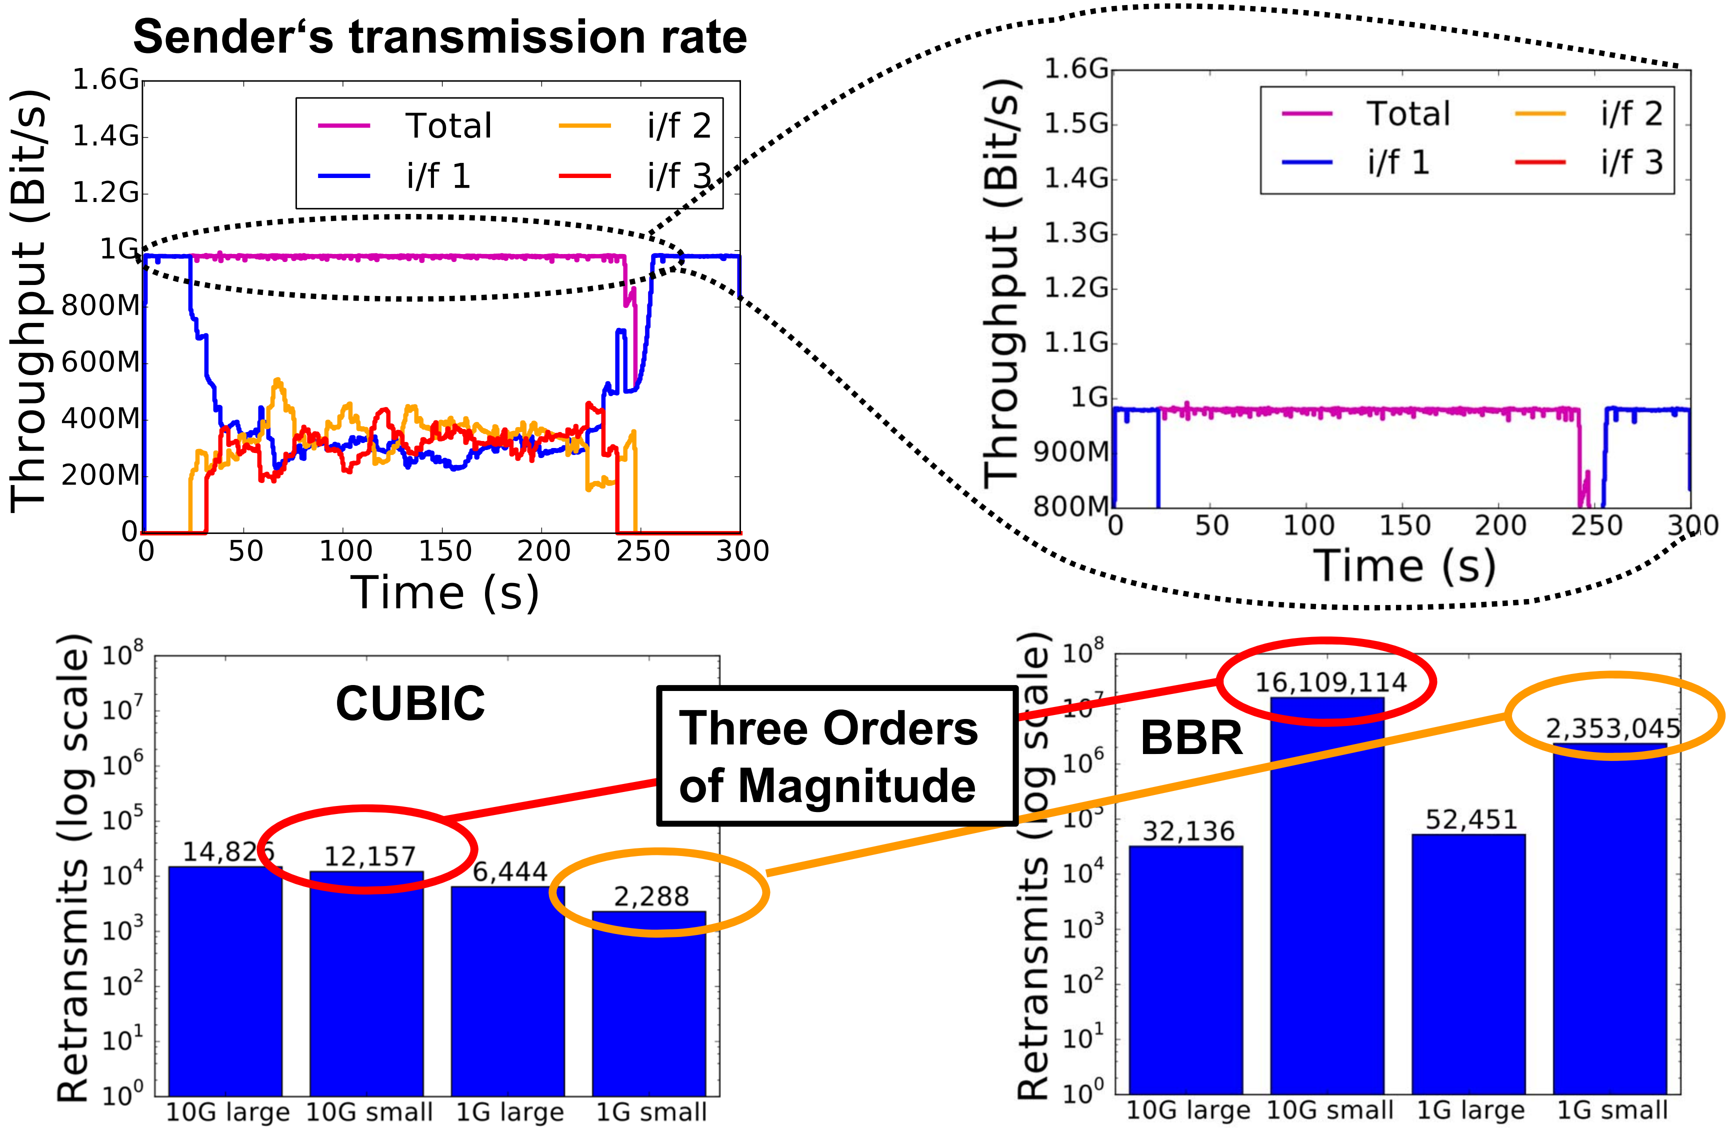

Inter-Protocol Fairness – BBR vs. CUBIC

- 1 Gbit/s, 1 BBR flow vs. 1 CUBIC flow
- Small buffers: BBR suppresses loss-based congestion control
- Single BBR flow works as intended

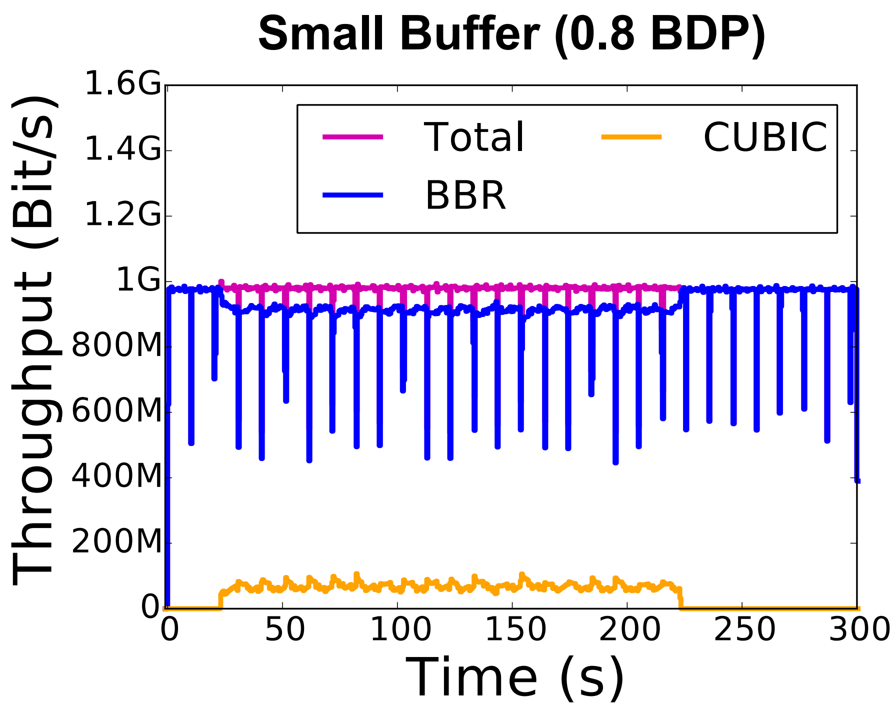

2 BBR vs. 2 CUBIC Flows

- Model mismatch: multiple BBR flows behave more aggressively
- Loss-based congestion control flows get severely suppressed

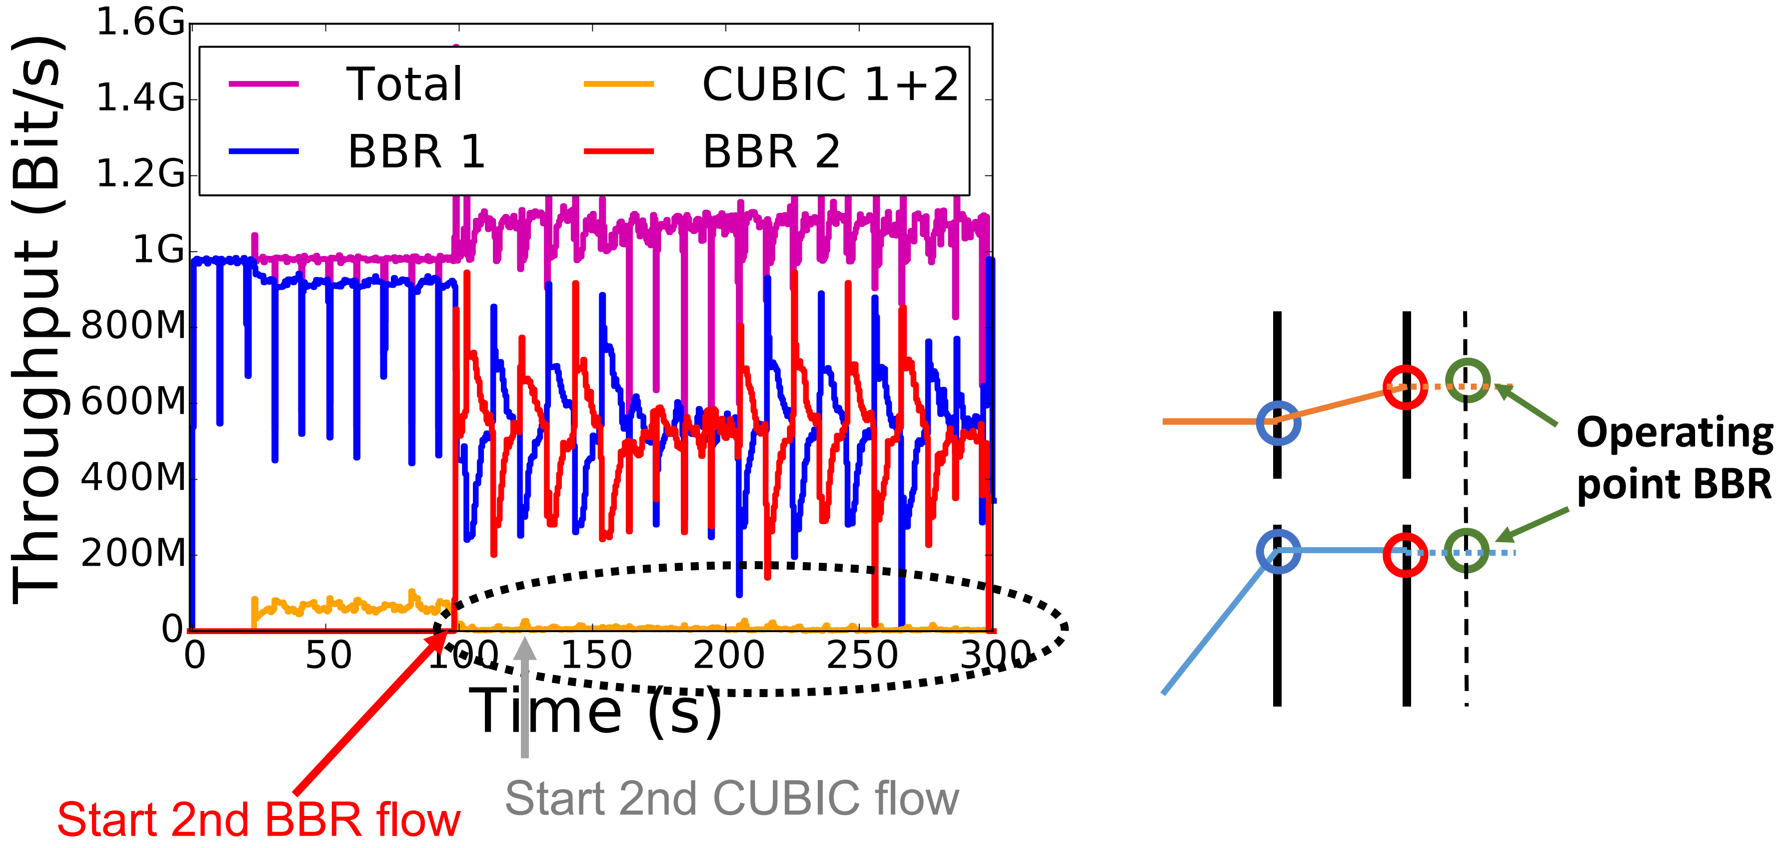

Intra-Protocol Fairness

- 6 flows (2 per interface), 20ms RTT_{min}
- No consistent fairness behavior

10Gbit/s, large buffer

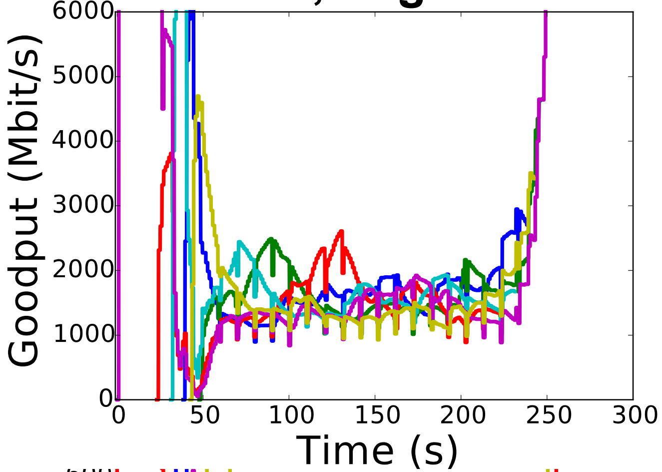

10Gbit/s, small buffer

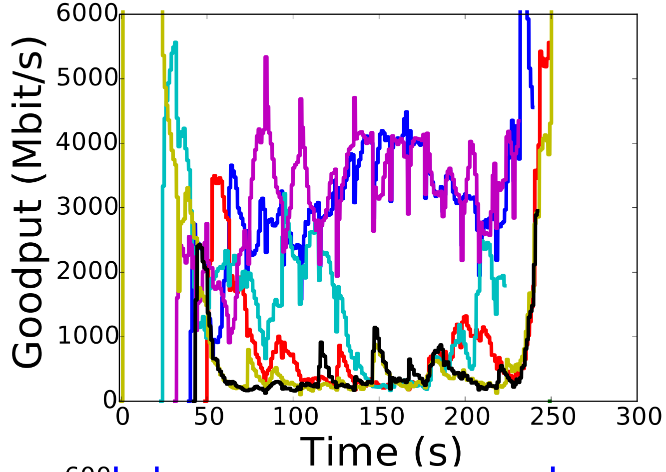

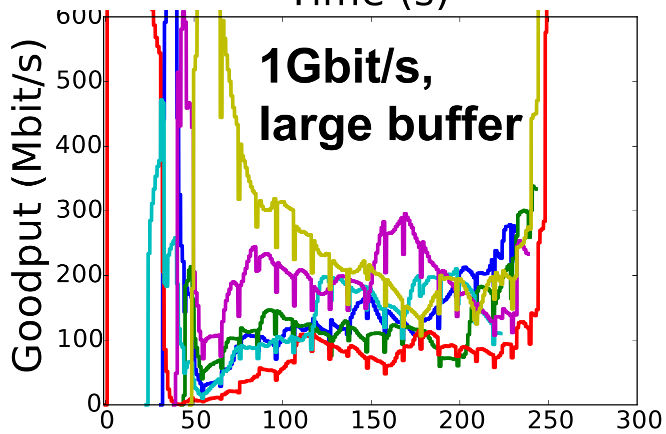

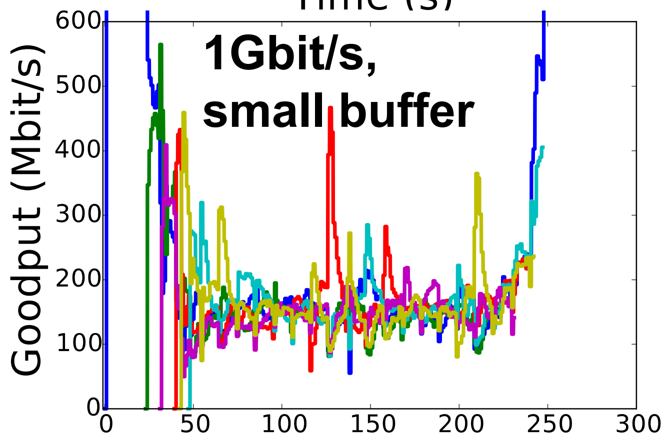

RTT Fairness

- 3 concurrent BBR flows with different RTT_{min}=20ms, 40ms, 80ms
- Each BBR flow operates at inflight cap of 2 bdp
- Larger RTT_{min} means more data inflight
  → Higher throughput at the bottleneck

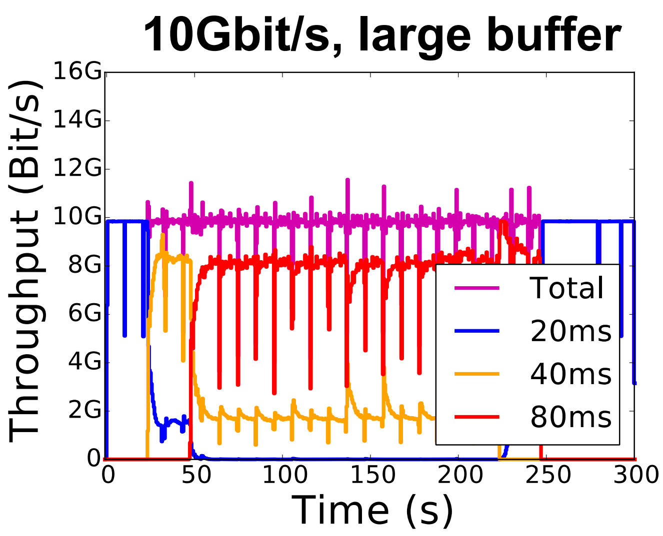

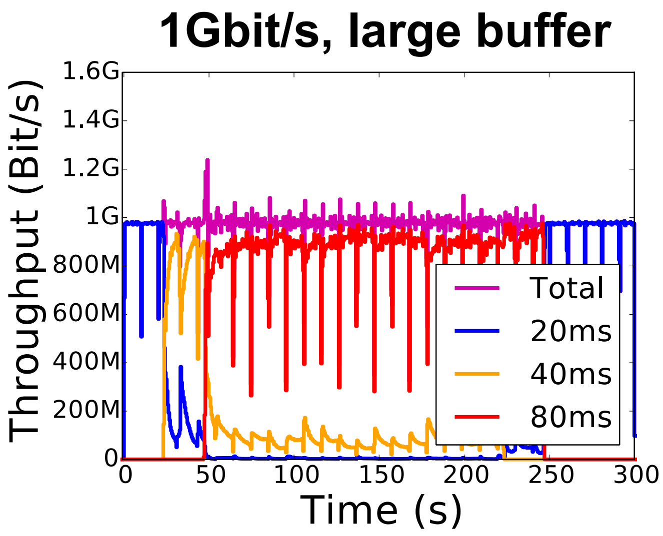

Summary

- BBR: model-based congestion control
  - Works well if no congestion present (e.g., single flow at the bottleneck)
- Multiple flows: BBR steadily increases the amount of inflight data
  - Large buffers: BBR operates at inflight cap, RTT unfairness
  - Small buffers: high amount of packet losses
- No consistent fairness behavior
- Unfairness to flows with loss-based congestion control, e.g., CUBIC
- BBR is already in use: but probably application-limited
- BBR is still under development

References

[1] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson, „BBR: Congestion-Based Congestion Control“, ACM Queue, vol. 14, no. 5, pp. 50:20–50:53, Oct. 2016.

[2] M. Hock, R. Bless, M. Zitterbart: „Experimental Evaluation of BBR Congestion Control”, Proceedings of IEEE ICNP 2017, Oct. 10–13, Toronto, Canada, http://doc.tm.kit.edu/2017-kit-icnp-bbr-authors-copy.pdf

BACKUP SLIDES

Multiple Flows – Large buffer

- 1 Gbit/s, 20ms RTT_{min}
- RTT is increased to 40ms → BBR operates at inflight cap of $2bdp$

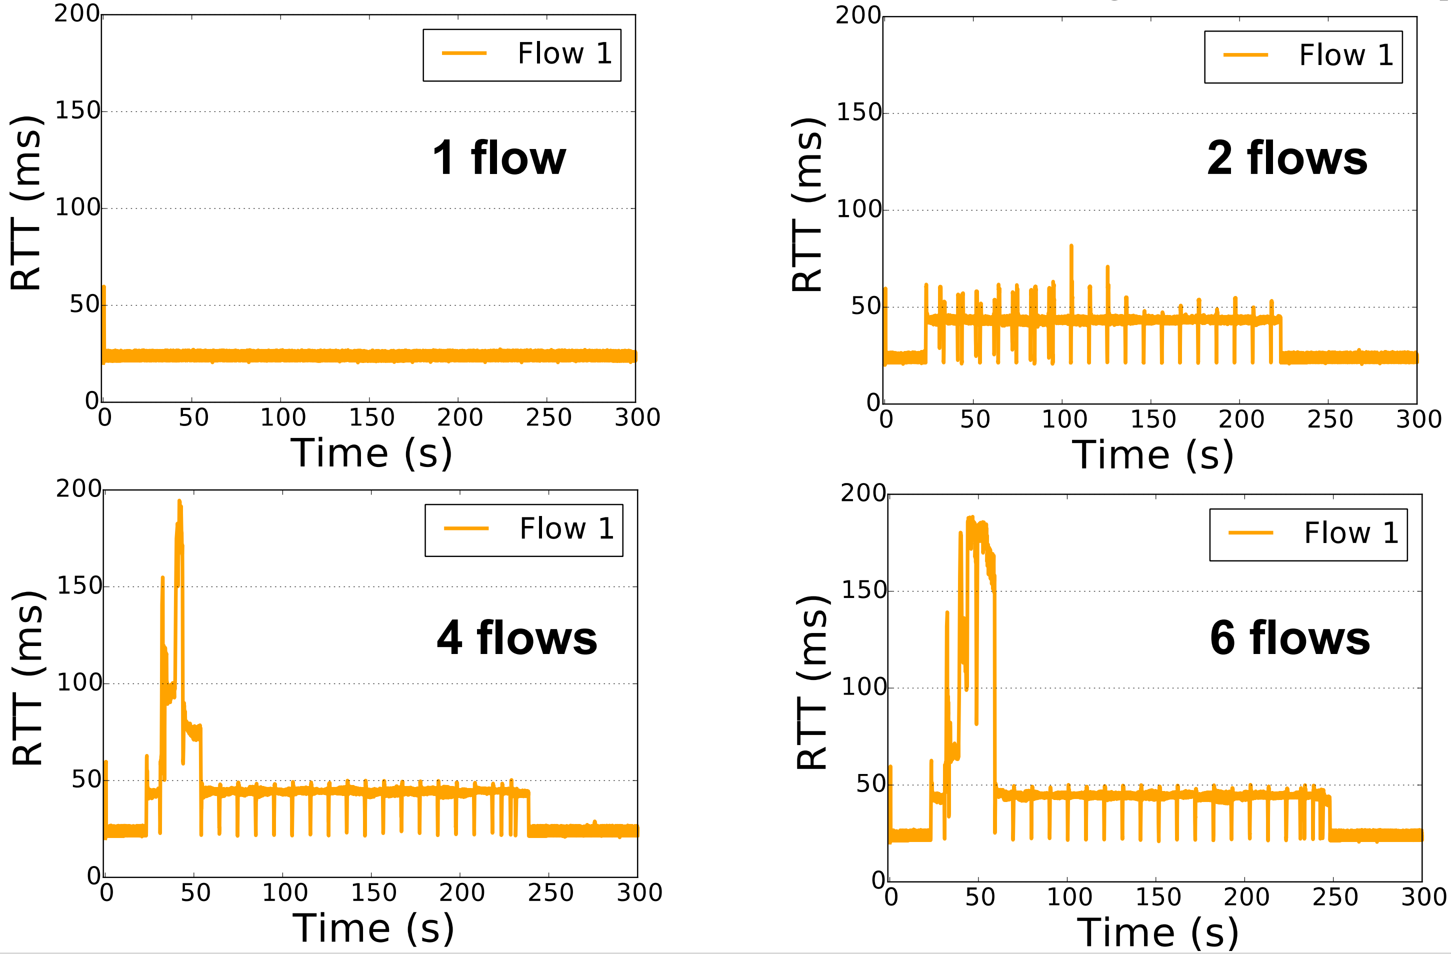

Inter-Protocol Fairness – BBR vs. CUBIC

- 1 Gbit/s, 1 BBR flow vs. 1 CUBIC flow
- Large buffers
  - BBR's inflight cap is larger due to present queuing delay
  - BBR may loose against loss-based congestion control
- Small buffers: BBR suppresses loss-based congestion control

Large Buffer

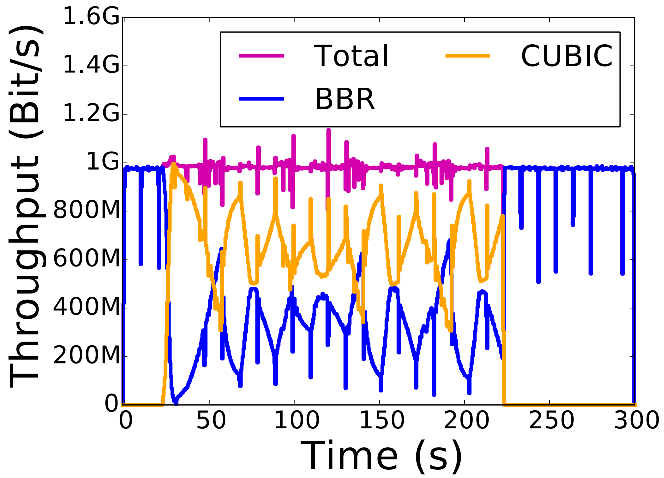

Small Buffer

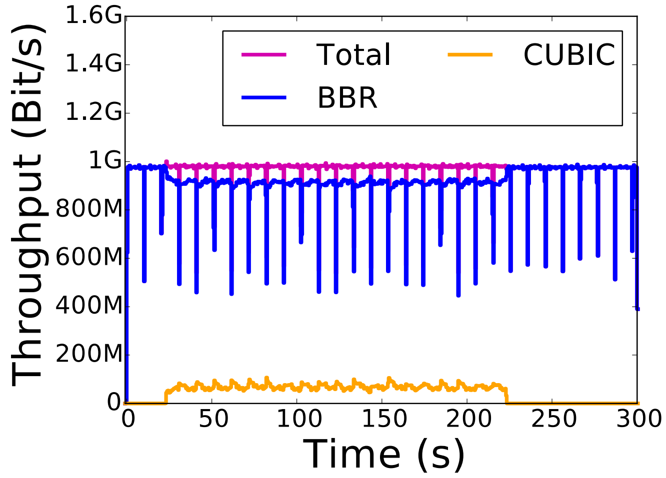

Interprotocol Fairness – BBR vs. CUBIC

- 1 BBR flow vs. 1 CUBIC flow

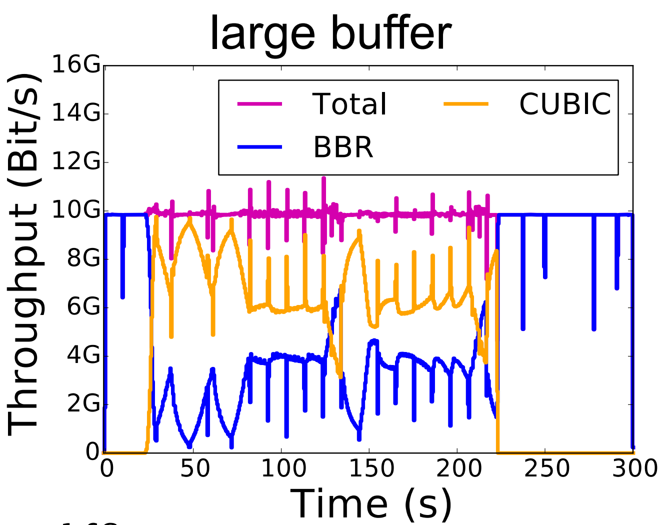

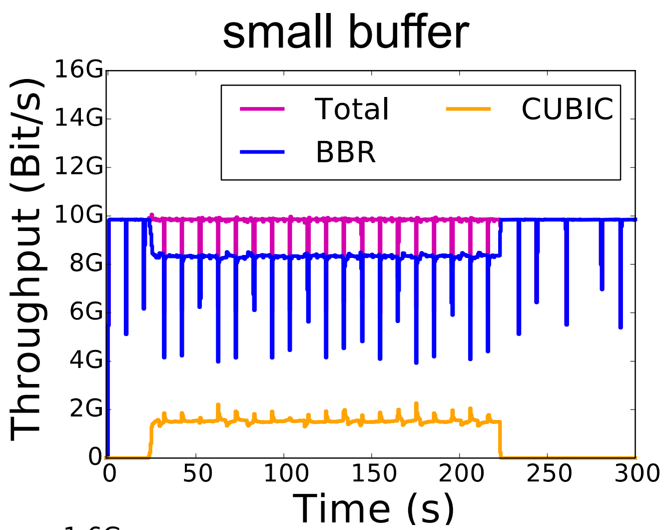

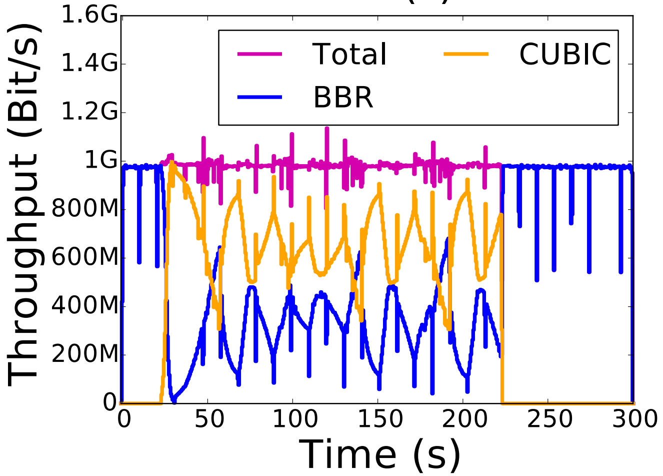

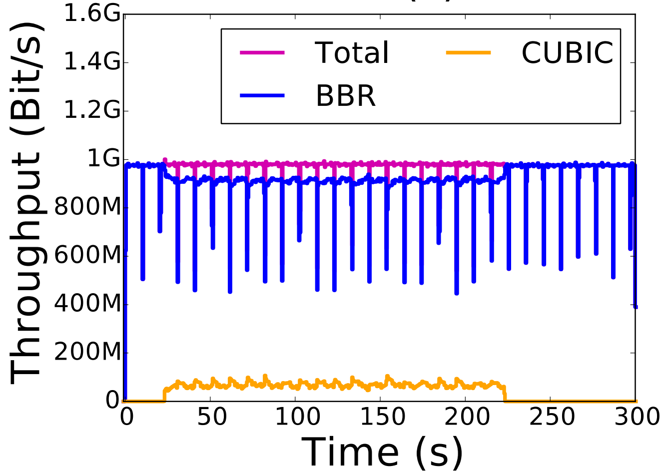

Packet Loss BBR – Outgoing data at sender

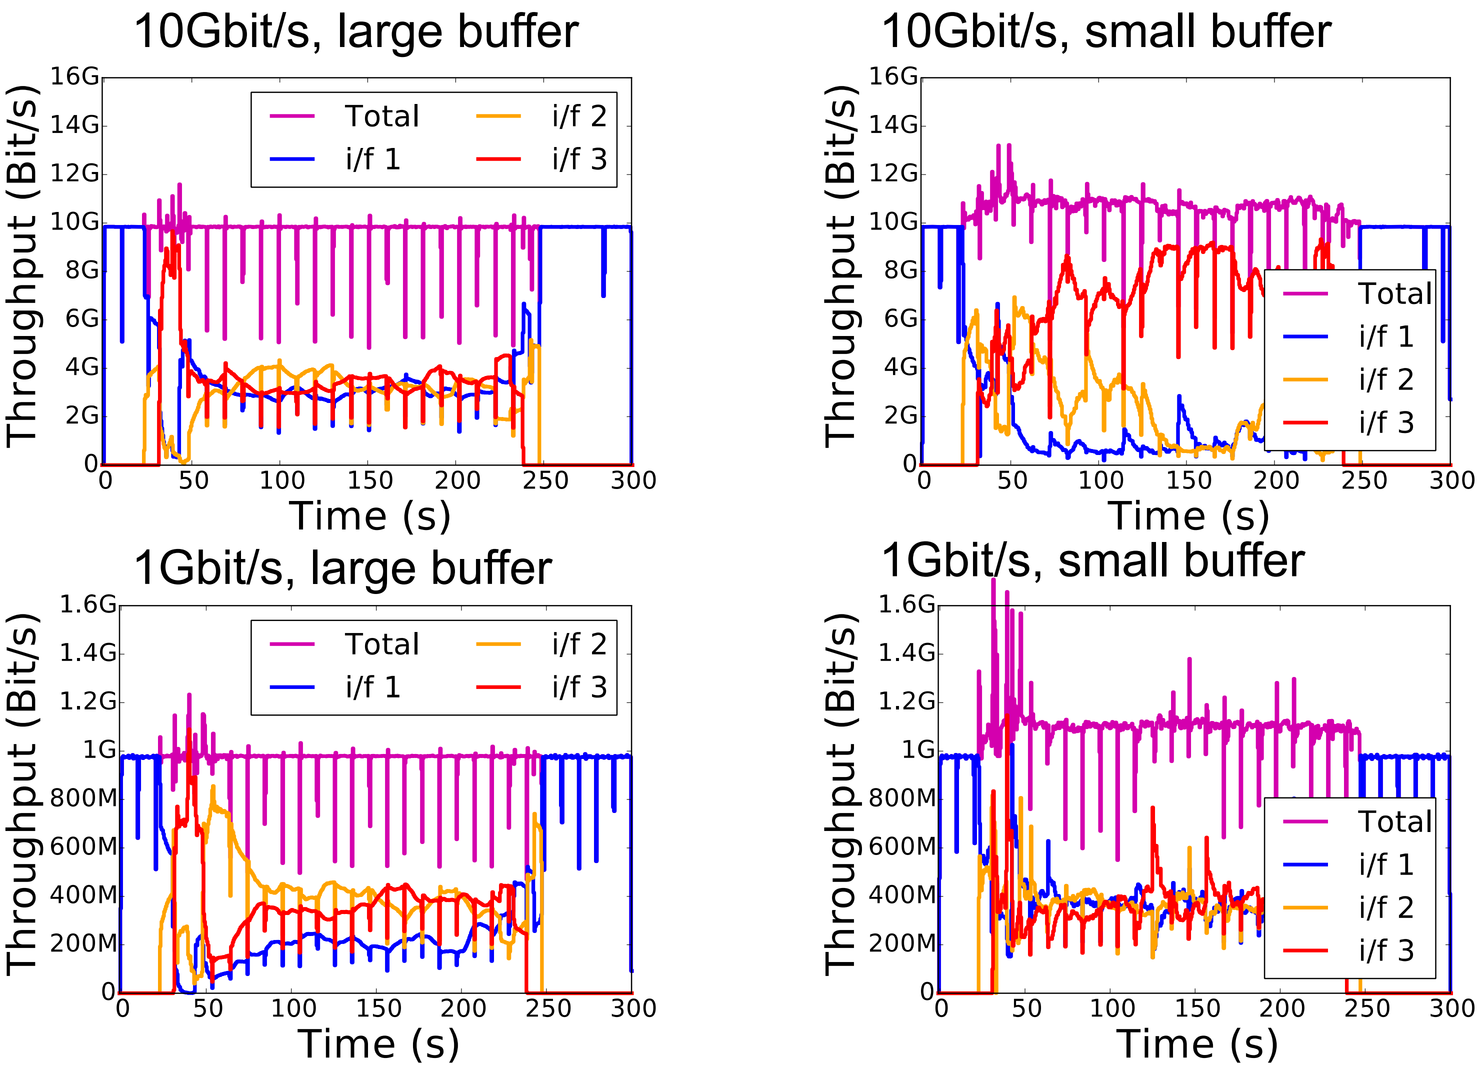

Packet Loss CUBIC – Outgoing data at sender