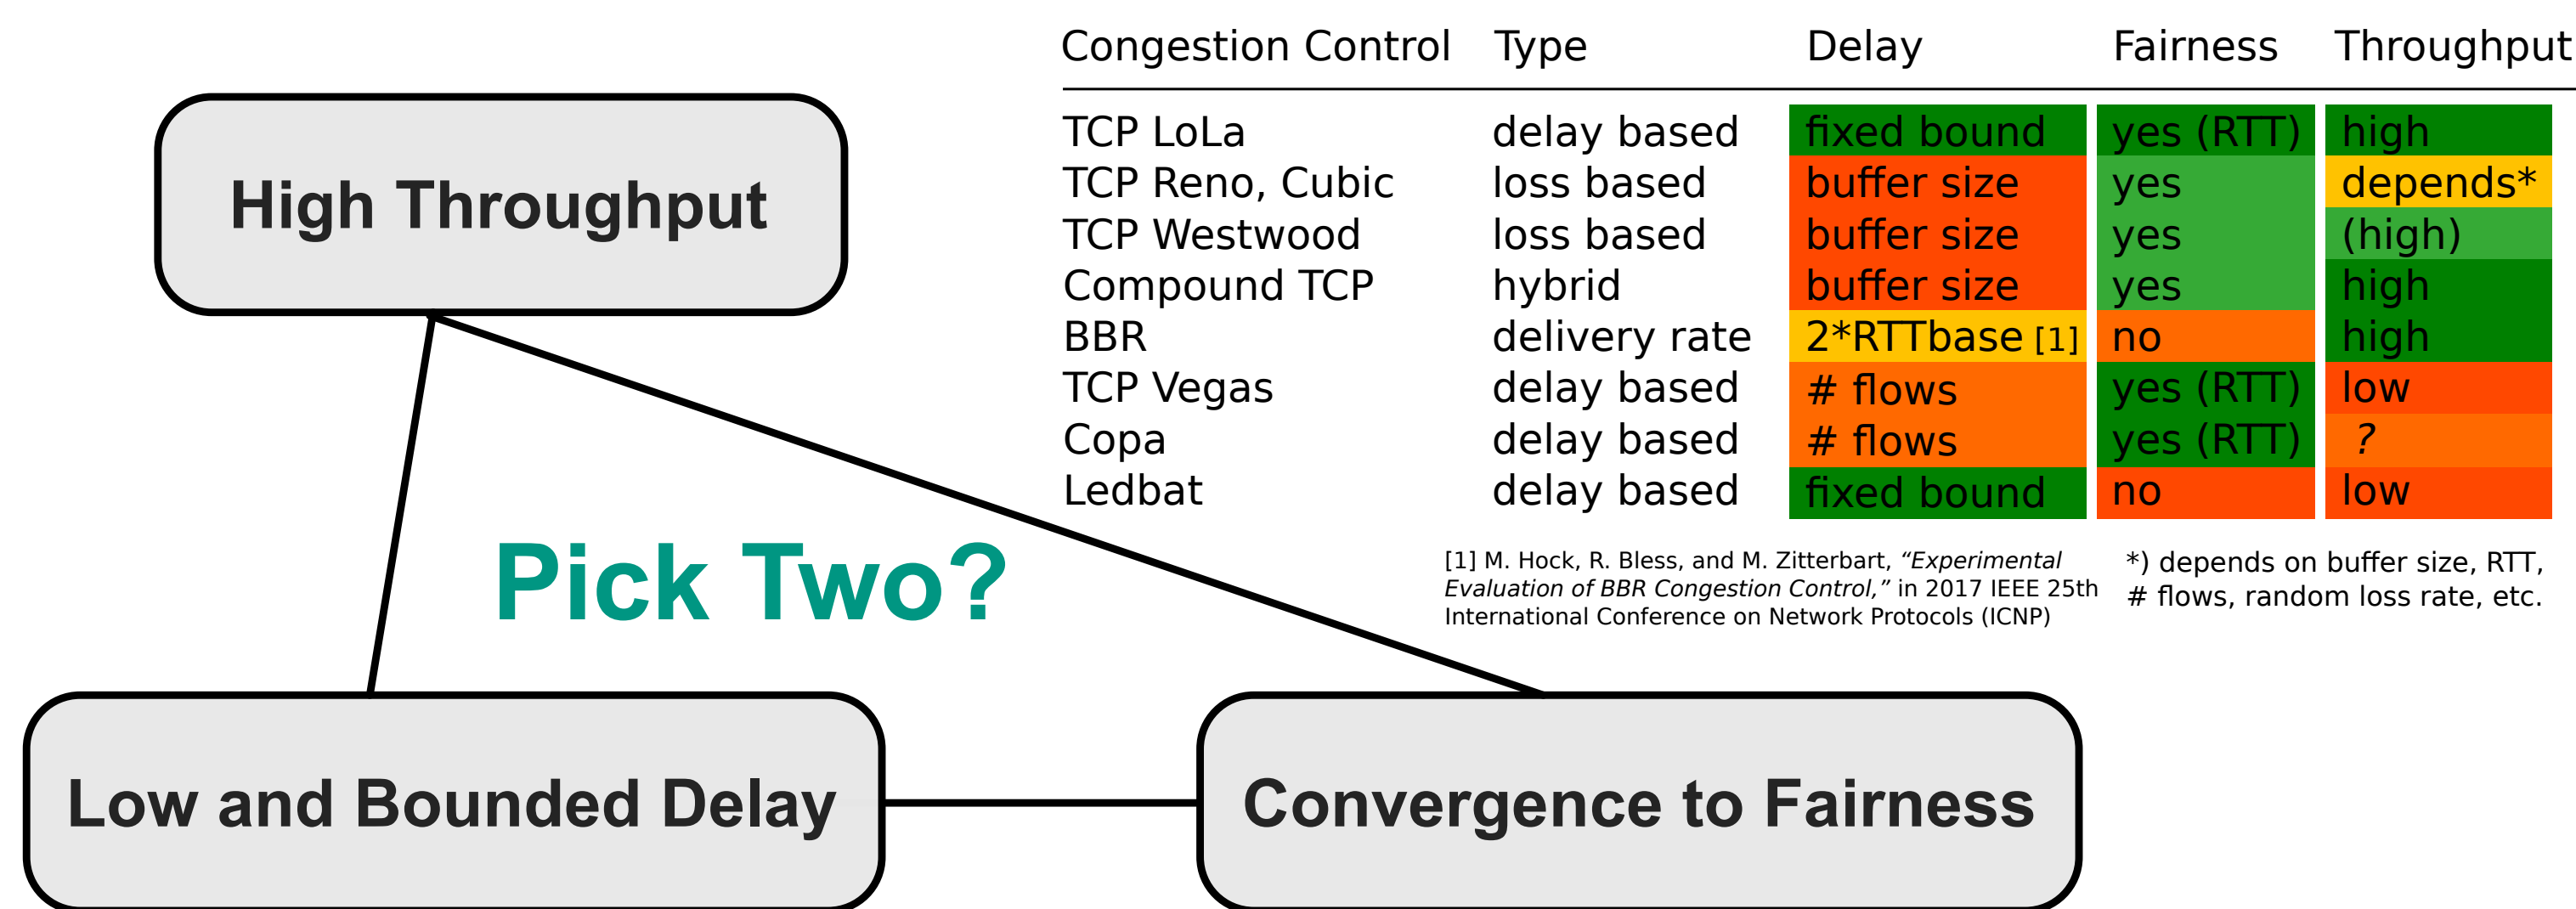


FFBquick: Fast Convergence to Fairness for Delay-bounded Congestion Controls

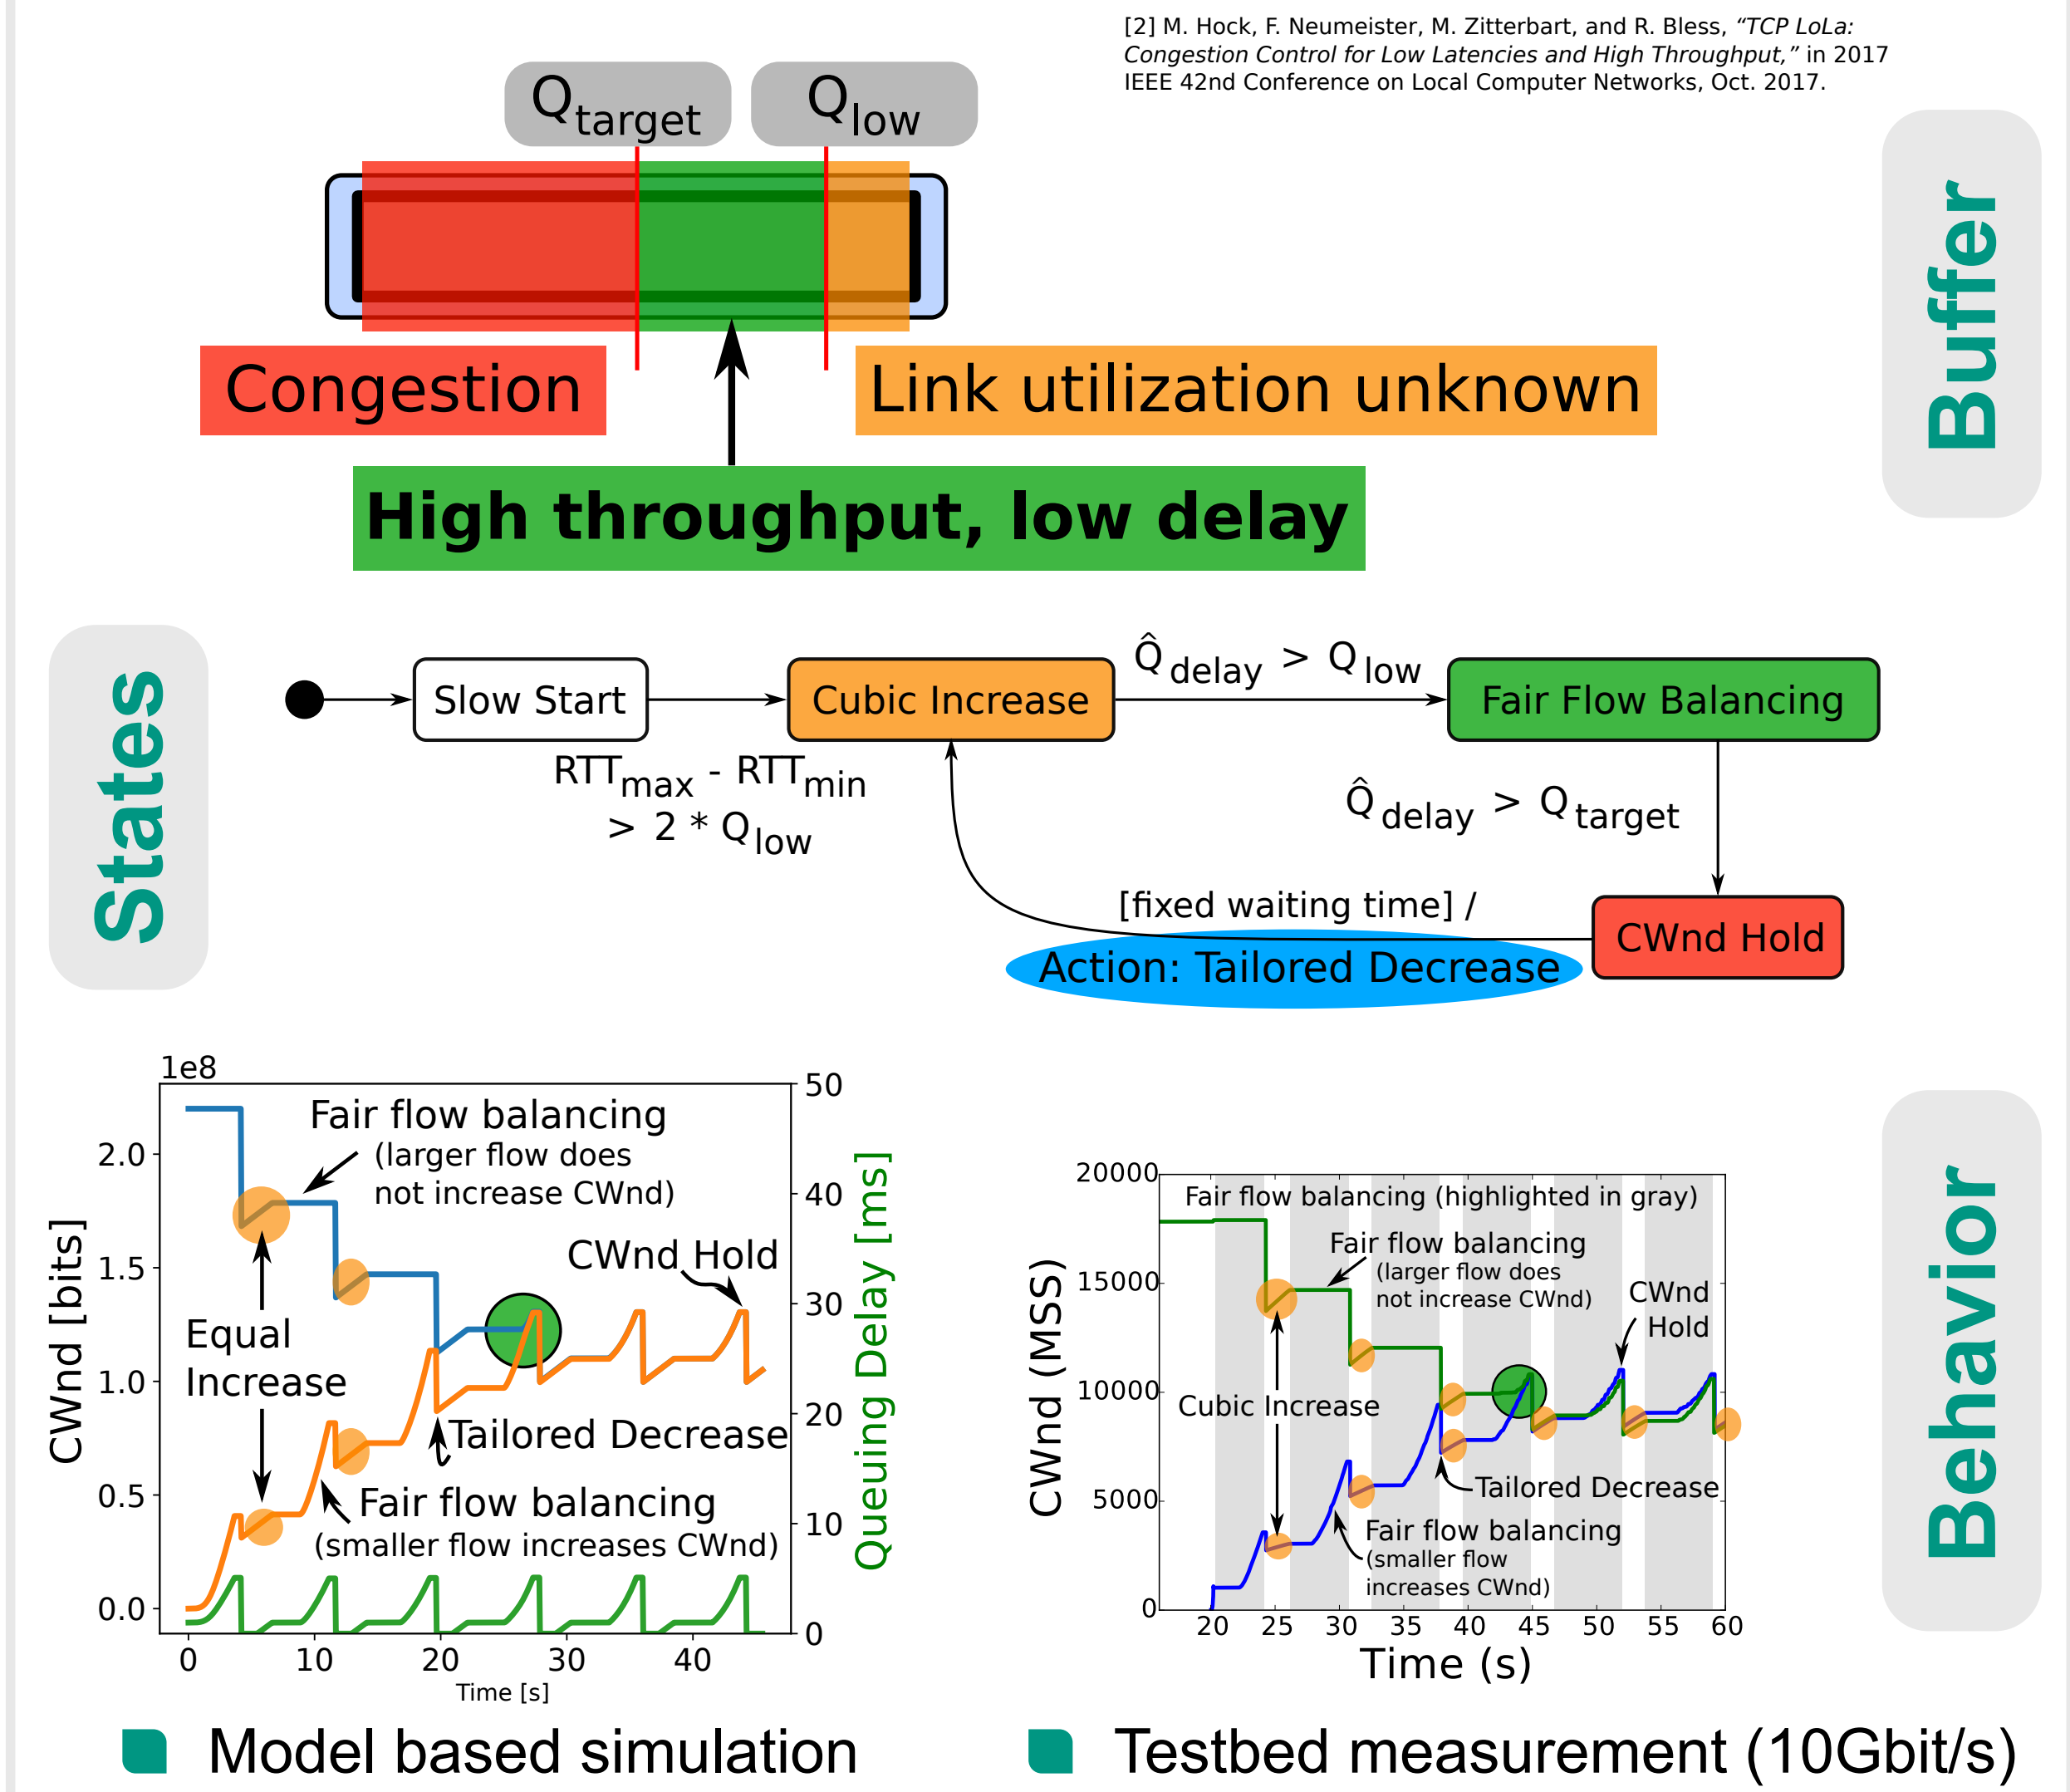
Mario Hock, Roland Bless, Felix Neumeister, Martina Zitterbart

Challenge



- How is a convergence to fairness typically achieved?
 - Fixed amount of packets allowed in buffer per flow
 - AIMD or similar dynamics
 - Additive Increase, Multiplicative Decrease
- What's new in Fair Flow Balancing?
 - Equal amount of packets allowed in buffer per flow
 - Dynamically scaled to maintain bounded delay!

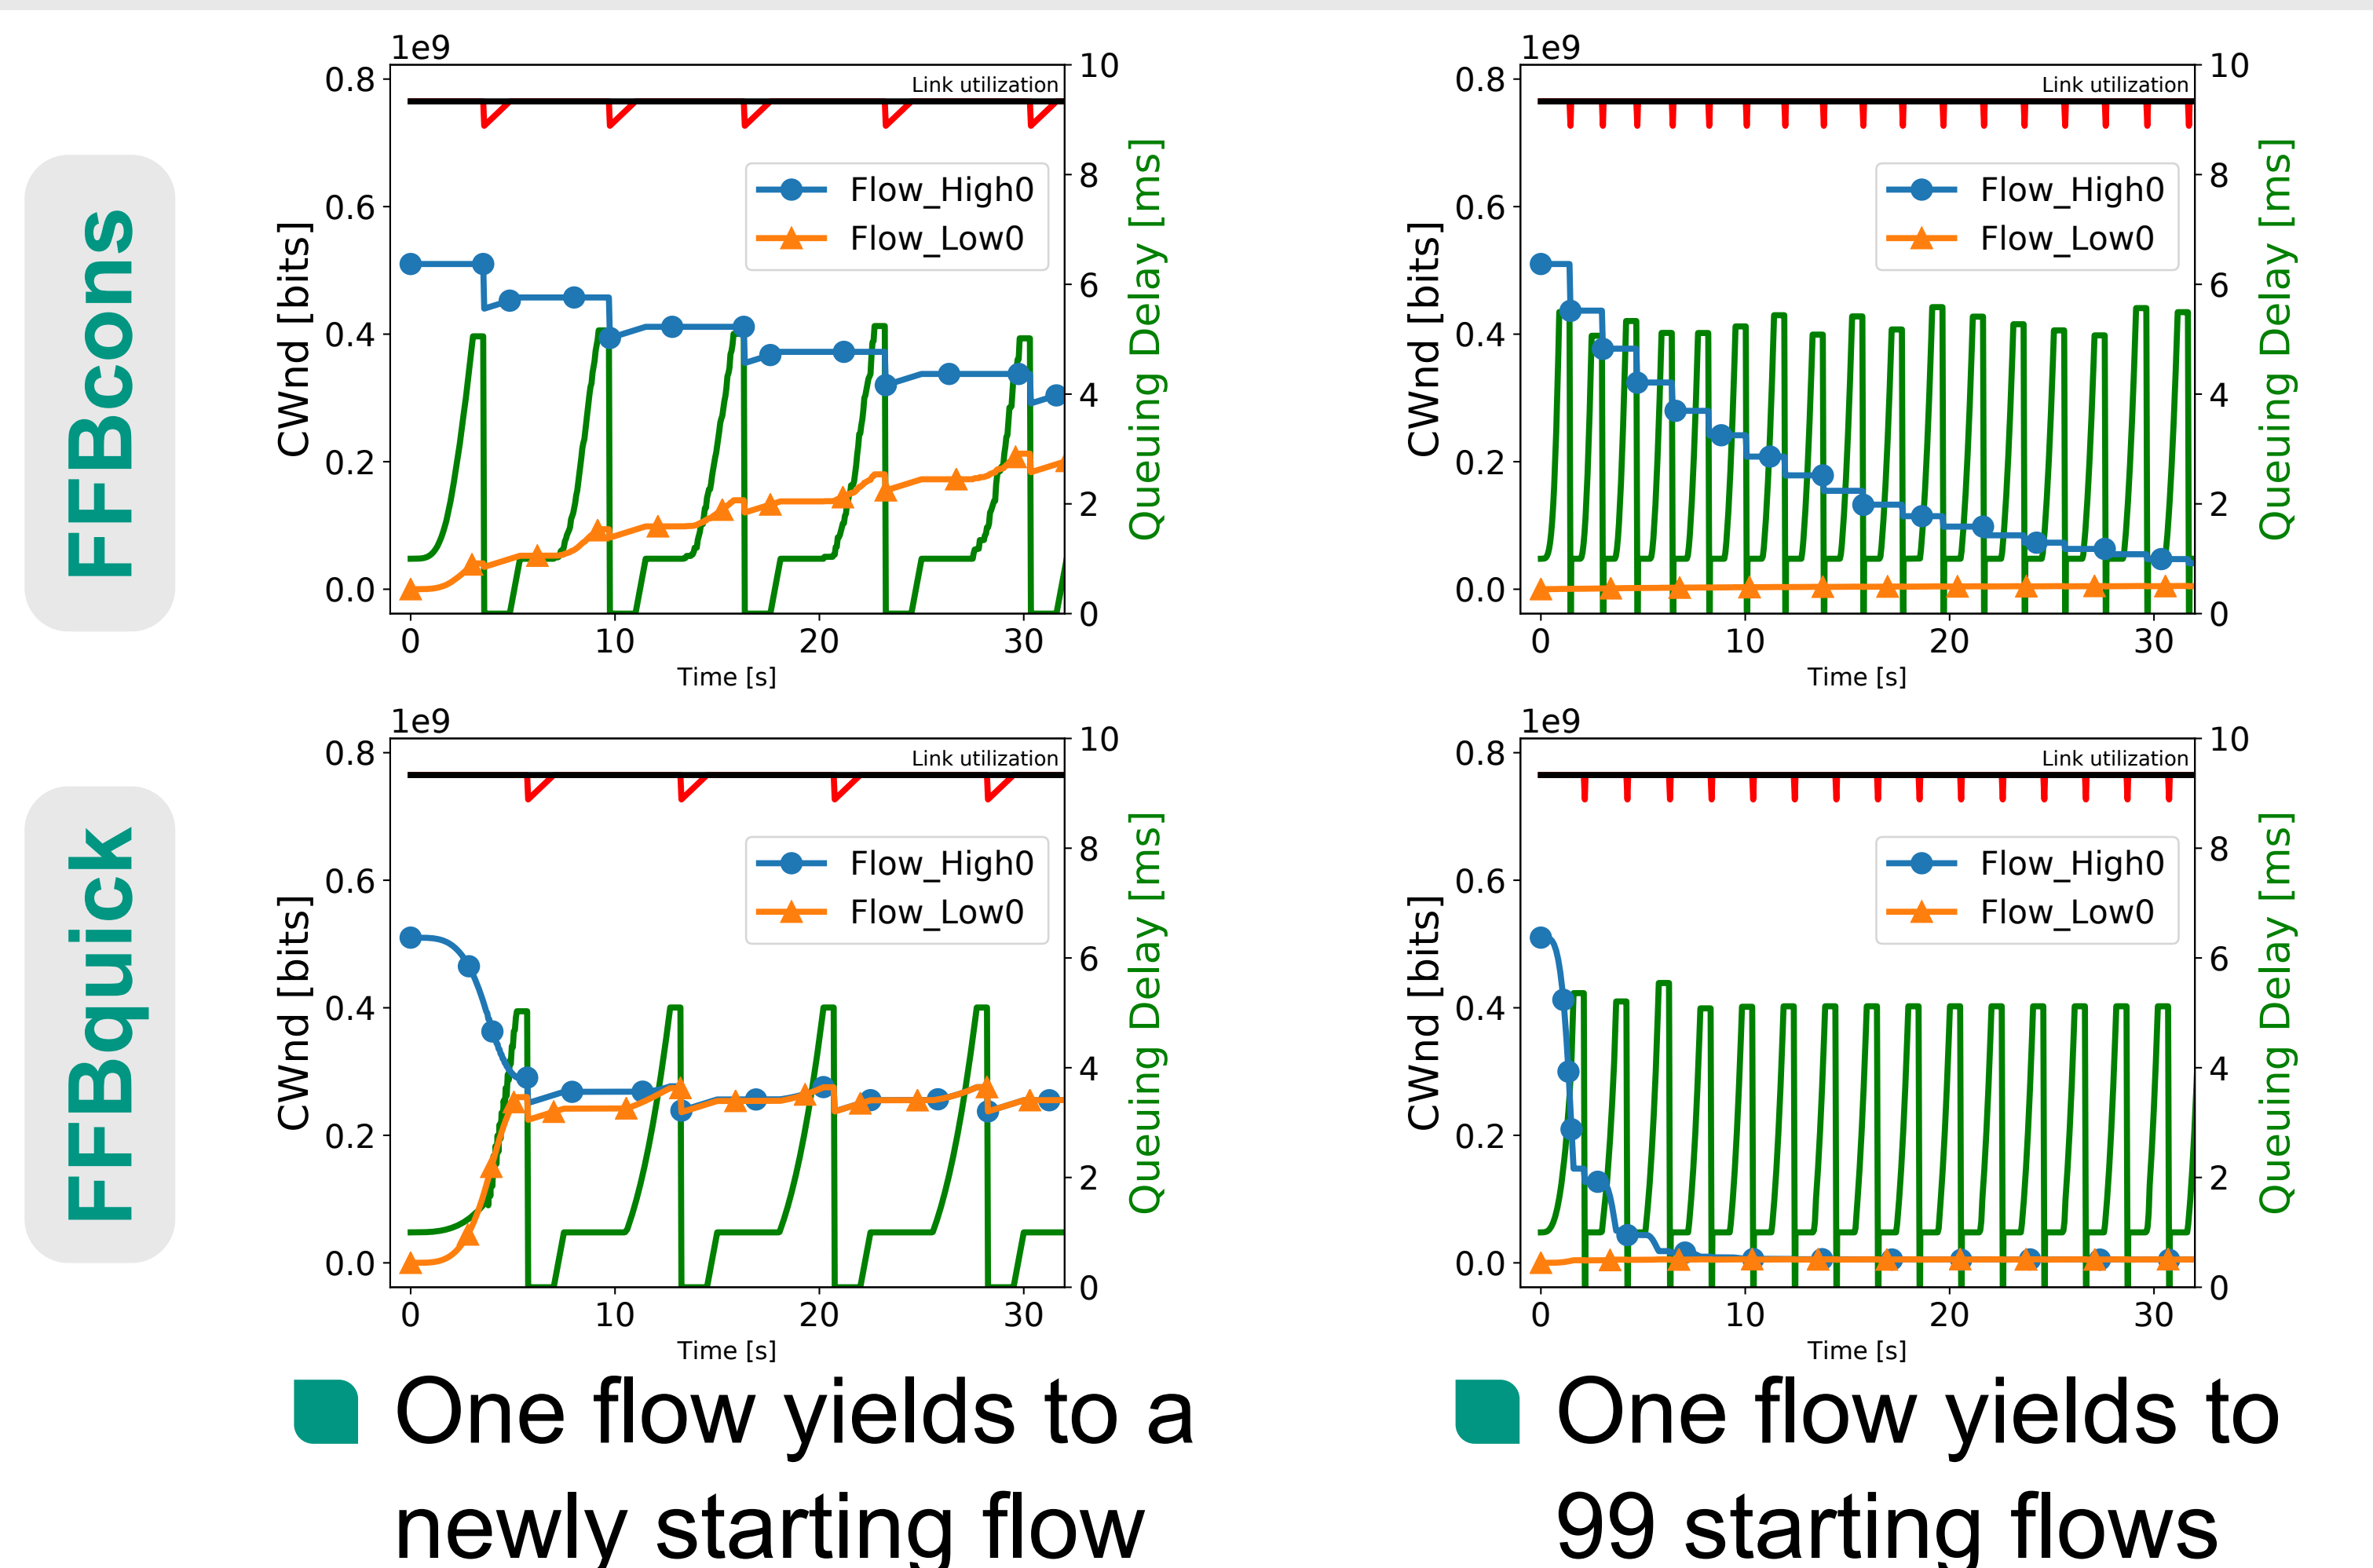
TCP LoLa [2]



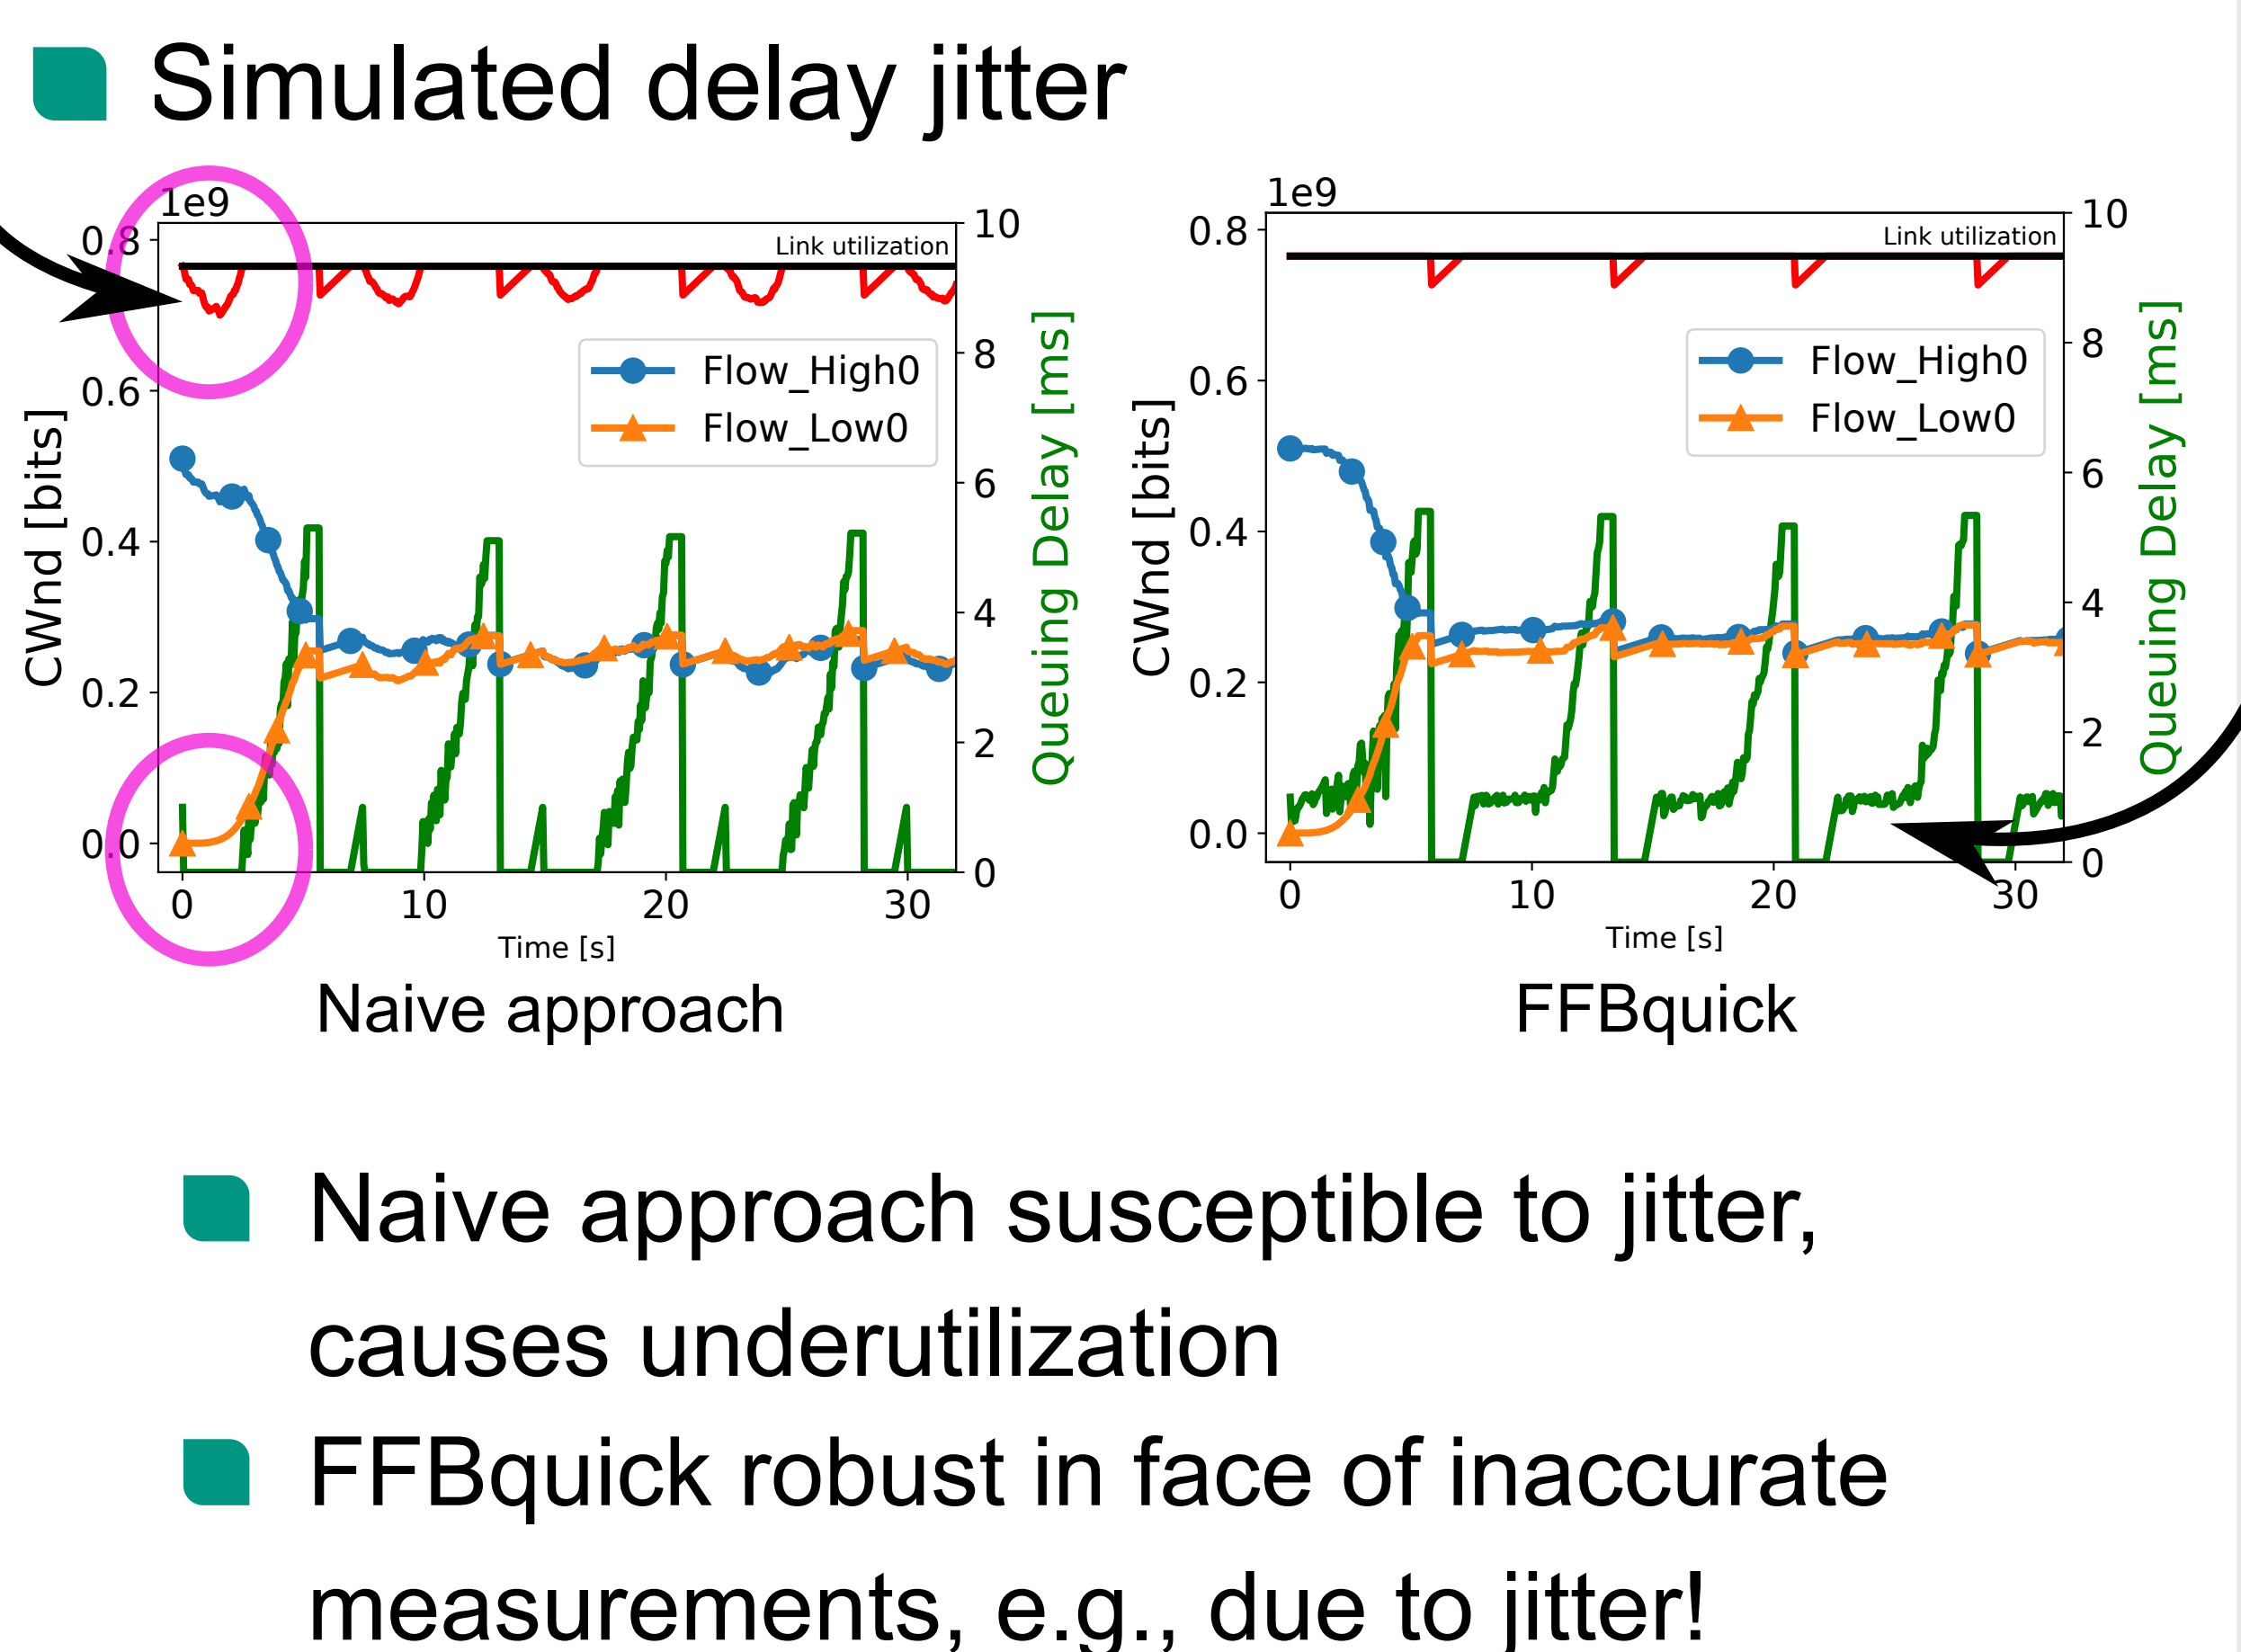
Evolution – Why is it done like this?

- $X(t) = (t \text{ [ms]} / \sigma)^3$
 - Allowed data in bottleneck queue
 - Q_{data}
 - Estimated amount of data in queue
 - $Q_{delay} \cdot \frac{CWnd}{RTT}$
 - $CWnd += X(t) - Q_{data}$
- Naive approach**
- Quick, but...
 - No safety margin
- FFBcons**
- Never decrease CWnd
 - Q_{data} "explodes"!
- FFBquick**
- Keep Q_{data} up
 - Quick convergence!

Significantly Improved Convergence Speed



Robust against Jitter / Inaccuracies

- Simulated delay jitter
- 
- Naive approach**
- FFBquick**
- Naive approach susceptible to jitter, causes underutilization
 - FFBquick robust in face of inaccurate measurements, e.g., due to jitter!