SWINBURNE UNIVERSITY OF TECHNOLOGY

Multimedia-unfriendly TCP Congestion Control and Home Gateway Queue Management

Lawrence Stewart[α], David Hayes[α], Grenville Armitage[α], Michael Welzl[β], Andreas Petlund[β]

[α] Centre for Advanced Internet Architectures (CAIA), Swinburne University of Technology

[β] Department of Informatics, University of Oslo

Outline

1. Introduction
2. Results
3. Conclusion

Detailed outline (section 1 of 3)

1. Introduction
2. Results
3. Conclusion

1. Introduction
   - Motivation
   - Not All Queues Are Created Equal
   - TCP CC (Congestion Control) Recap
   - Methodology

Motivation

- Gateway devices often have excessive drop-tail queue buffering[1,2]
- TCP CC algorithms each interact differently with network buffers
- Real-time multimedia traffic sharing the gateway bottleneck experiences collateral damage: latency/jitter and loss
- How does choice of CC algorithm and queue implementation impact the inflicted collateral damage?

[1] M. Claypool, R. Kinicki, M. Li, J. Nichols and H. Wu, “Inferring Queue Sizes in Access Networks by Active Measurement”, PAM Workshop '04

[2] M. Dischinger, A. Haeberlen, K. P. Gummadi and S. Saroiu, “Characterizing residential broadband networks”, SIGCOMM IMC '07

Not All Queues Are Created Equal

- When is a drop-tail queue full?

Packet/Slot based	drop if occupancy ≥ qlen
Fixed-byte strict	drop if occupancy + pkt ≥ qlen
Fixed-byte loose	drop if occupancy > qlen

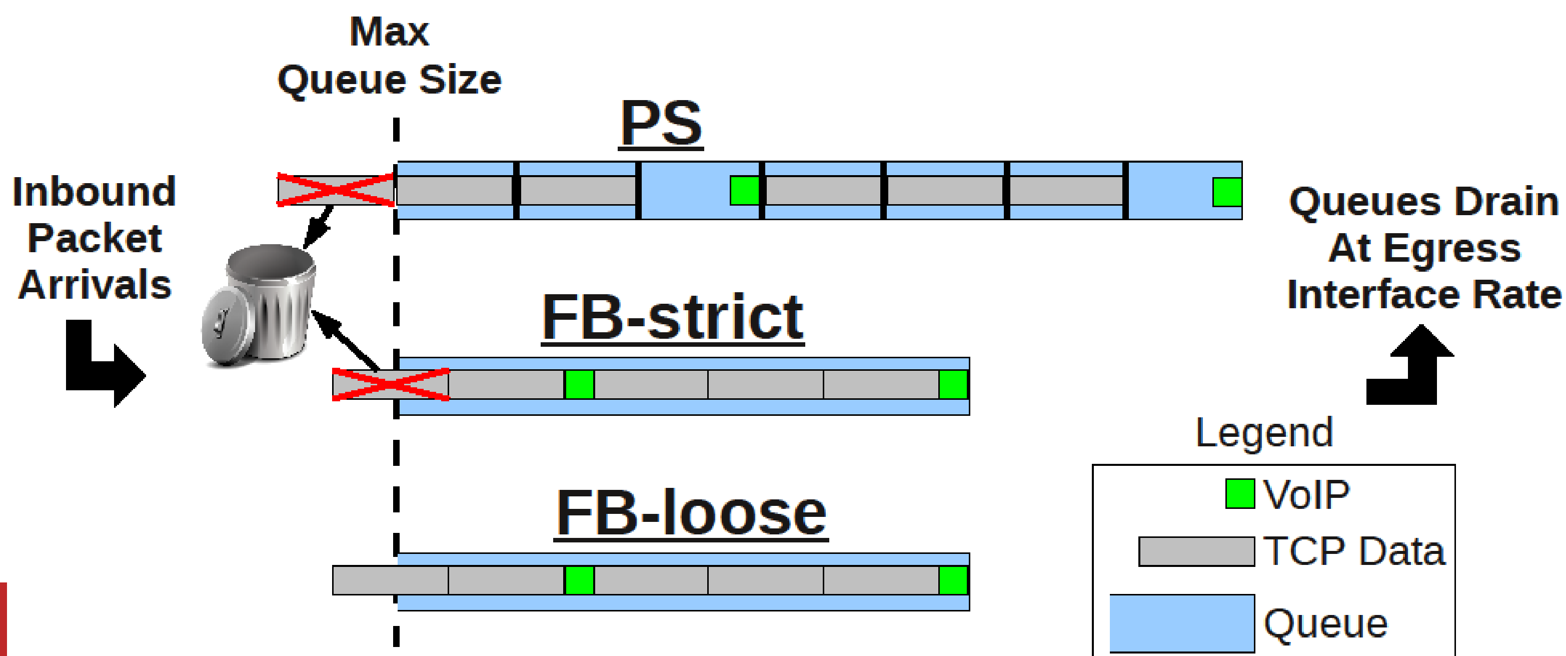

TCP Congestion Control Recap

- NewReno is ailing defacto standard (high latency/speed, wireless
- Many new algorithm proposals: Compound, CUBIC, HTCP, ...
- BSD still uses NewReno
- Linux uses CUBIC
- Windows Vista / 7 have Compound TCP available for testing

TCP Congestion Control Recap

Downstream NewReno Behaviour

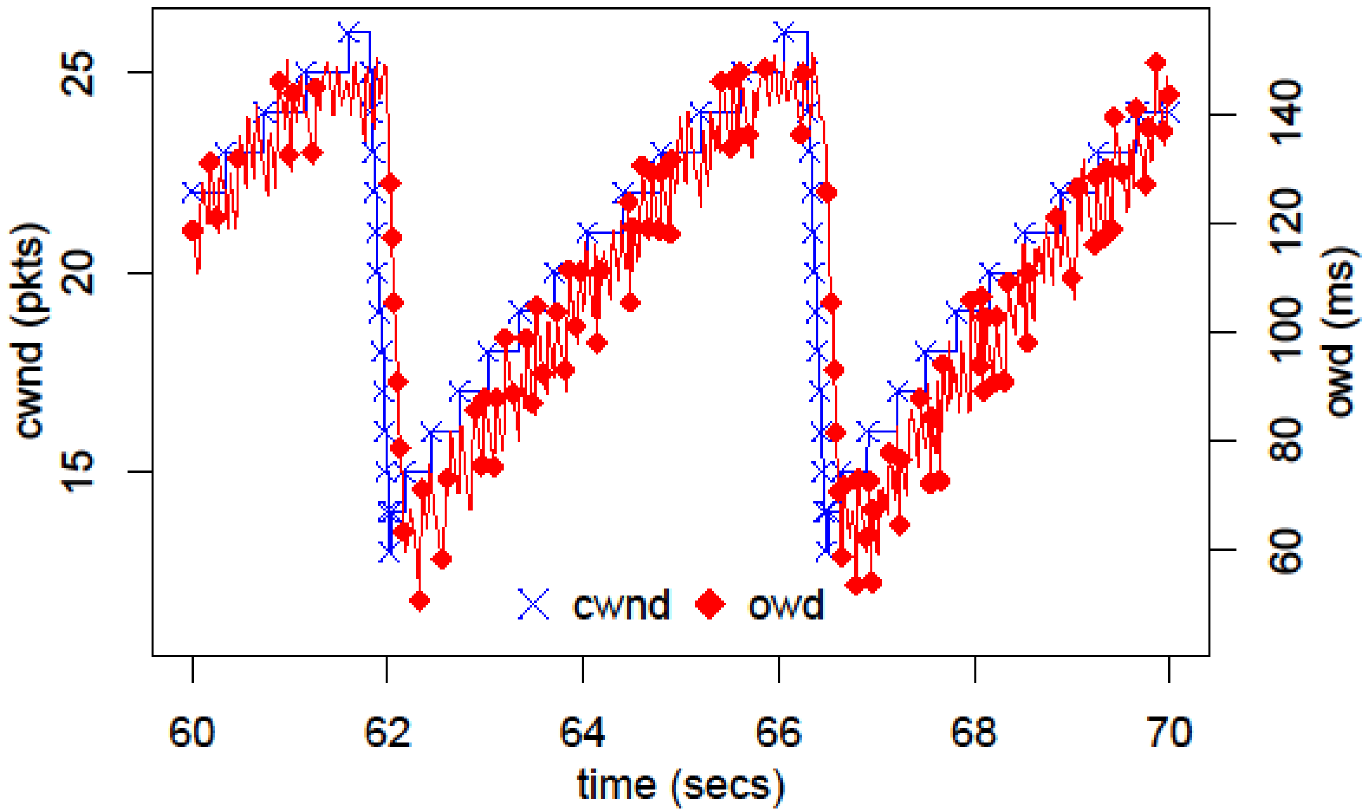

Congestion window oscillation coupled with induced queuing delay.

ADSL 1 type link: 1500/256 kbps, min 100ms RTT, 20 000 byte queue.

Methodology

- Testbed & NS-2
- 1500/256kbps & 24/1Mbps (ADSL 1 & 2 speeds)
- Byte-based [10k, 100k] & slot-based [7, 67] queues
- RTTs of 24, 50, 100 and 200ms
- Dummynet for bandwidth, queue size & RTT emulation
- Single bulk TCP flow A->C
- 186 byte, N(μ=20ms, σ =1ms) CBR UDP flow B<->D

Testbed

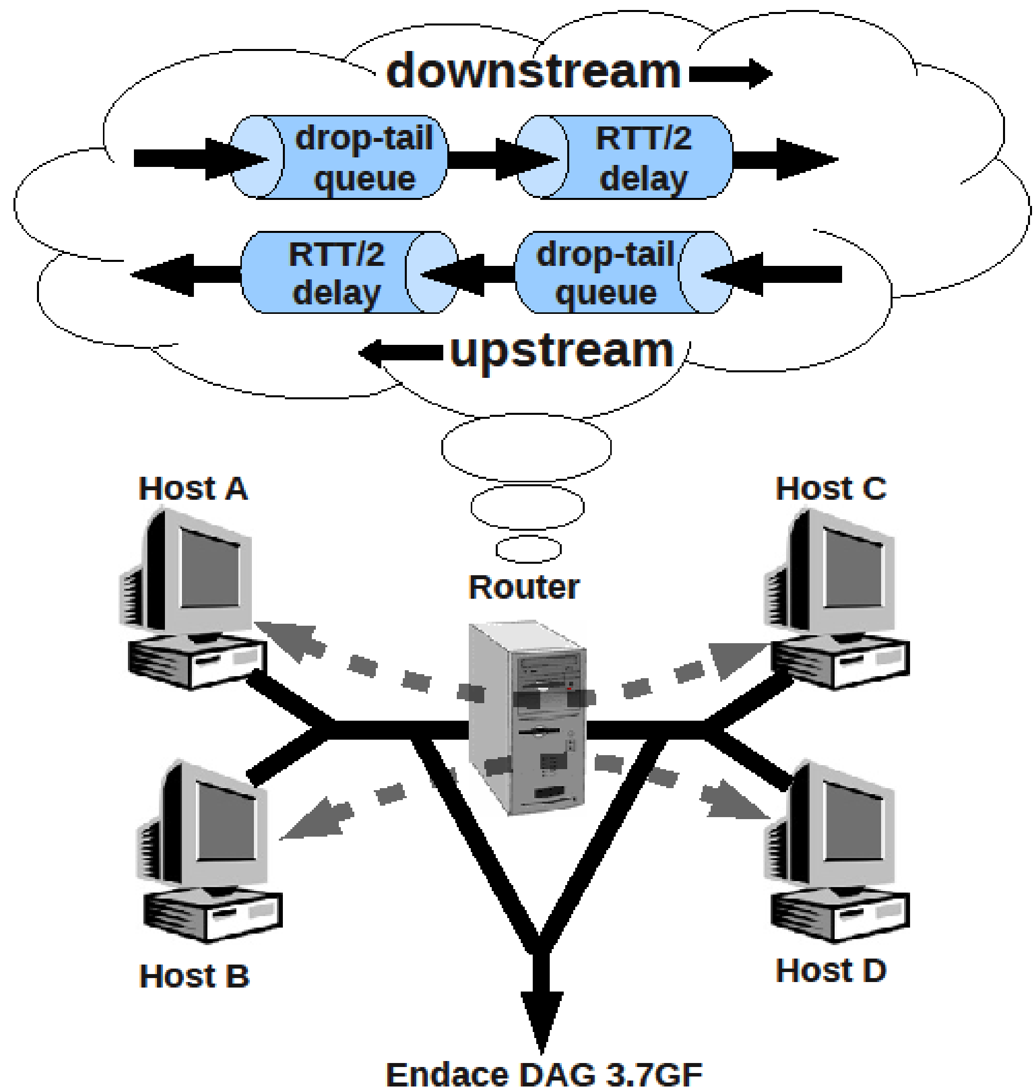

Detailed outline (section 2 of 3)

1. Introduction
2. Results
3. Conclusion

2. Results
   - Goodput
   - Collateral Damage to Voice Stream – Latency
   - Collateral Damage to Voice Stream - Loss

Goodput

ADSL 1 (1500/256kbps)

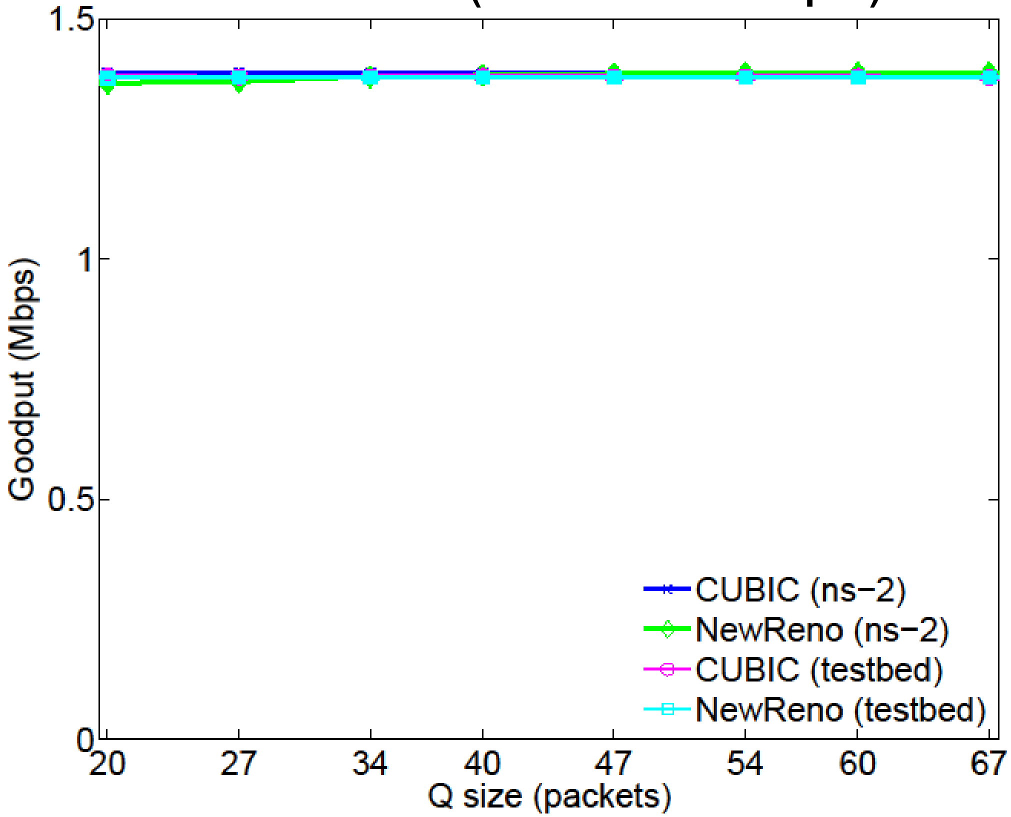

ADSL 2 (24/1Mbps)

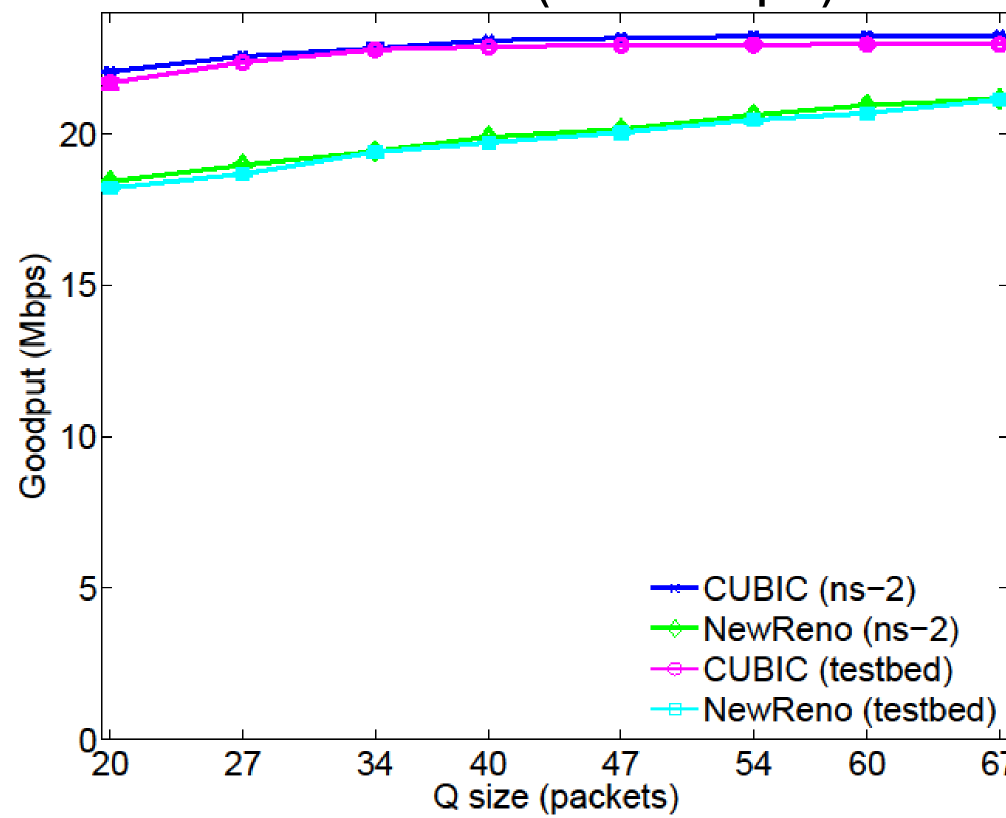

Downstream goodput vs PS queue length - measured (testbed) and simulated (NS-2), 1500/256 kbps & 24/1 Mbps links, 100ms min RTT.

CUBIC Induces More OWD Than NewReno

Downstream NewReno

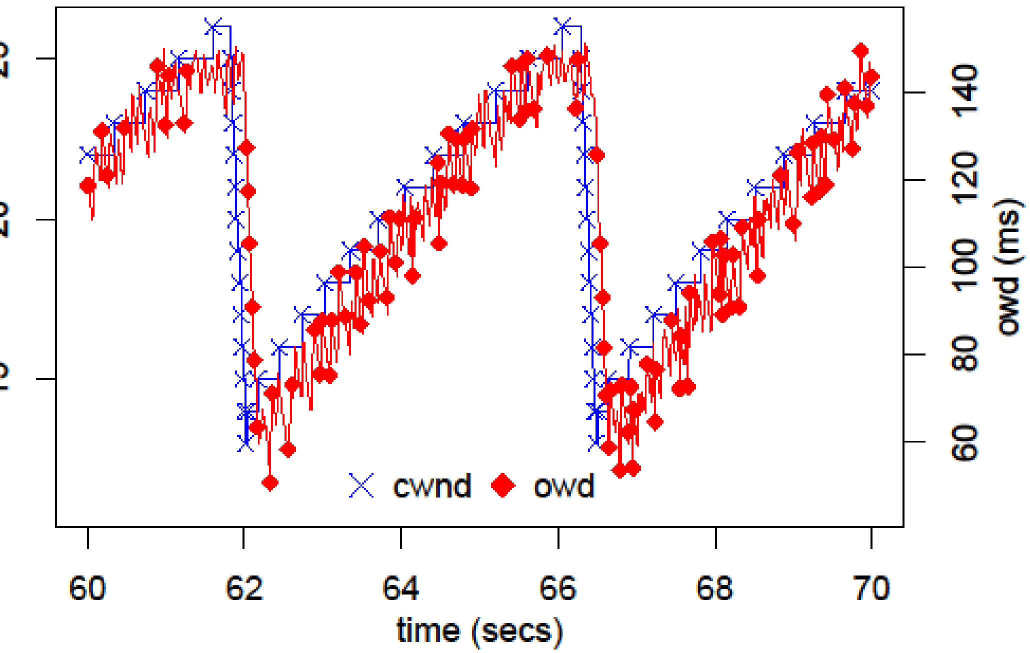

Downstream CUBIC

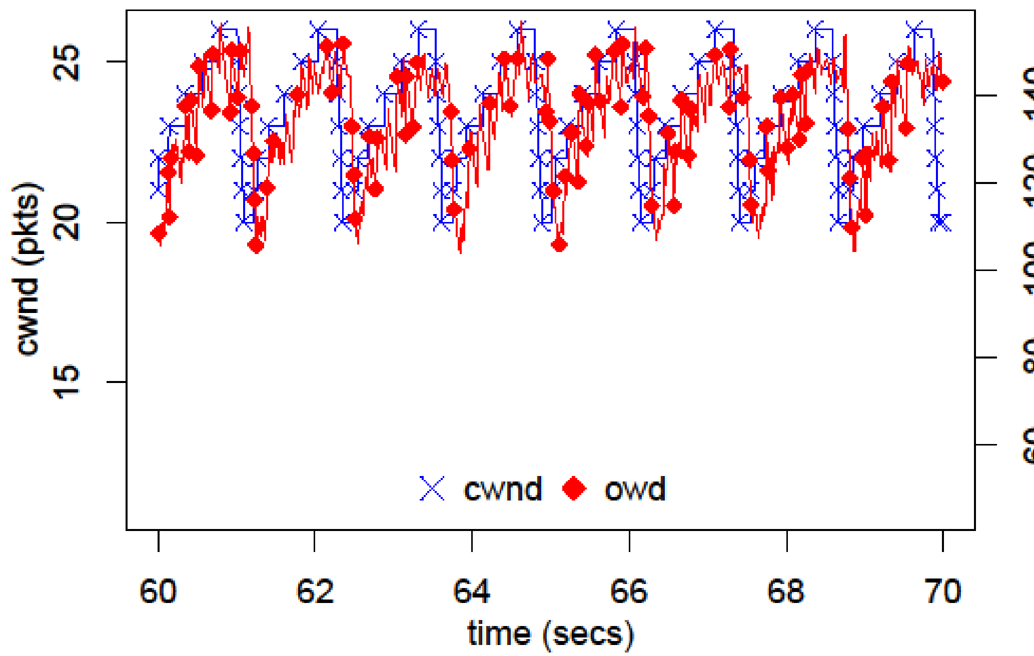

Congestion window oscillation coupled with induced queuing delay.

ADSL1 type link: 1500/256kbps, min 100ms RTT, 20 000byte queue.

Collateral Damage to Voice Stream - Latency

ADSL 1 (1500/256kbps)

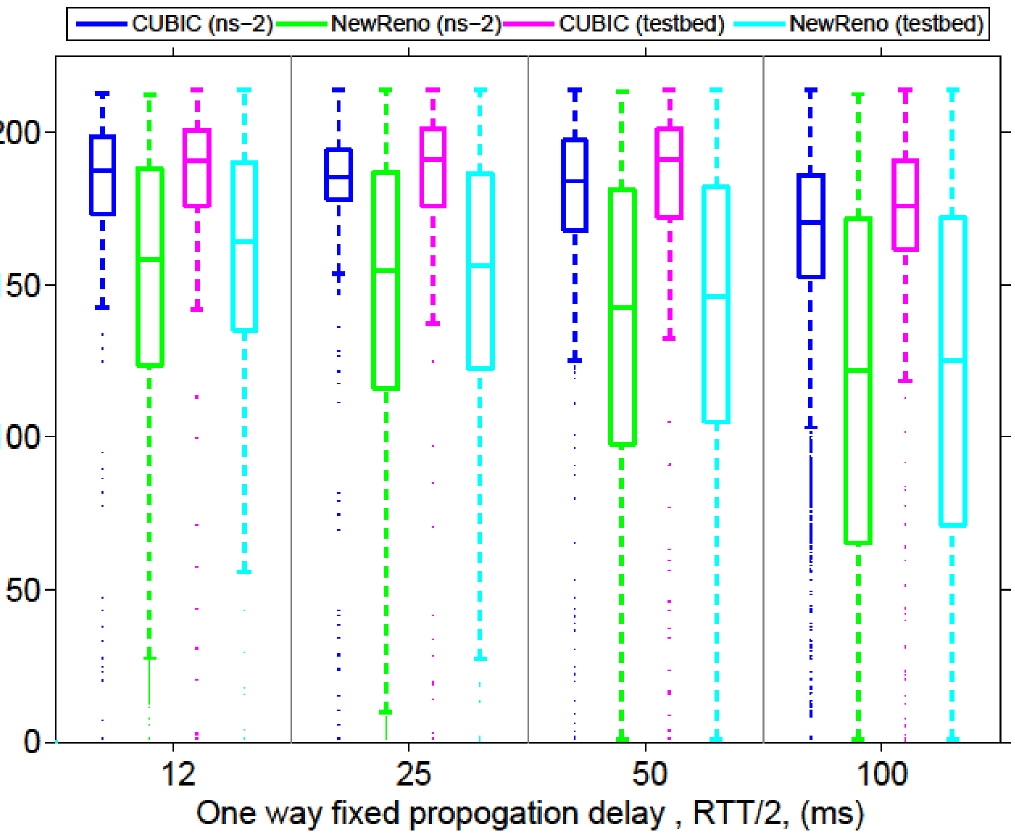

ADSL 2 (24/1Mbps)

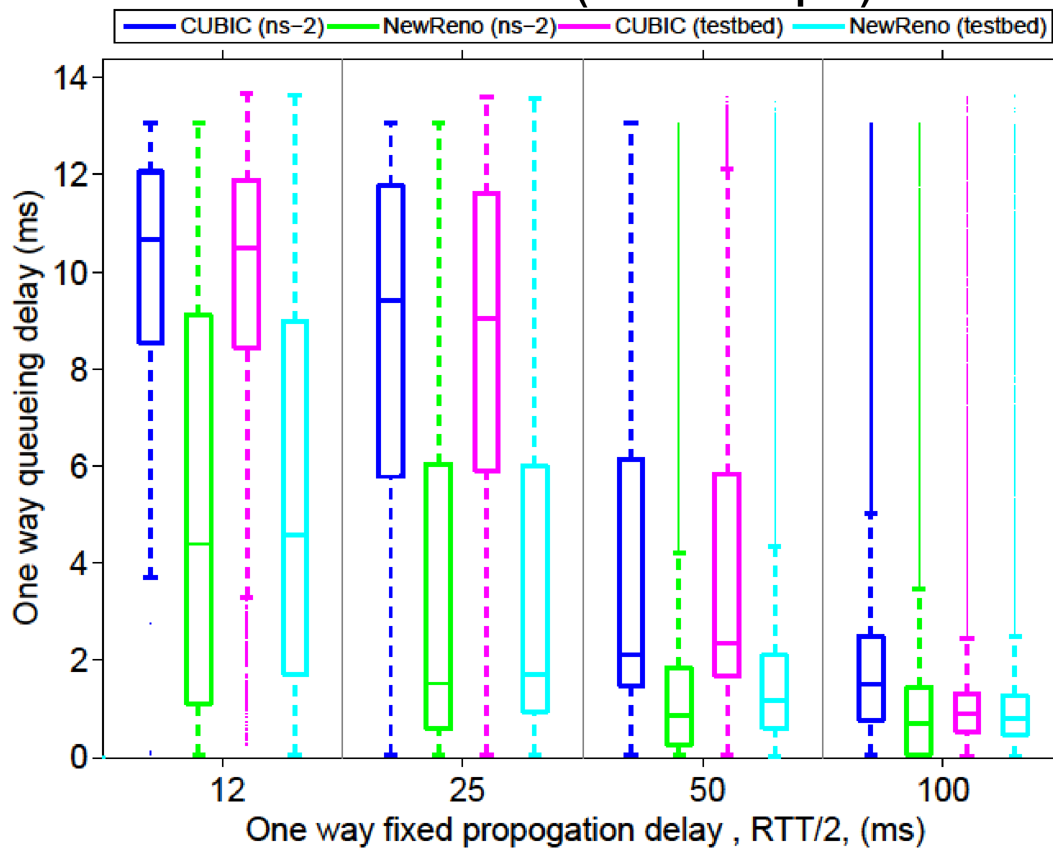

Induced downstream queuing delay - measured (testbed) and simulated (NS-2), 40 000byte FB-loose queue.

Collateral Damage to Voice Stream - Loss

FB-loose, FB-strict & PS

NS-2

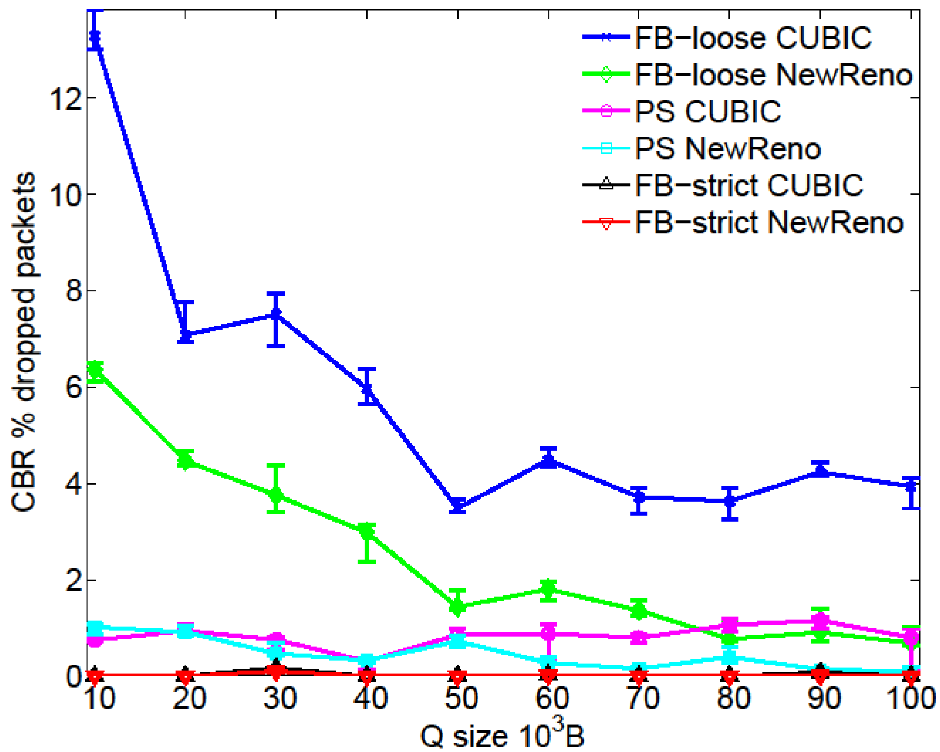

FB-loose

Testbed & NS-2

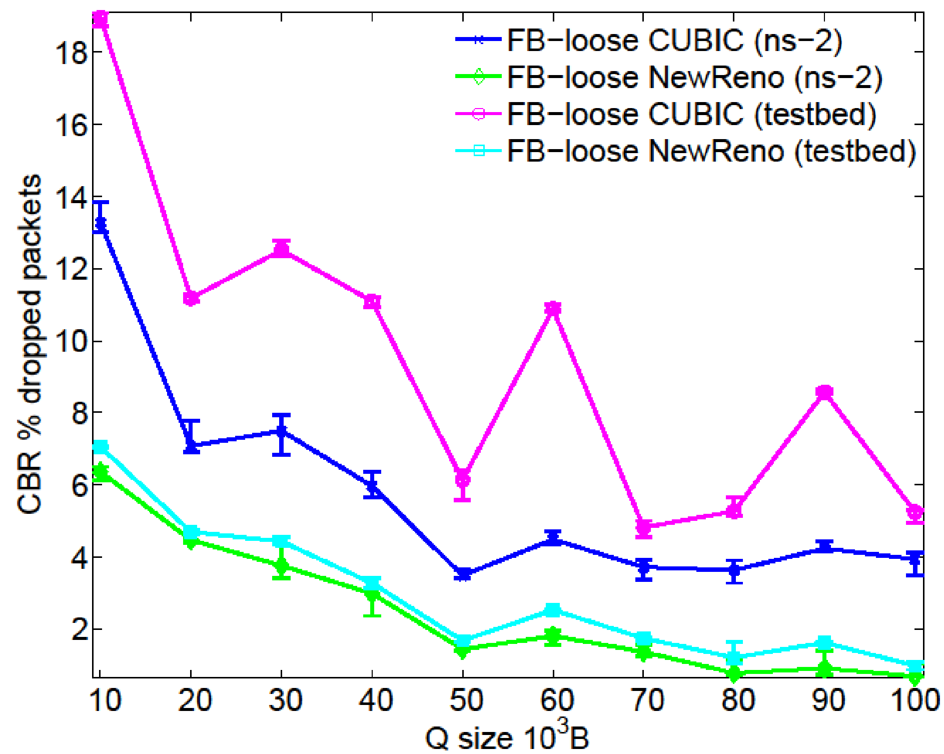

Proportion of CBR packets lost on the downstream for different packet drop mechanisms - 1500/256kbps link, 100ms minimum RTT.

Detailed Simulation: Impact of Queue Size on Loss

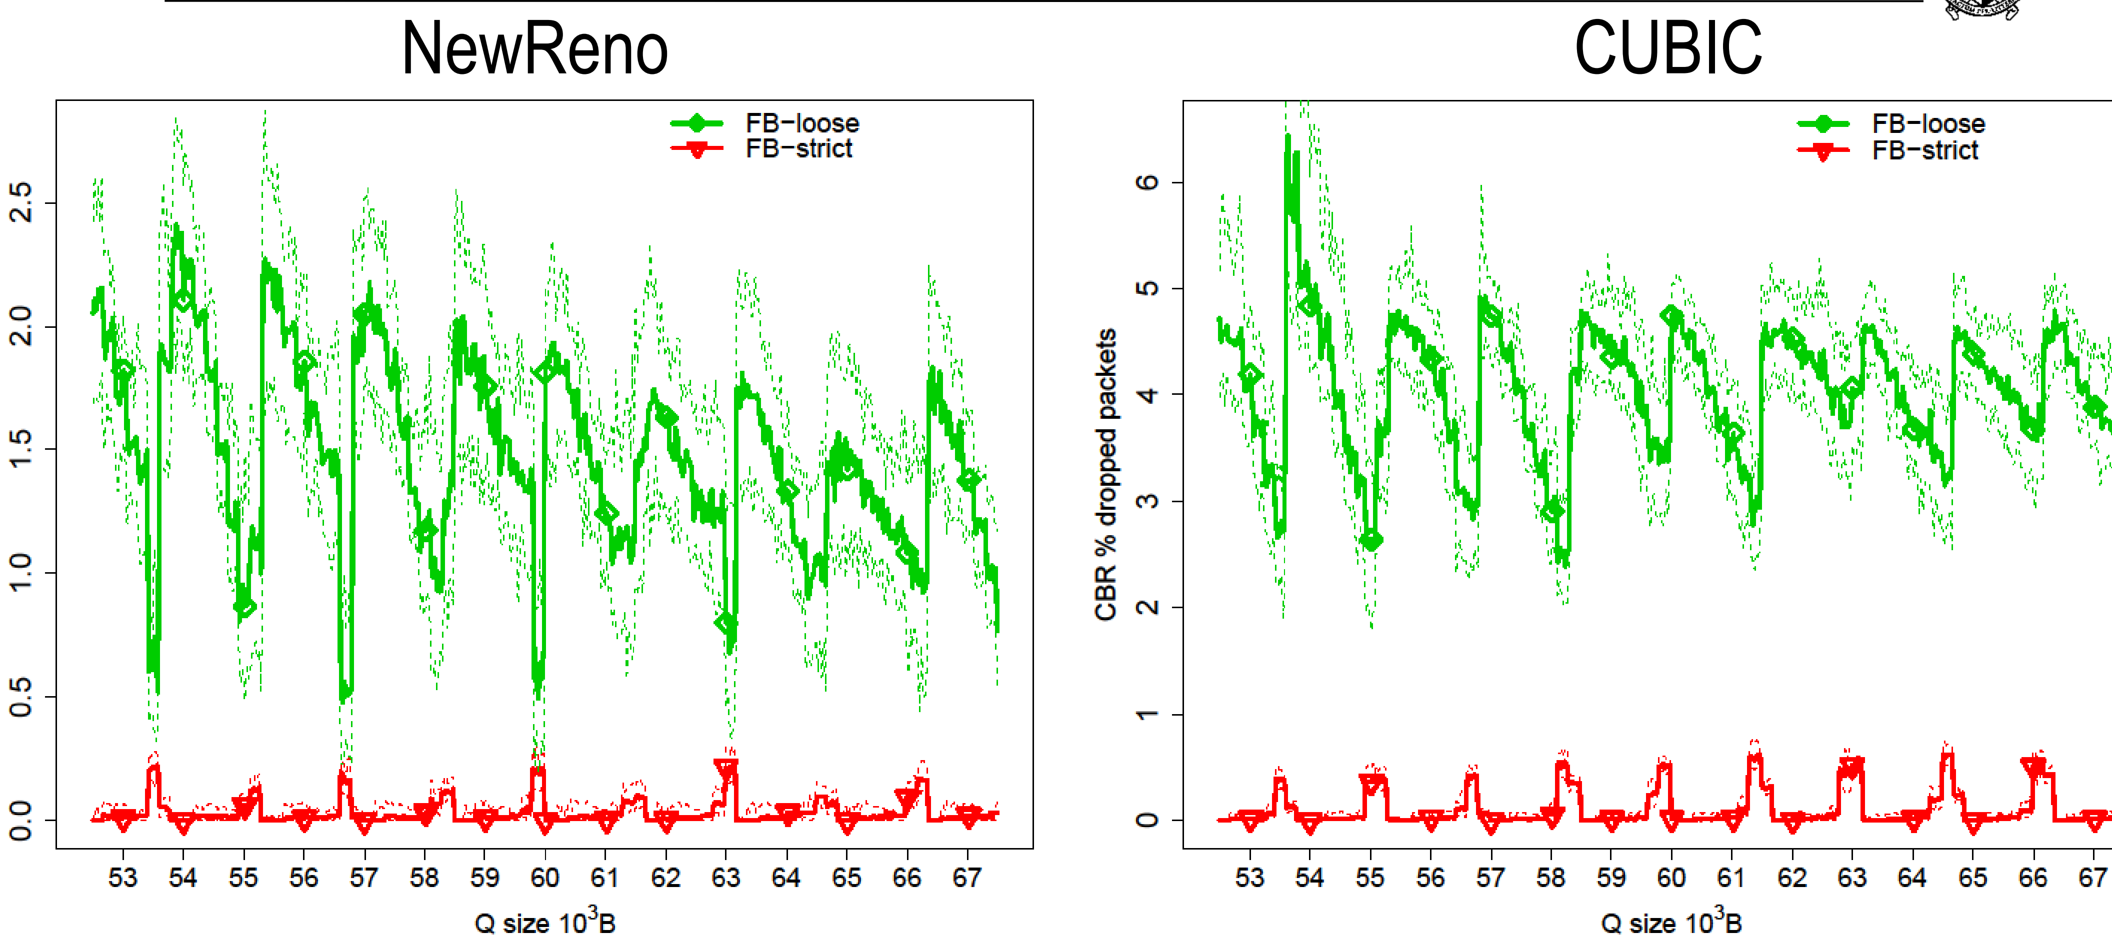

[52500, 67500] in 30 byte increments, 1500/256kbps, 100ms min RTT.

Dashed line envelope - 5% & 95% confidence intervals.

Time Series Analysis of FB-Strict

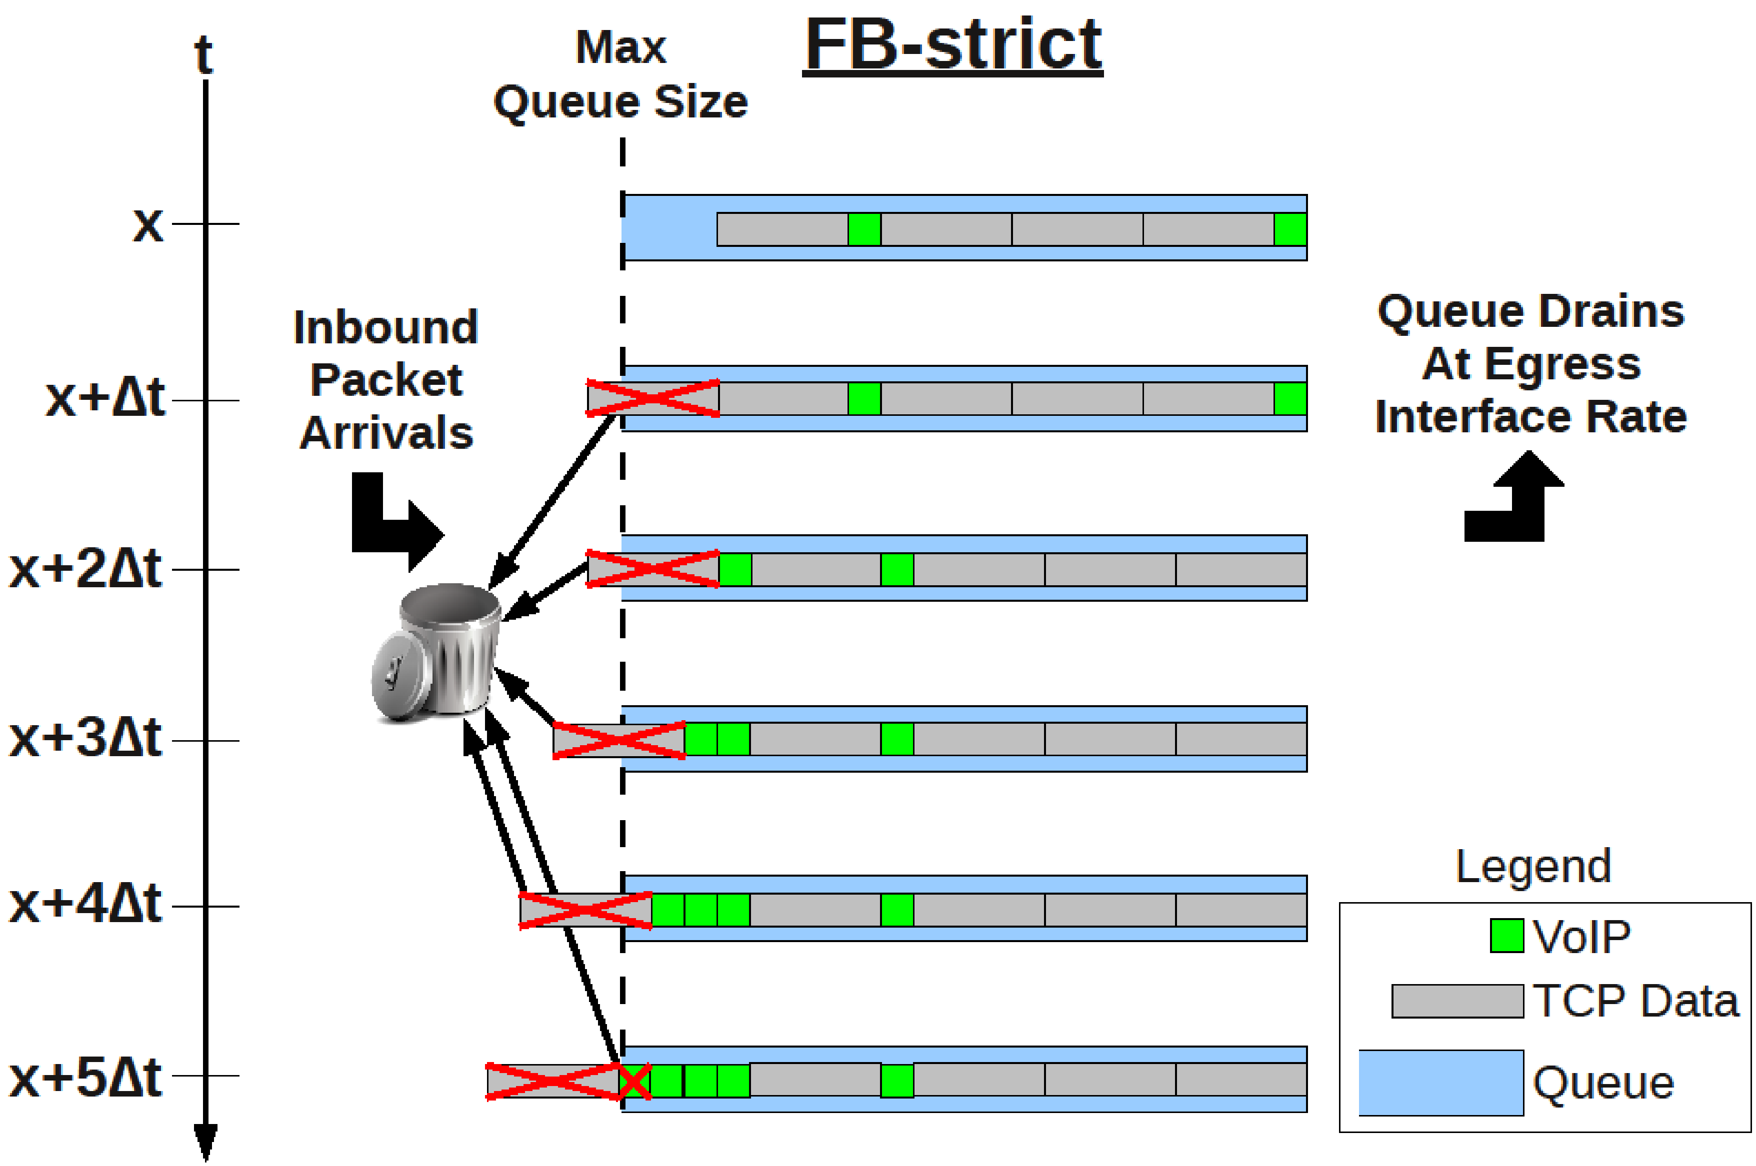

Time Series Analysis of PS

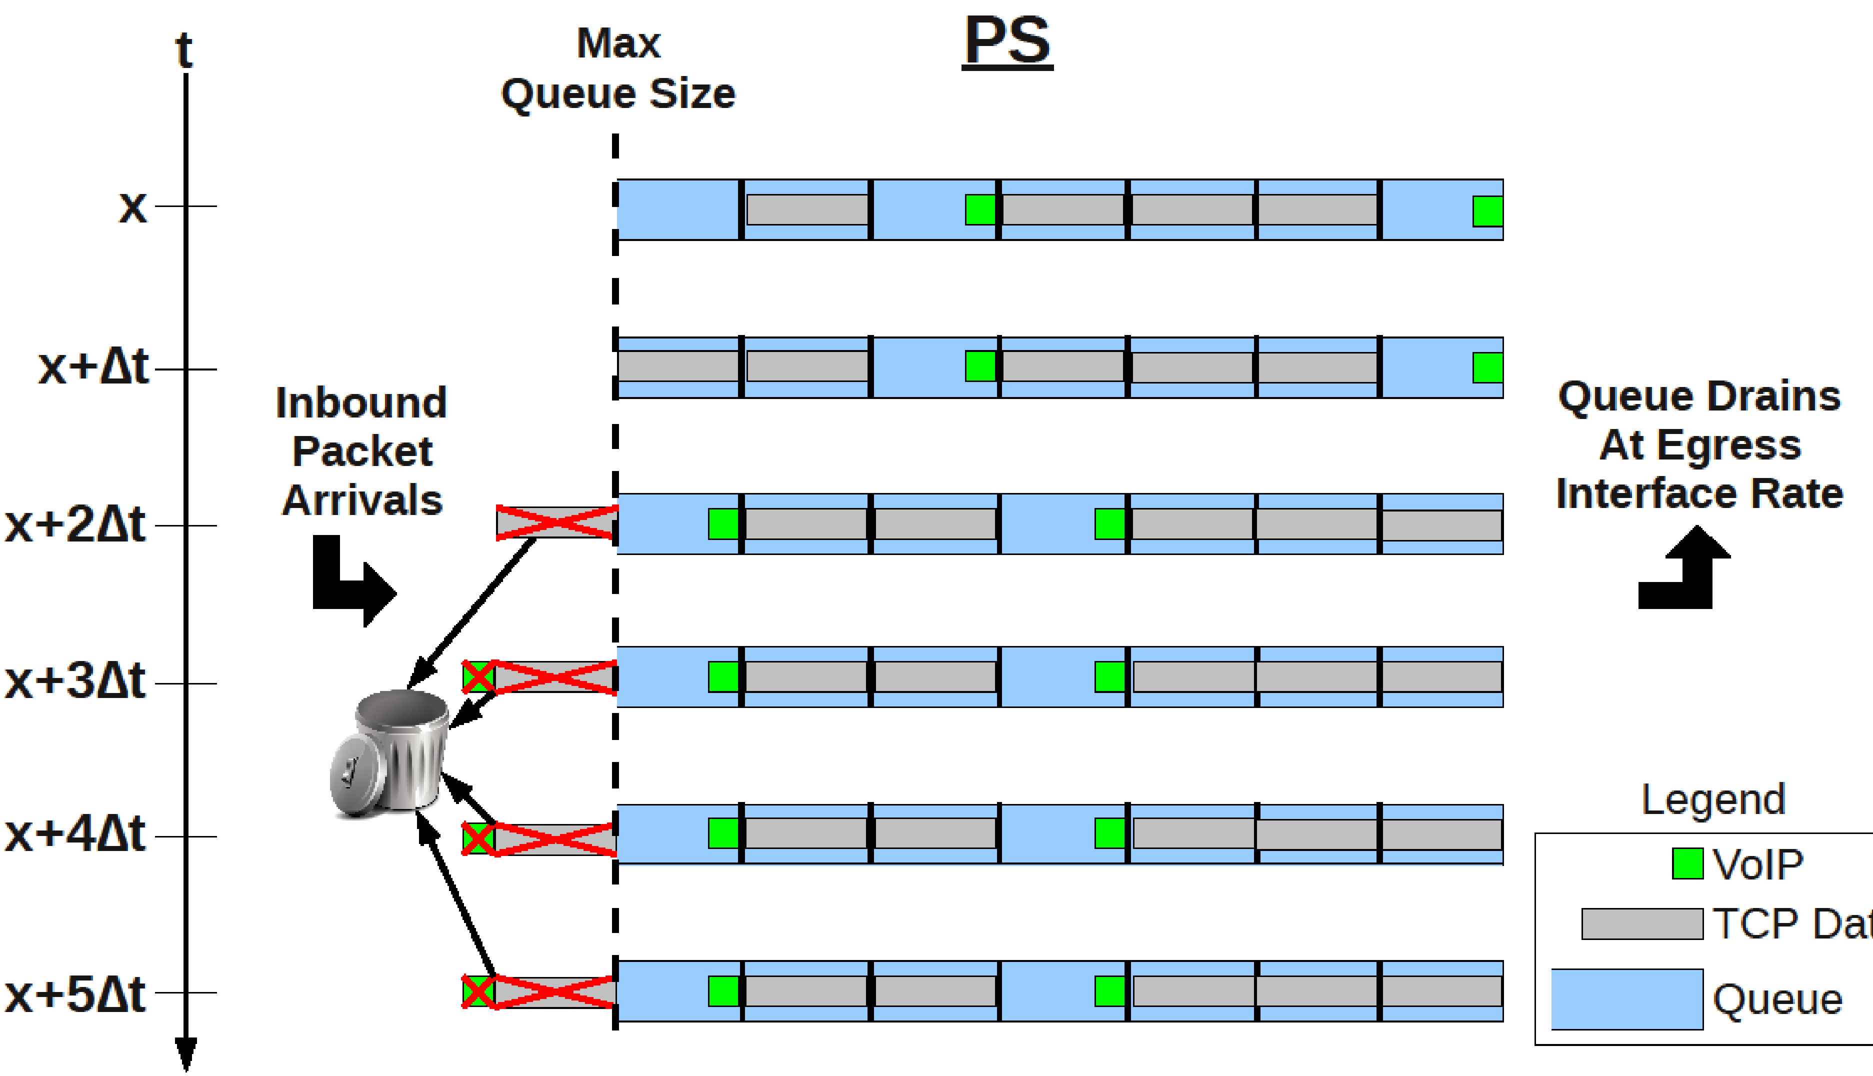

Time Series Analysis of FB-Loose

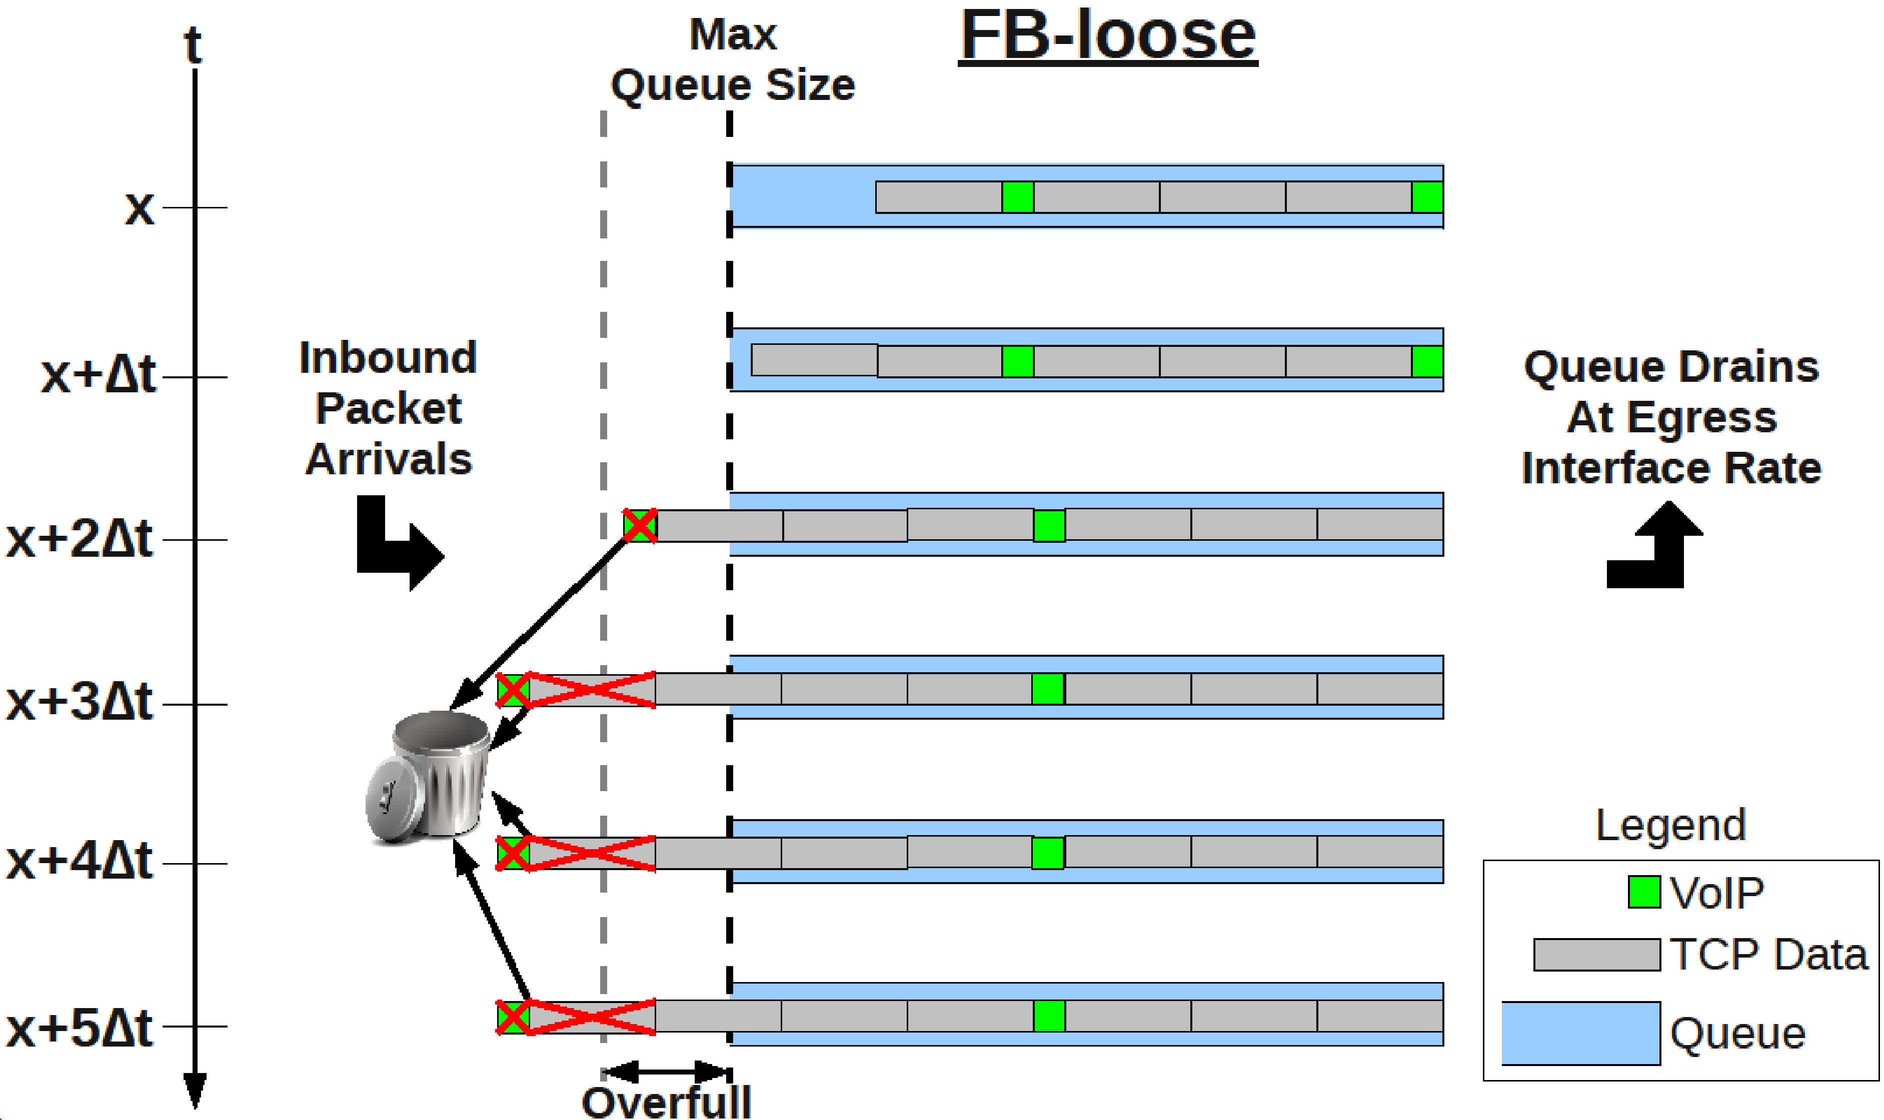

Detailed outline (section 3 of 3)

1. Introduction
2. Results
3. Conclusion

3. Conclusion
   - Summary & Open Questions

Summary & Open Questions

- Summary
  - CUBIC keeps buffers fuller: latency & loss increase
  - FB-Loose disadvantages small, real-time multimedia packets
  - Byte-based queue sizes around integer multiples of TCP packet size correspond with lowest FB-Loose loss & highest FB-Strict loss
  - Dummynet implements FB-Loose byte-based queues
- Open Questions
  - Does using TCP for interactive traffic change things?
  - Can we develop guidelines for home gateway manufacturers?