ICCRG 99

Steady state tensions: Problems and Solutions

Koen De Schepper
koen.de_schepper@nokia-bell-labs.com

Bob Briscoe
ietf@bobbriscoe.net

Ongoing (r)evolutions

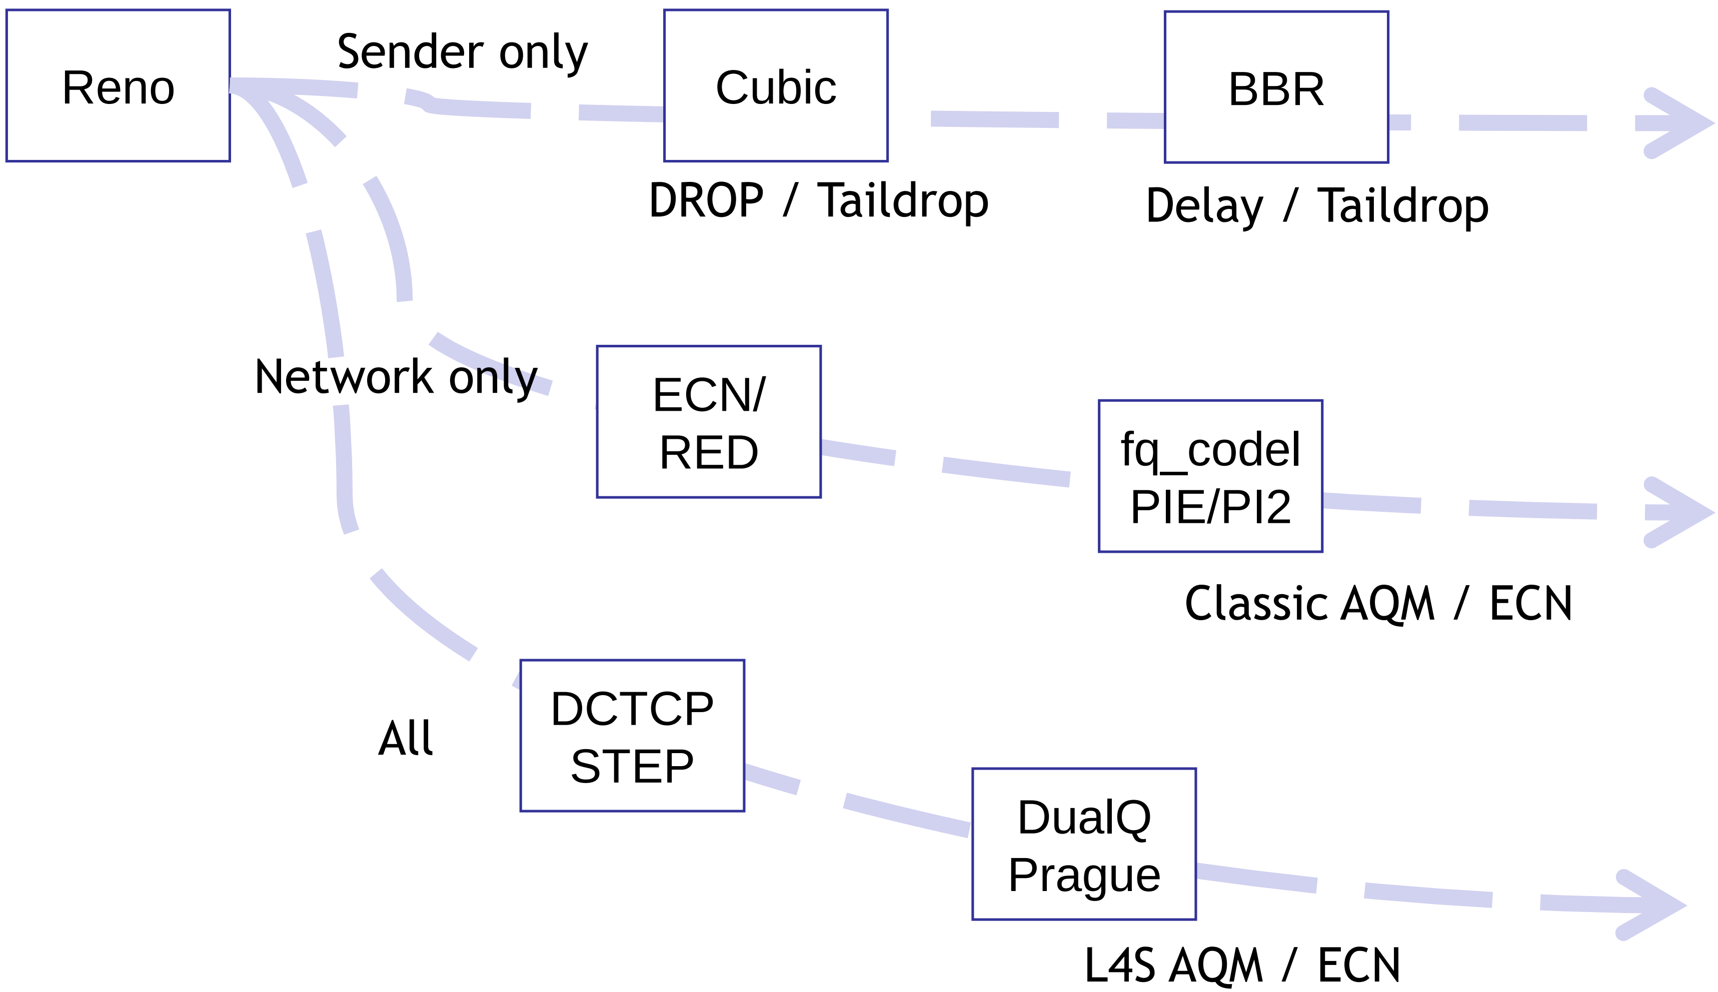

Ongoing (r)evolutions

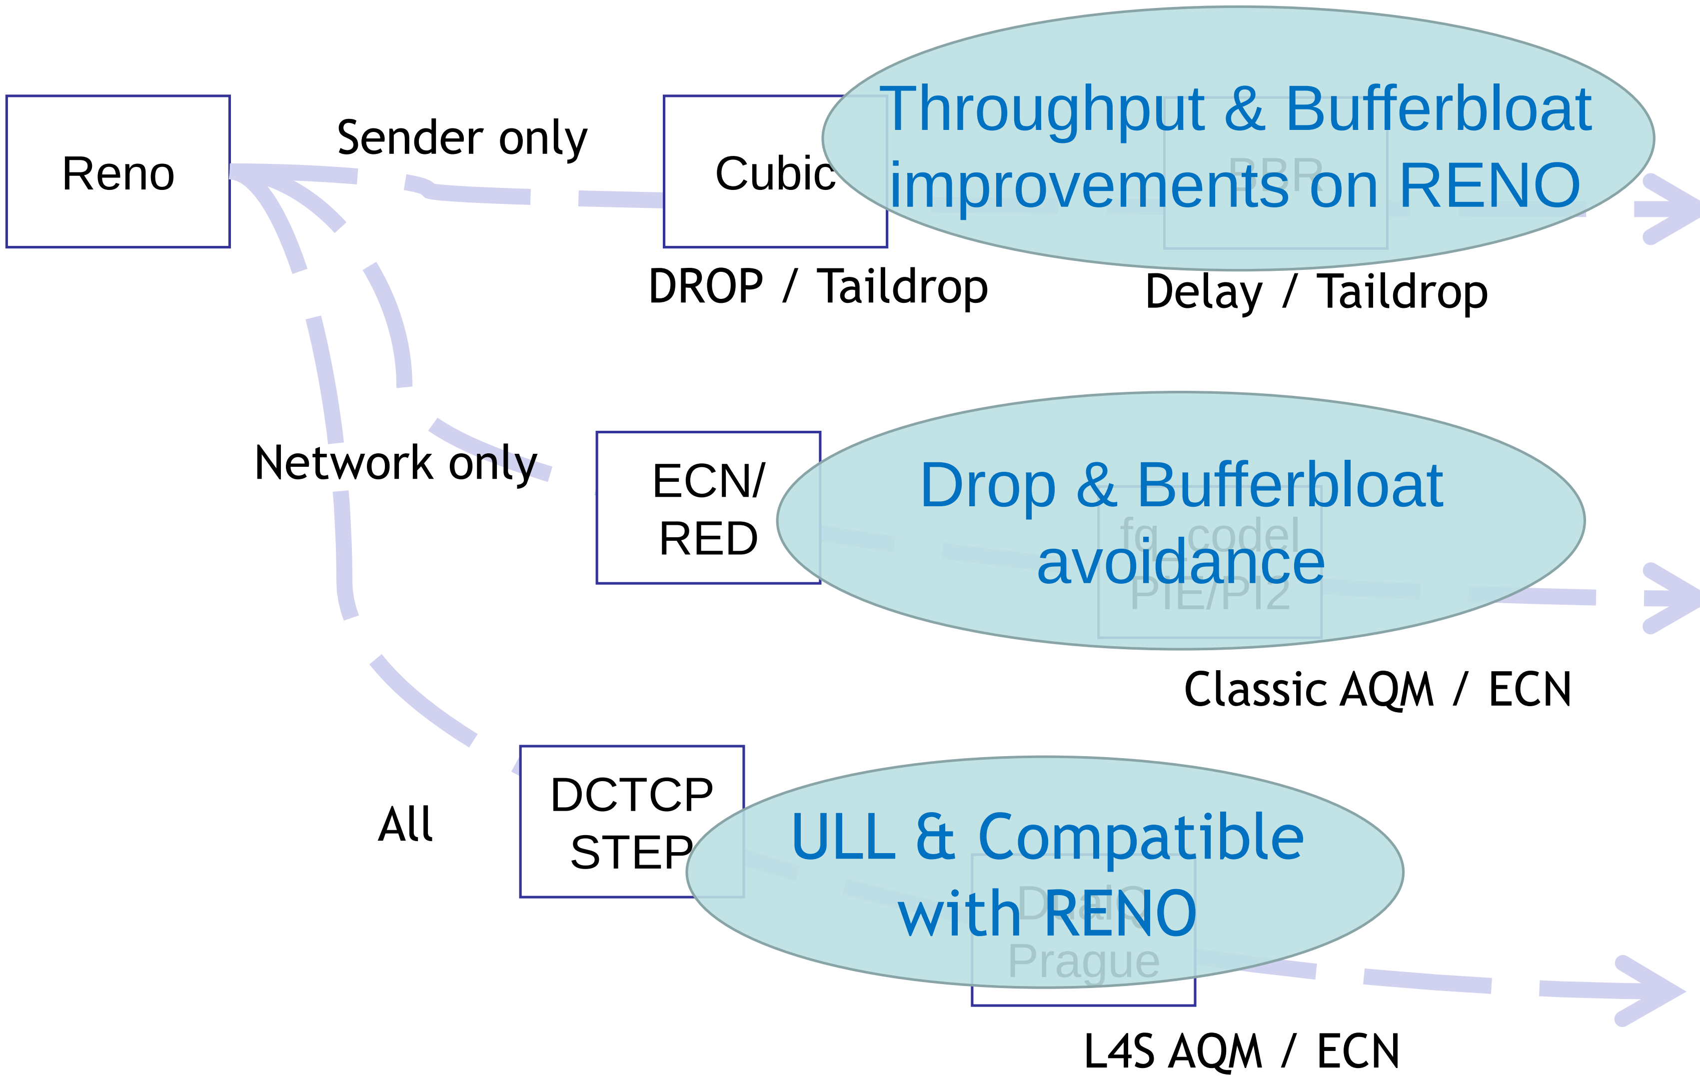

Sender only evolution: BBR & Taildrop / Bufferbloat

BBR controls the queue on 20ms on 100Mbps & 20ms RTT link

High throughput and full link for long flows

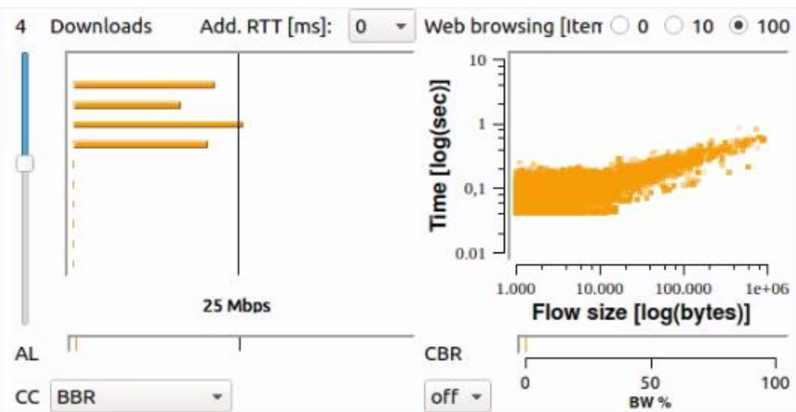

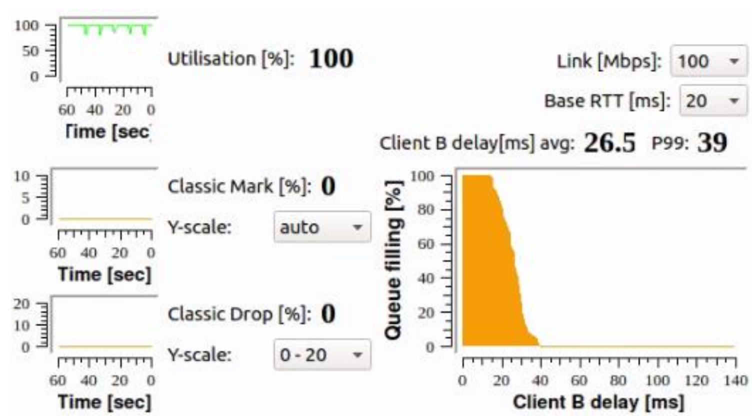

No drop and limited delay for short flows

Short flow variations build on top of queue target

Network only evolution:
Cubic & AQM

Cubic and PIE under good working conditions 100Mbps & 20ms RTT

PIE AQM controls queue target to 15ms using only 0.01% drop

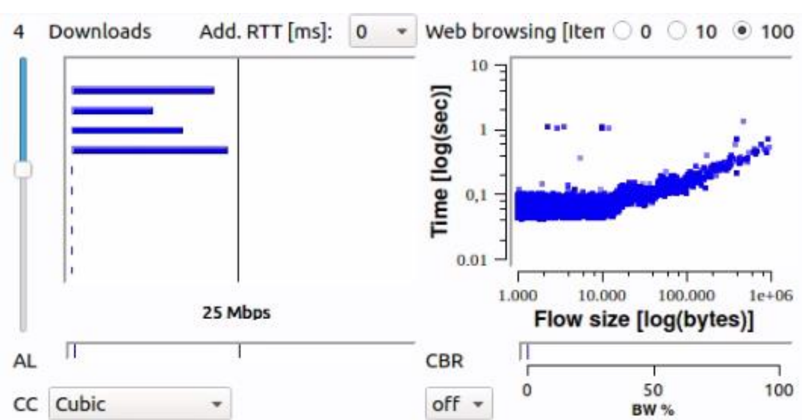

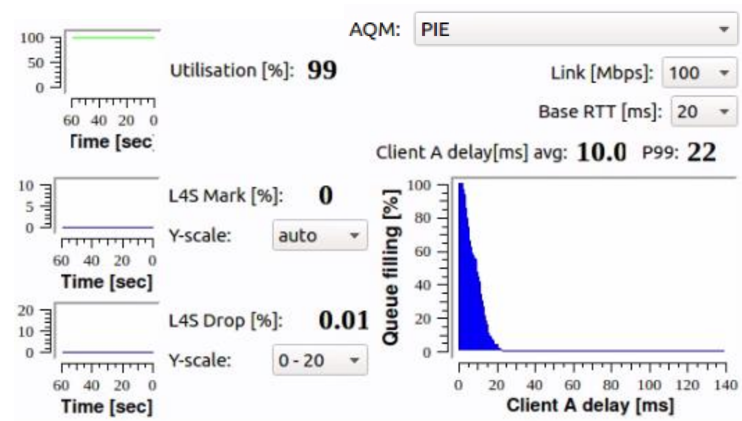

Limited drop and delay for short flows

Short flow variations vary around queue target

Network AND end-system evolution: DCTCP & STEP (DualQ)

DCTCP controls queue below step threshold on 100Mbps & 20ms RTT

Limited link utilization to keep ultra-LL

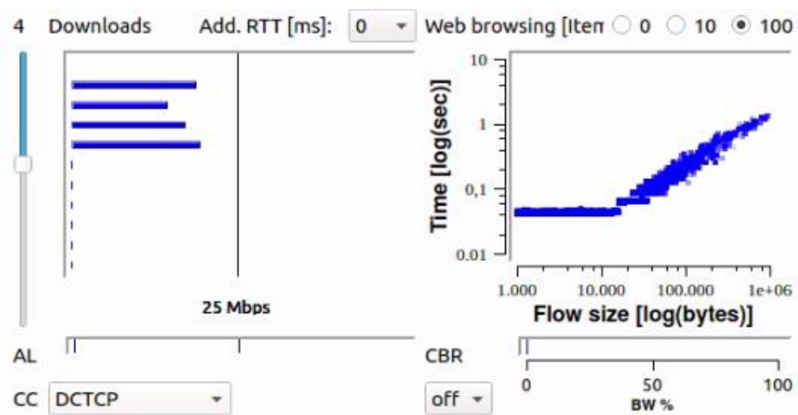

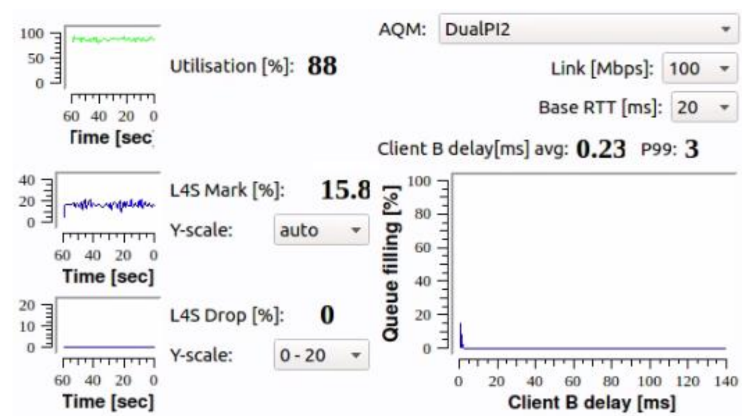

No drop and ultra low delay for short flows

Short flow variations are below queue target

Concerns with combinations

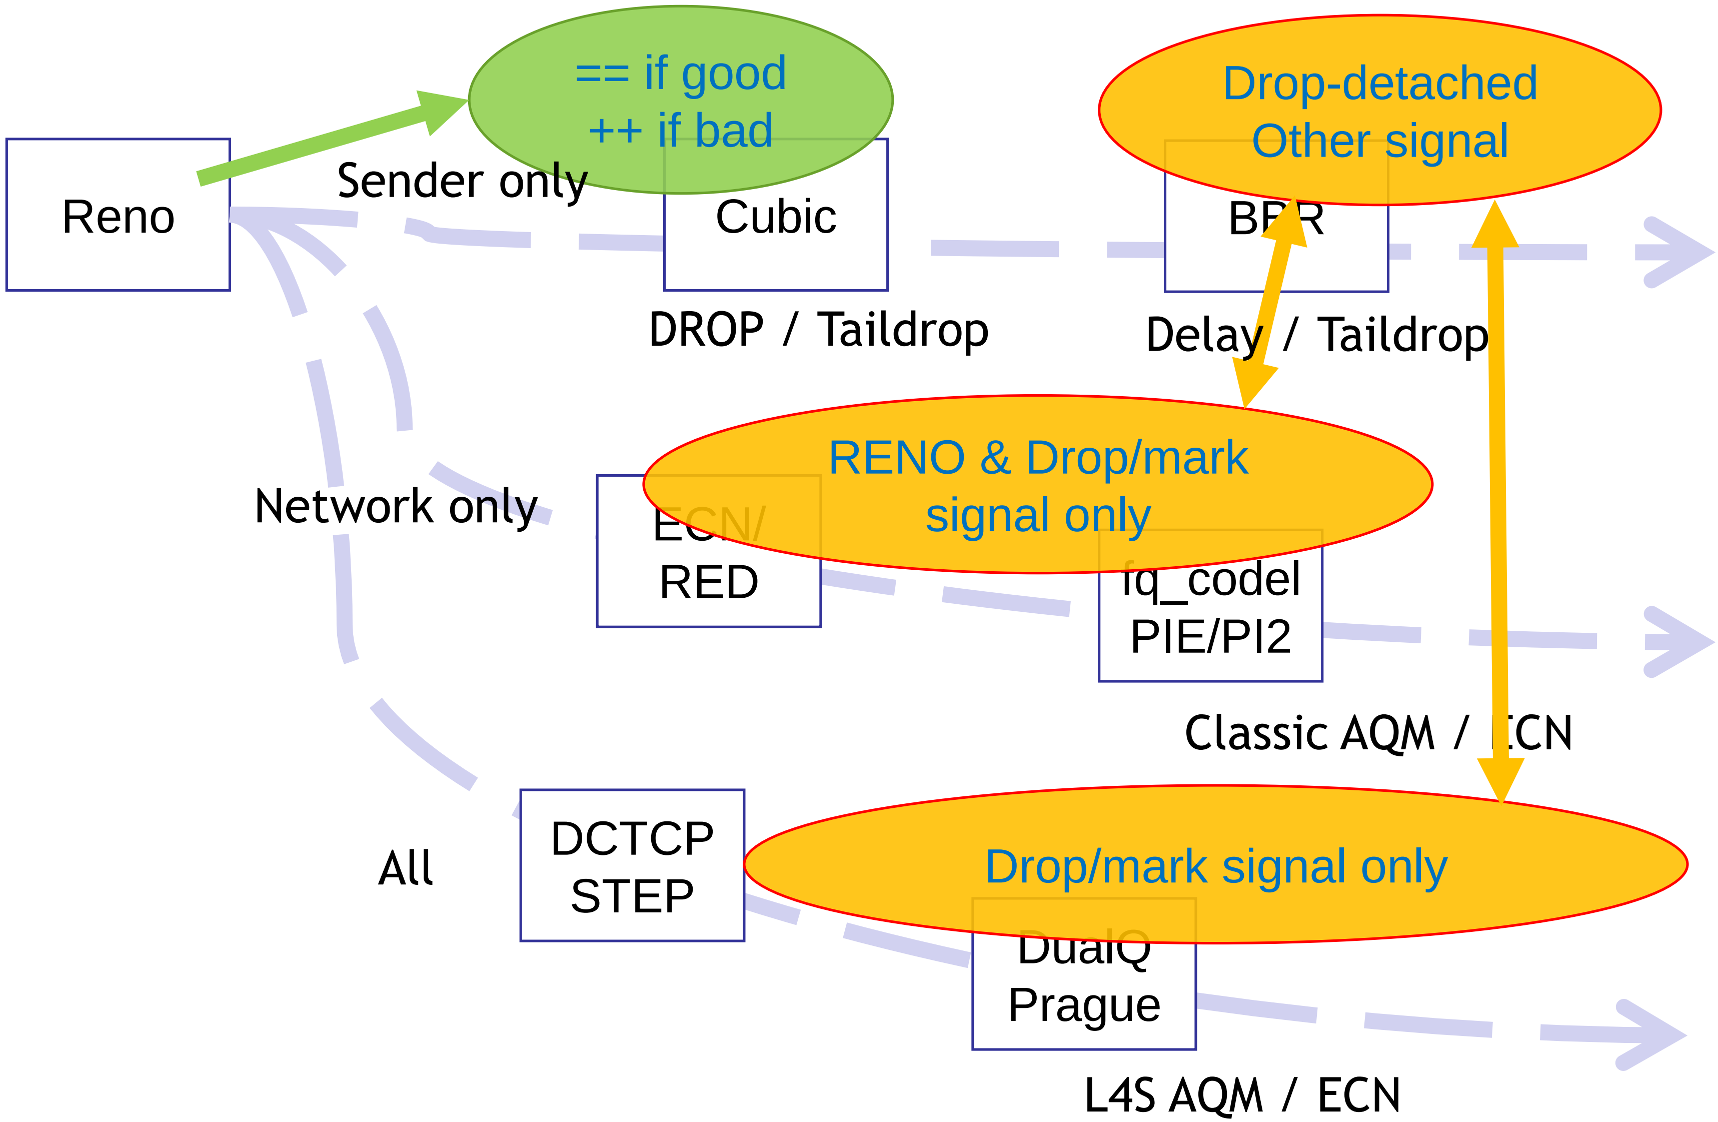

Network-only and Sender-only clash: BBR & AQM

PIE AQM tries to limit queue to 15ms, needs 5 to 20% drop

BBR enforces bigger Q and does not respond to high drop probability

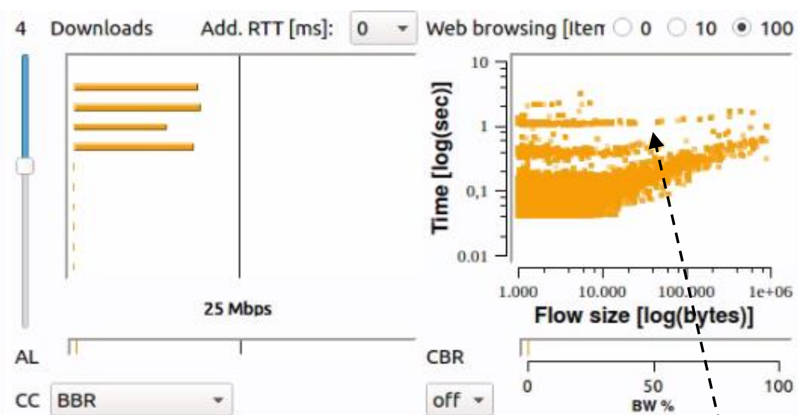

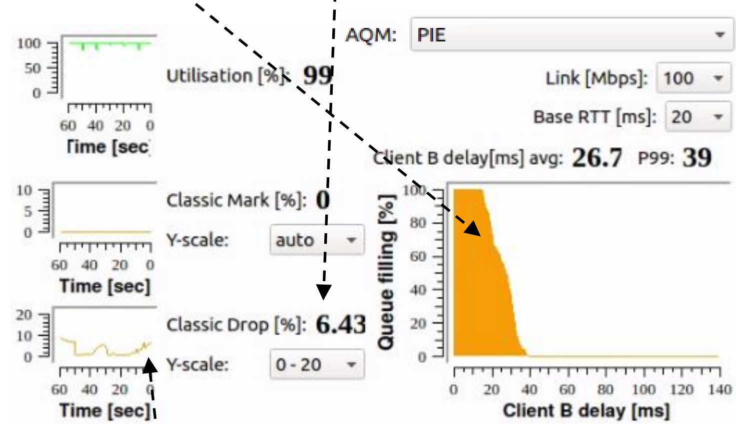

Very bad short flows completion times

10-20s very high drop 5-20%

10-20s policing detection 1%

Network-only and Sender-only clash: BBR & AQM

PIE AQM tries to limit queue to 15ms, needs 5 to 20% drop

BBR enforces bigger Q and does not respond to high drop probability

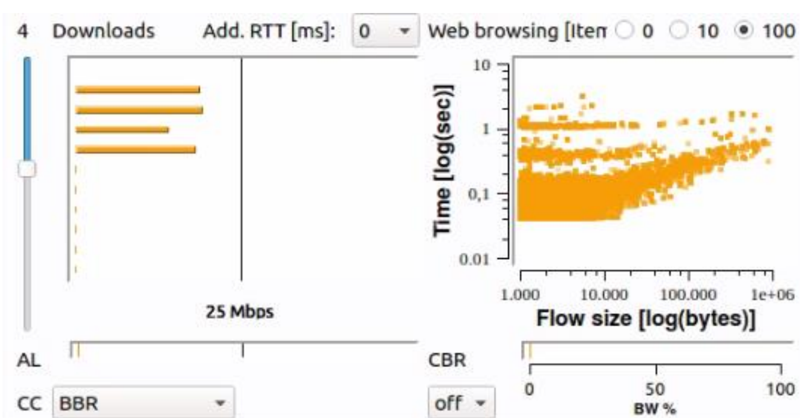

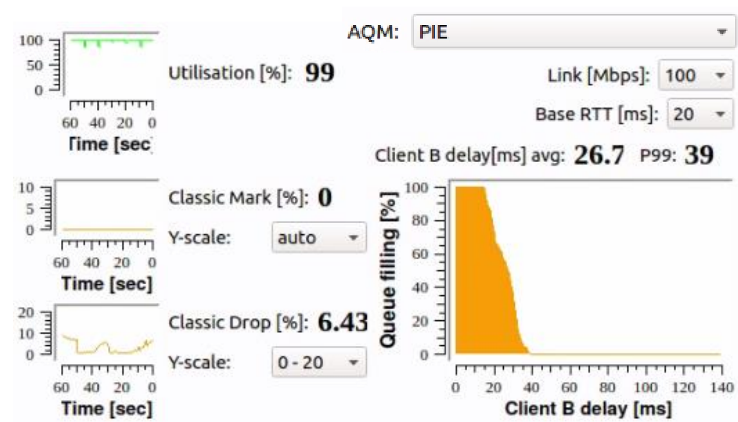

Will using BBR (as is) force operators to disable AQMs (with ECN) ?

Solution: Rate to drop_p & RTT relation

BBR:

- solve low throughput under high RTT and high loss cases of Classic flows

L4S:

- compatible with Classic, but not under pathological conditions (as it will also get low throughput)

Solution:

- both require corrected Classic throughput behavior under high RTT/Loss conditions
- allows compatibility between L4S and Classic (BBR++) in all cases

RTT independence and Scalability for Classic TCP

Change rules for Classic TCP:

- RTT independent above 5ms? ➔ 30Mbps at 1% drop
- Scalable for $p < 1\%$? ➔ 24 drop/marks per second

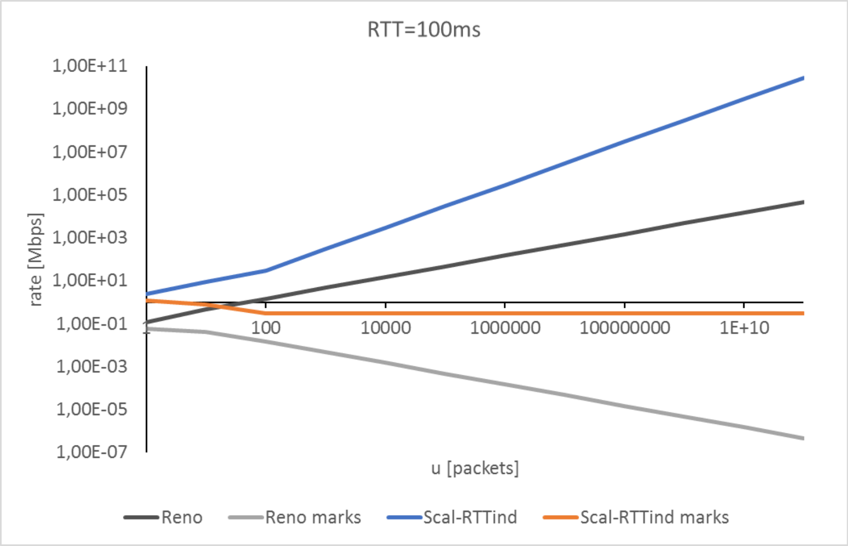

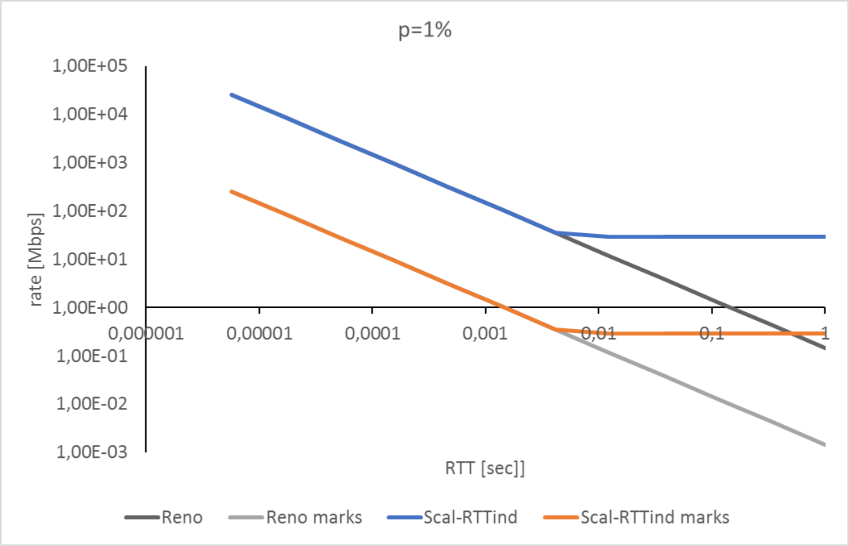

RTT independence for Classic TCP

RTT independent above 5ms? ➔ 30Mbps at 1% drop

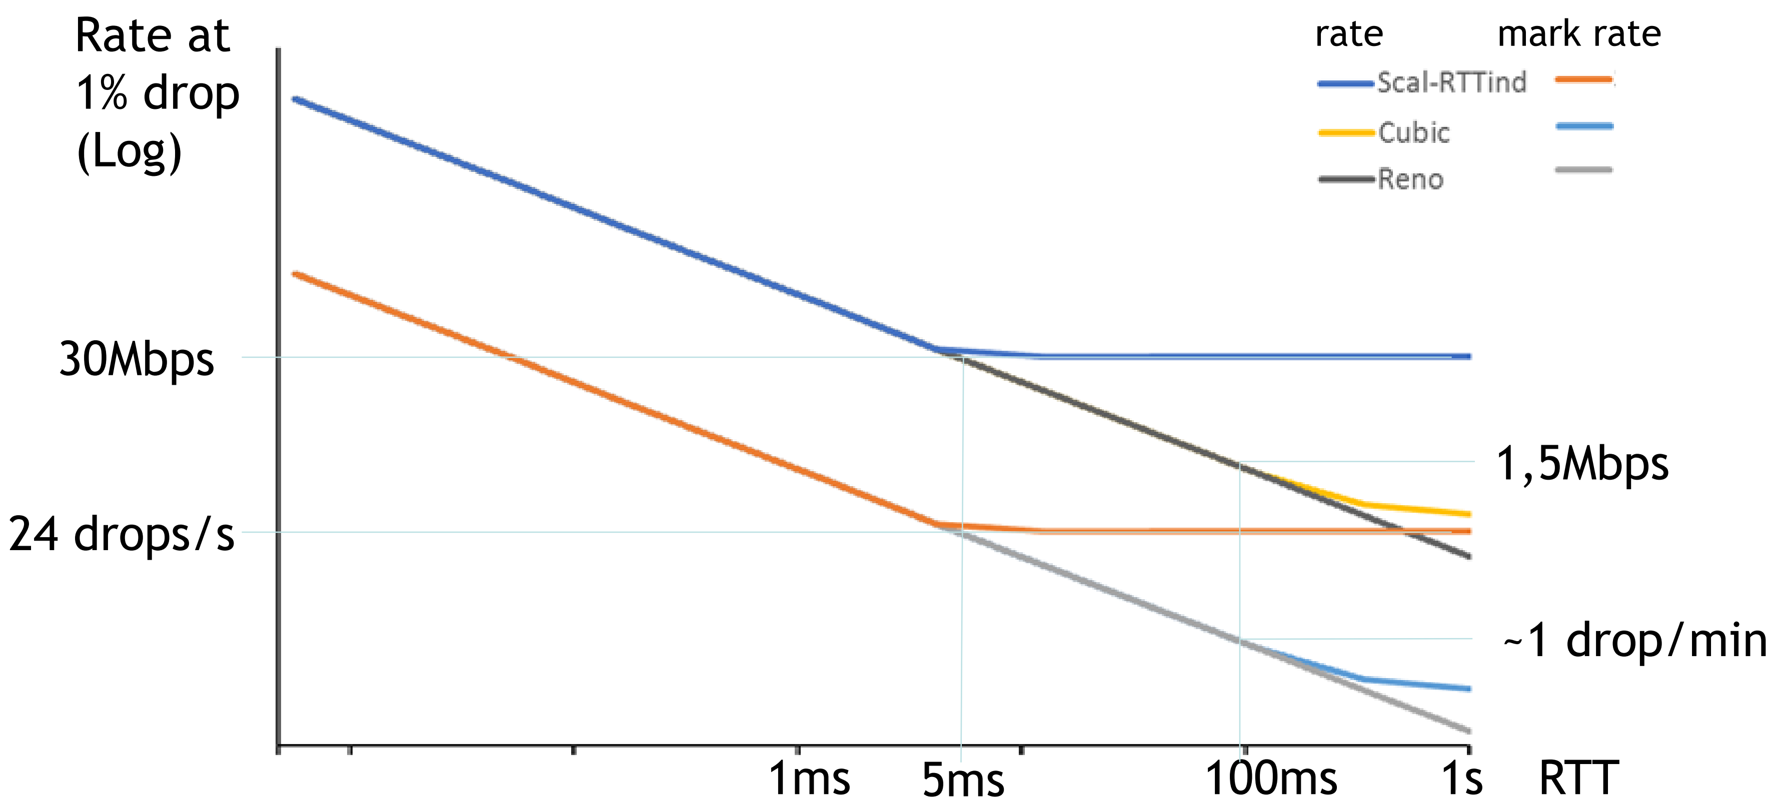

Scalable for Classic TCP

Scalable for p < 1%? ➔ 24 drop/marks per second

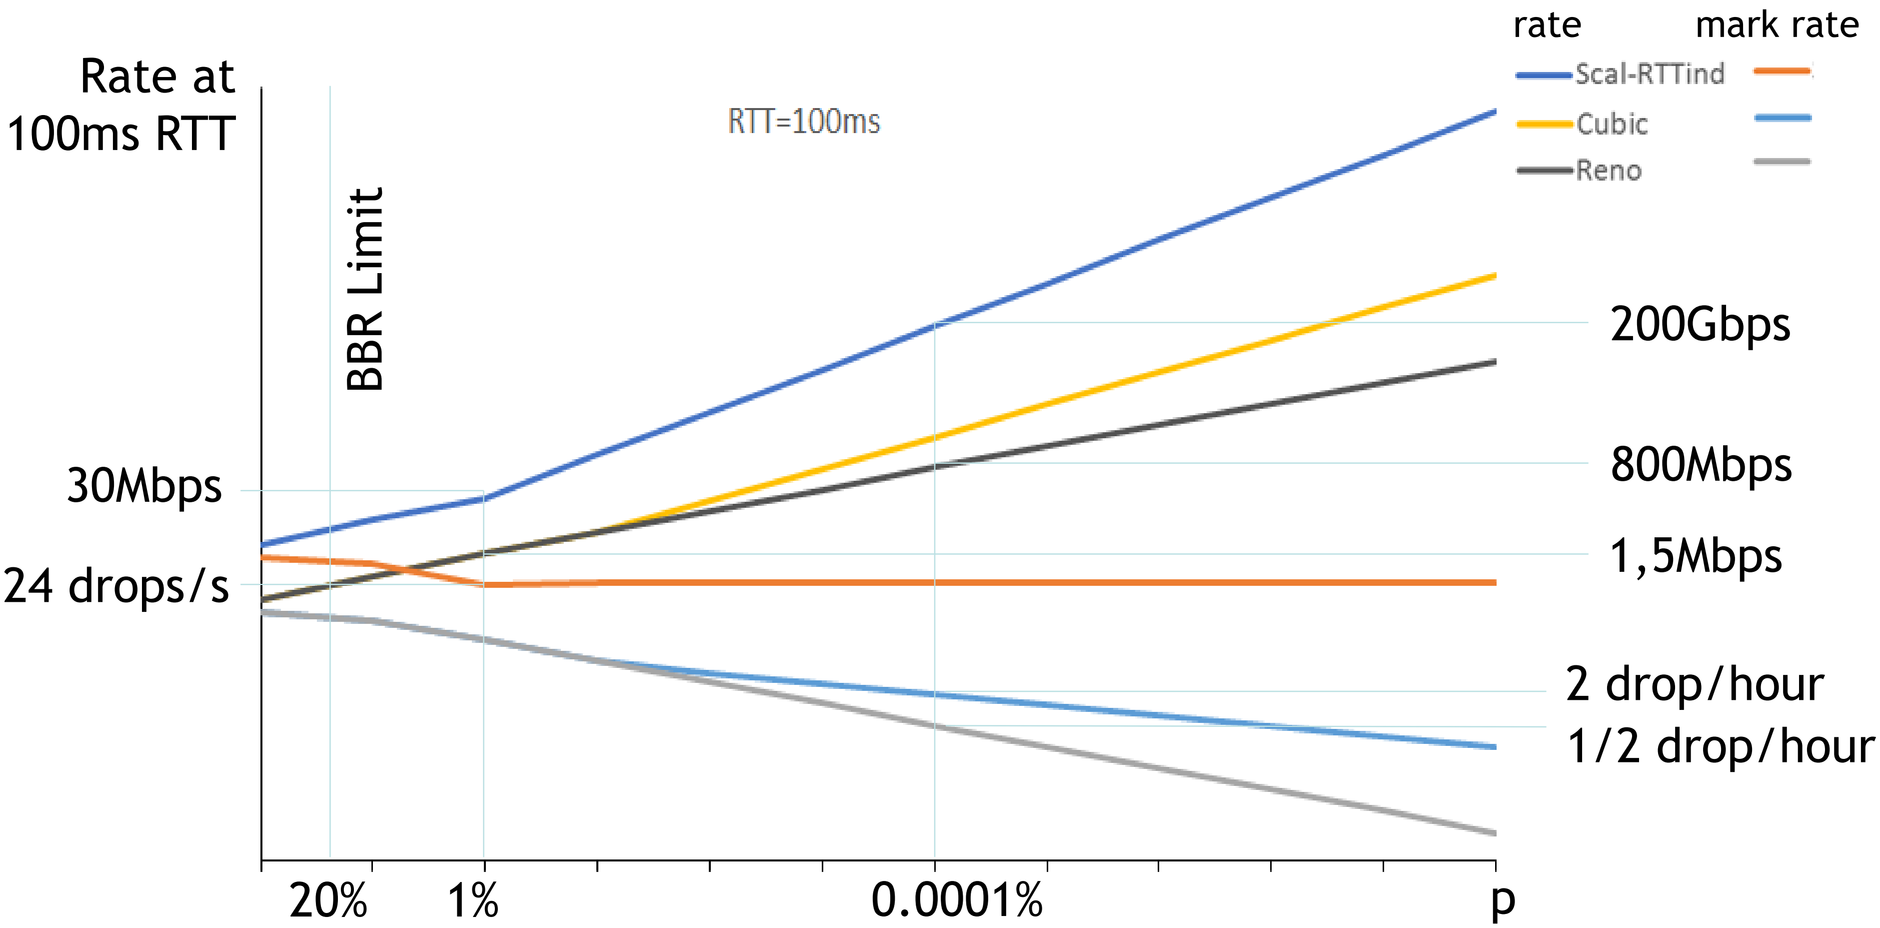

Resolving Tensions between Congestion Control Scaling Requirements

6 scalability requirements:

1. Scalable congestion signaling
2. Limited RTT-dependence
3. Unlimited responsiveness
4. Low relative queuing delay
5. Unsaturated signaling (previous talk)
6. Coexistence with Classic TCP

Link to paper:

Resolving Tensions Between Congestion Control Scaling Requirements

Link to experiment videos:

BBR with AQM: https://youtu.be/4eYfyKYe9nM

BBR with Cubic: https://youtu.be/akO1HN2ey48

Questions

koen_de_schepper@nokia-bell-labs.com