Demo: L4S in action

- Koen De Schepper, Inton Tsang, Bob Briscoe, Olga Bondarenko

  NOKIA Bell Labs [simula . research laboratory]

- 19-07-2016

R / T E

Reducing Internet Transport Latency

Demonstration

- Interactive demo/test GUI
- Compare:
  - PIE → State of the art single Q non-L4 AQM
  - DualPI2 → DualQ using PIE as Classic AQM

    → Linux open source version: https://github.com/olgabo/dualpi2

- DualQ is validated and compared with other AQMs and congestion controls in extensive tests on a DSL fixed access testbed
  - Identified safety & performance improvements for DCTCP
    → TCP-Prague requirements

Testbed

DSLAM=DSL Access Multiplexer
SR=Service Router
BNG=Broadband Network Gateway
RGW=Residential Gateway

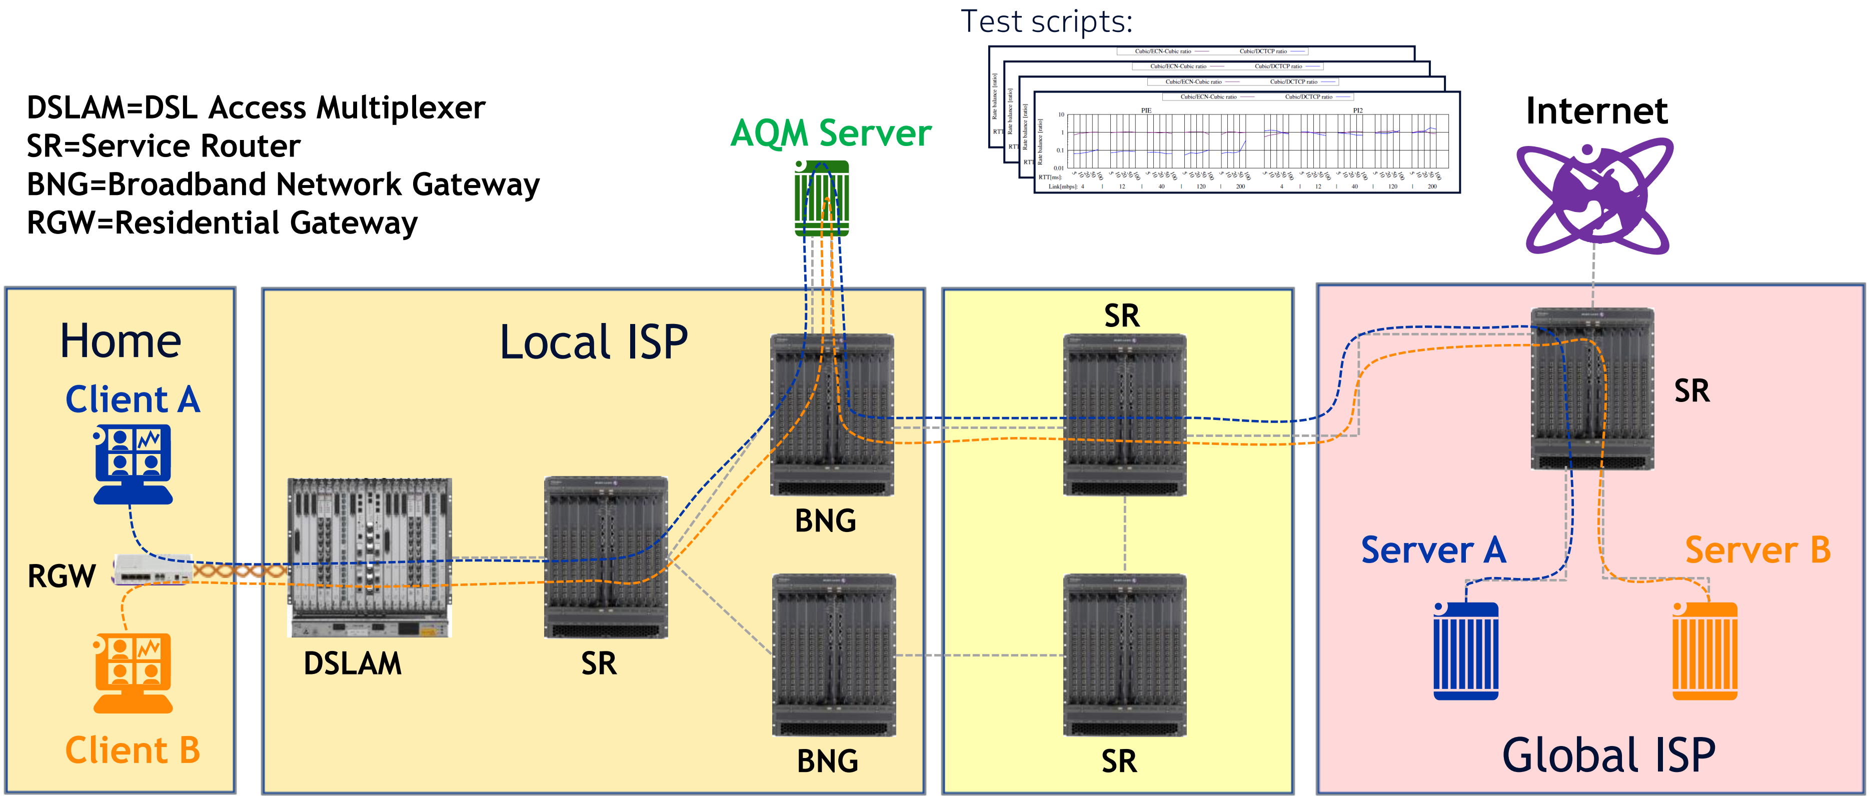

Demo/Test GUI

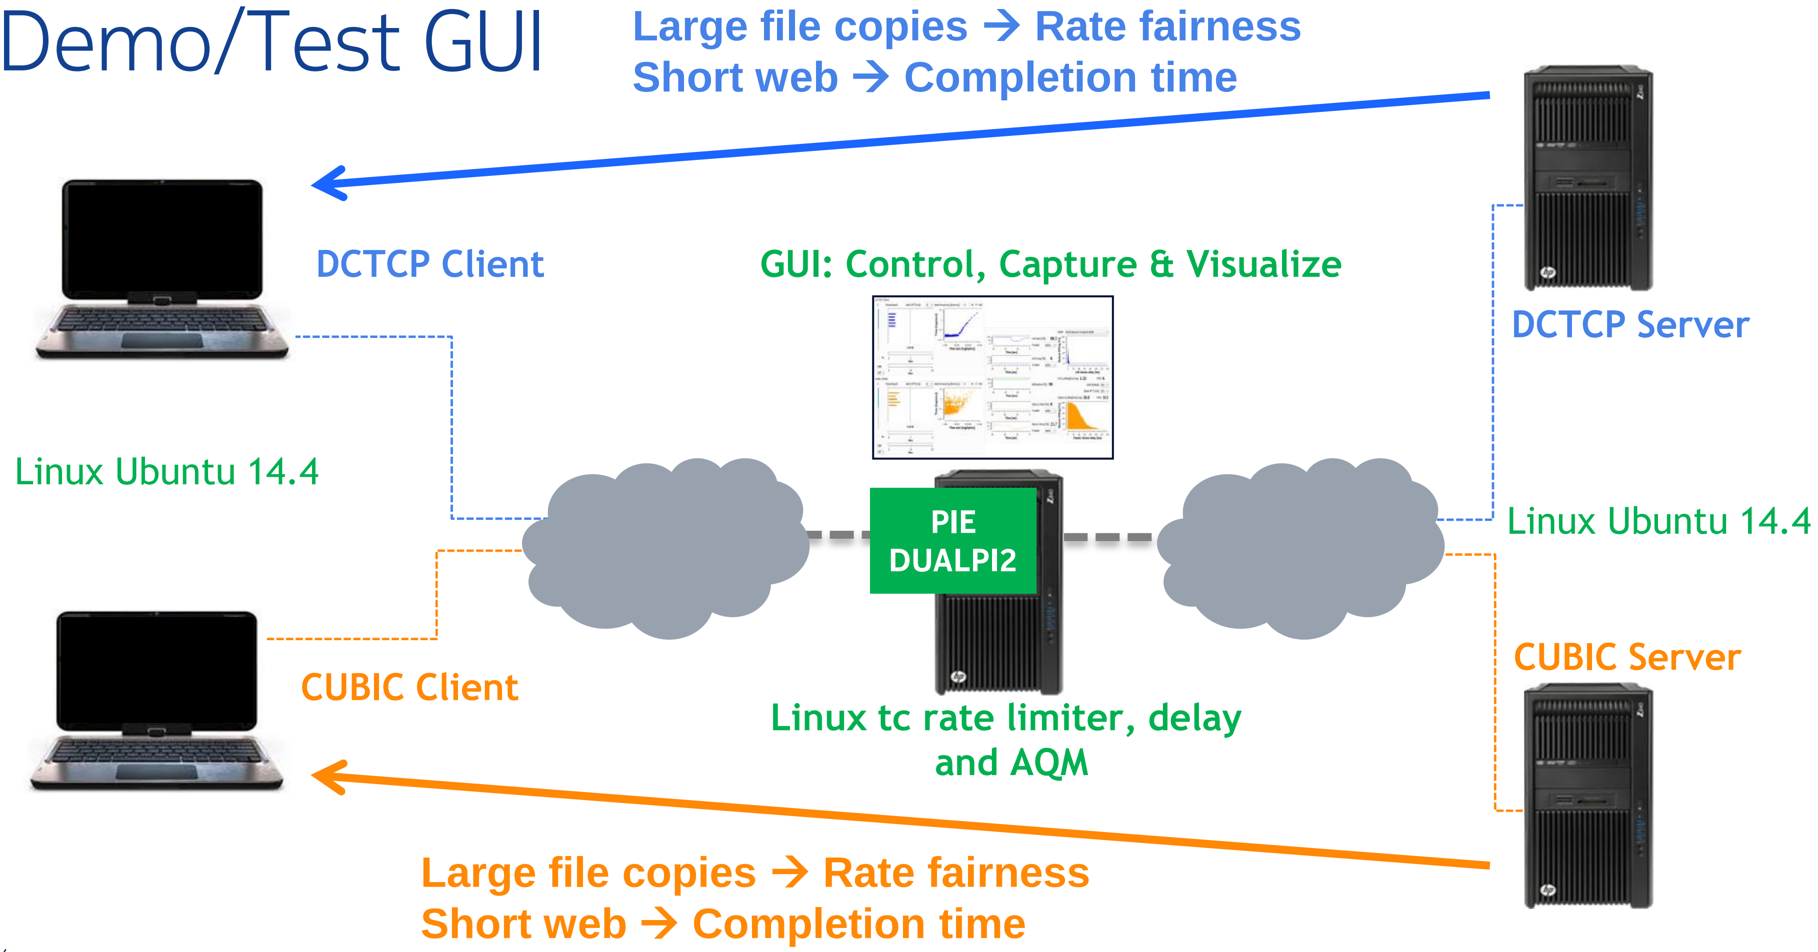

Demo GUI

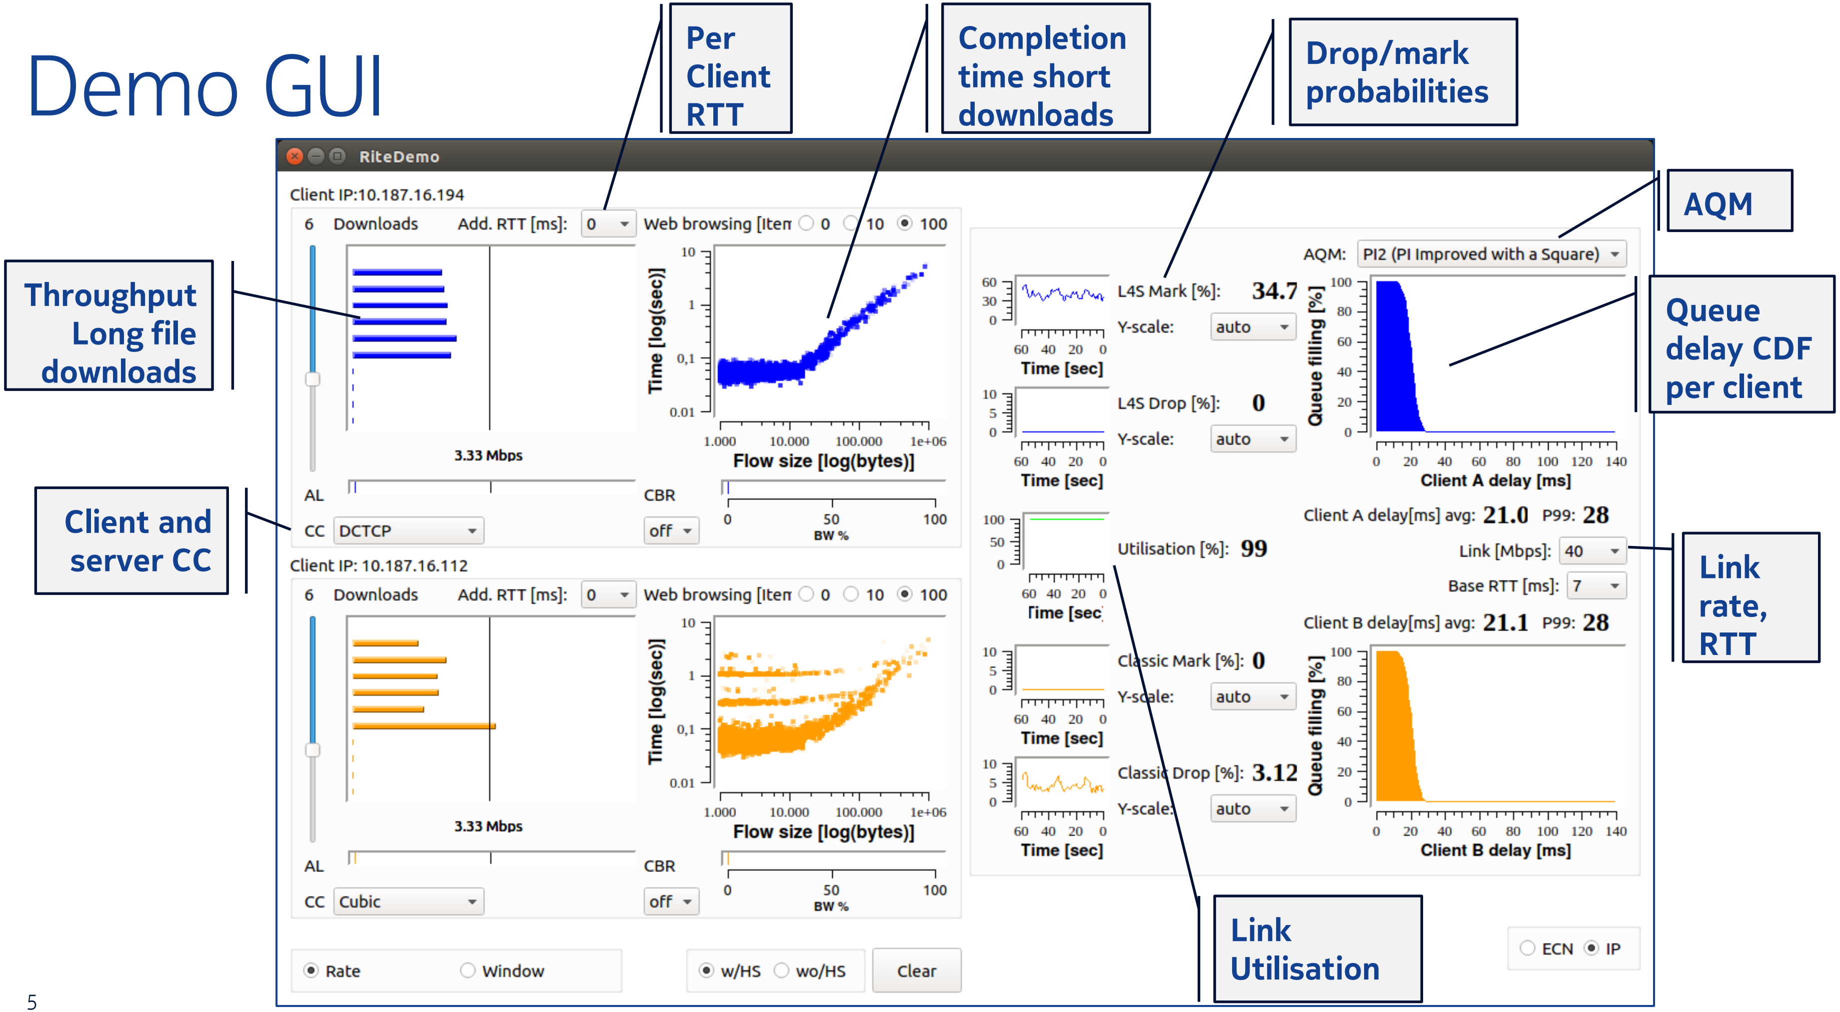

Per Client RTT
Completion time short downloads
Drop/mark probabilities
AQM
Throughput Long file downloads
Queue delay CDF per client
Client and server CC
Link rate, RTT
Link Utilisation
RiteDemo
Client IP:10.187.16.194
6 Downloads
Add. RTT [ms]: 0
Web browsing [Iten 0 10 100
3.33 Mbps
Time [log(sec)]
Flow size [log(bytes)]
AL
CBR
CC DCTCP
off
BW %
Client IP: 10.187.16.112
CC Cubic
AQM: PI2 (PI Improved with a Square)
L4S Mark [%]: 34.7
Y-scale: auto
L4S Drop [%]: 0
Time [sec]
Queue filling [%]
Client A delay [ms]
Client A delay[ms] avg: 21.0 P99: 28
Utilisation [%]: 99
Link [Mbps]: 40
Base RTT [ms]: 7
Client B delay[ms] avg: 21.1 P99: 28
Classic Mark [%]: 0
Classic Drop [%]: 3.12
Client B delay [ms]
Rate
Window
w/HS
wo/HS
Clear
ECN
IP

Access technologies evolve: new opportunities

- 5G, G.Fast, GPON, …
  - High throughputs
  - 1ms latency requirements

→ Classic TCP becomes a big bottleneck

→ L4S can exploit the lower latency without the classic compromises

- Nokia believes that standardization of L4S is important

Questions

koen.de_schepper@nokia.com

Backup

Reason for RTT independent TCP-Prague requirement

- One of the sixlemmas is that big queues enhance RTT fairness:

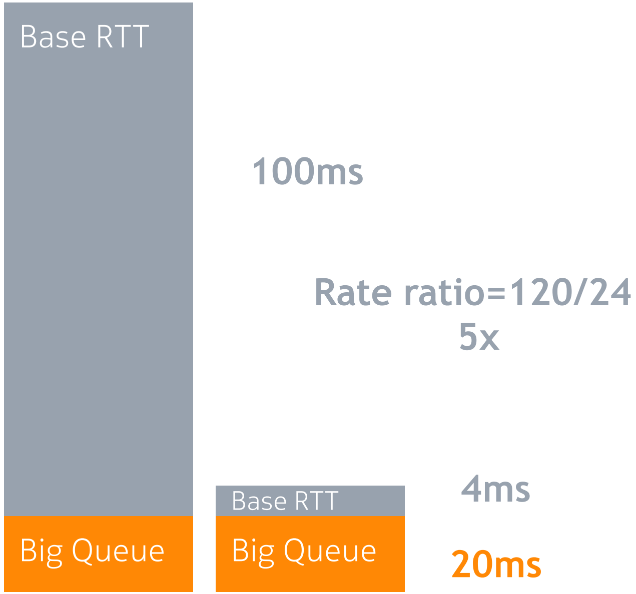

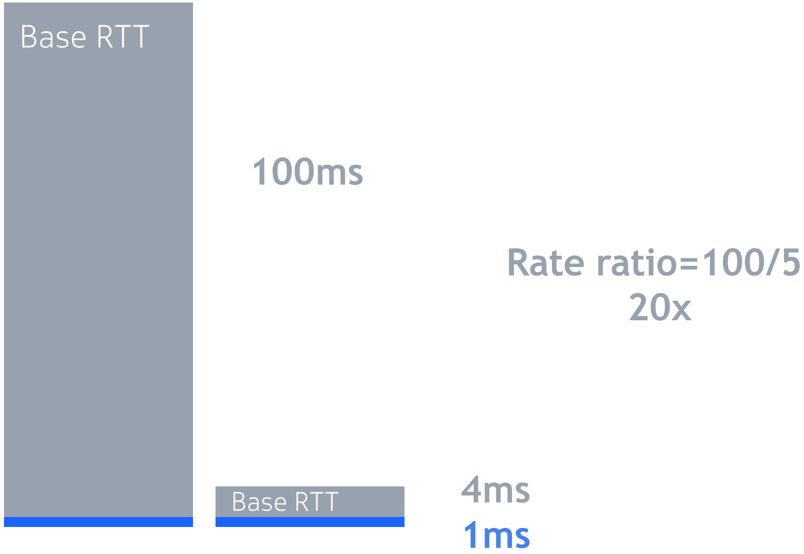

PIE autotune enhancement

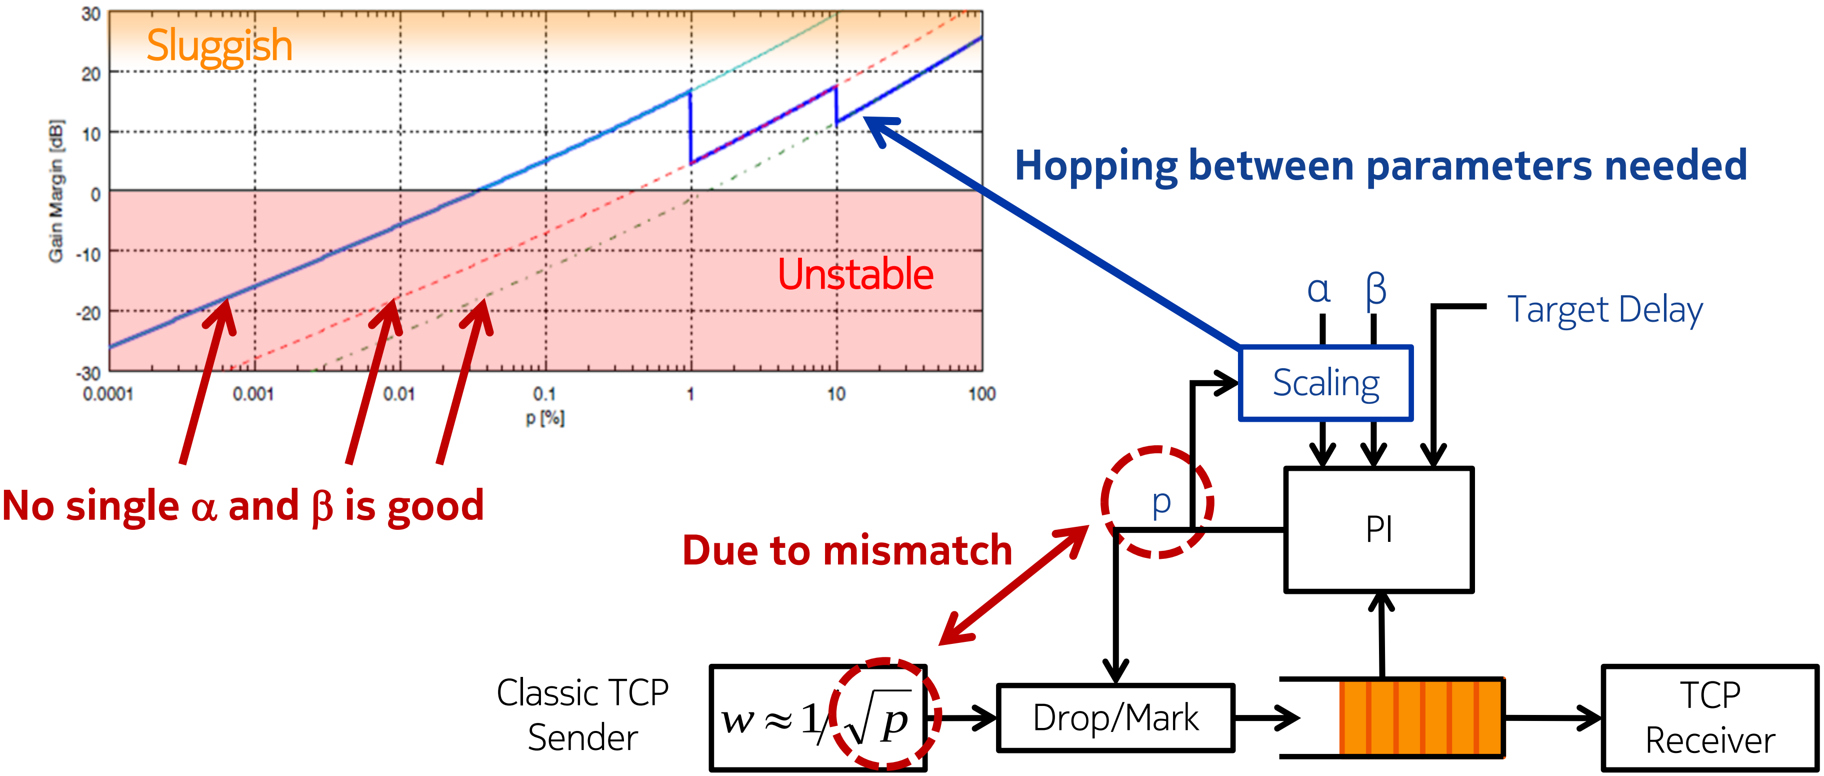

PI2: One PI to rule them all

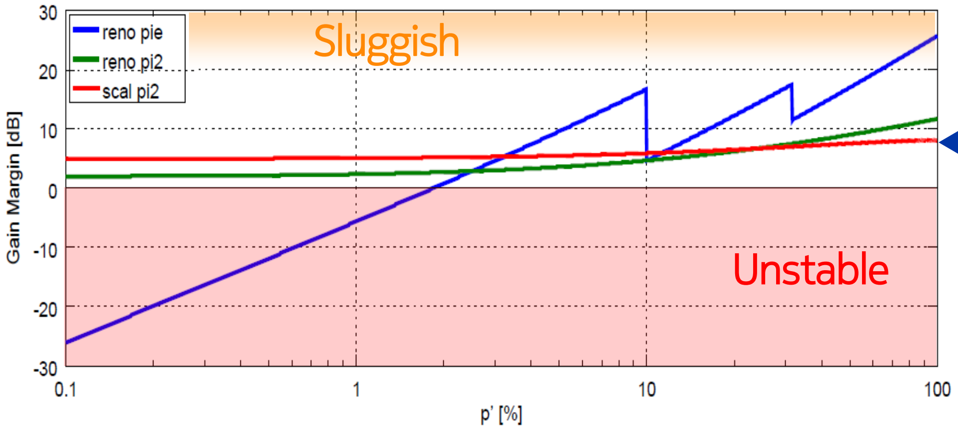

No hopping needed
Fixed parameters for $w \approx 1/p'$

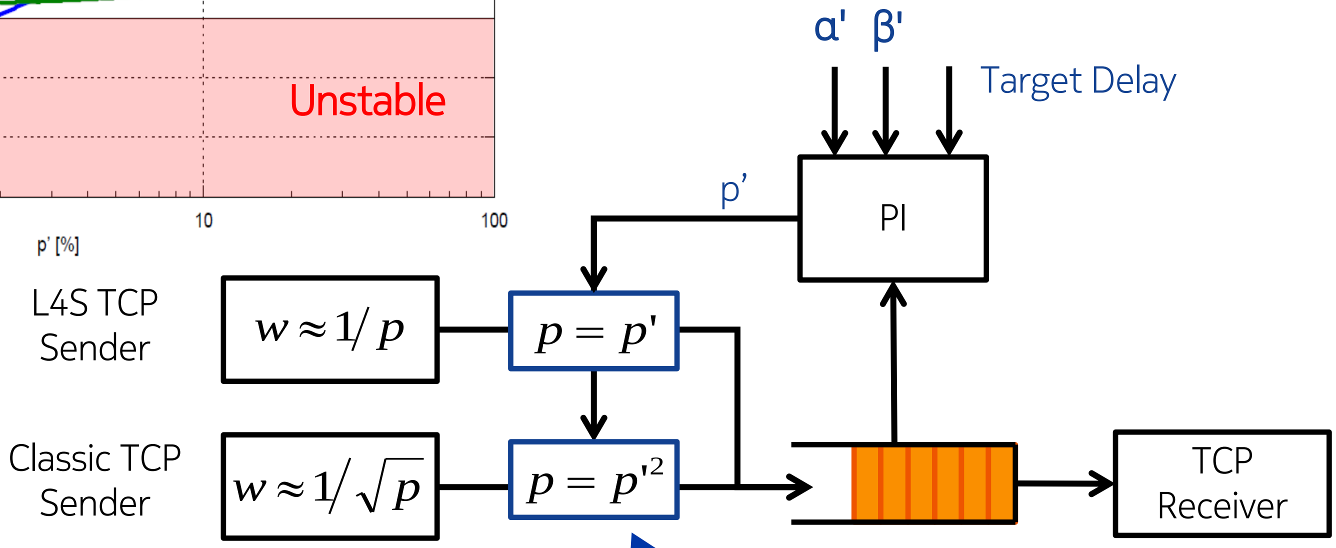

Turnes classic into $w \approx 1/p'$

Dual Queue Coupled AQM

DualQ to preserve low latency for L4S traffic

L4S (DCTCP)

$$r \approx 1/p$$

Classic (Reno / Cubic)

$$r \approx 1/\sqrt{p}$$

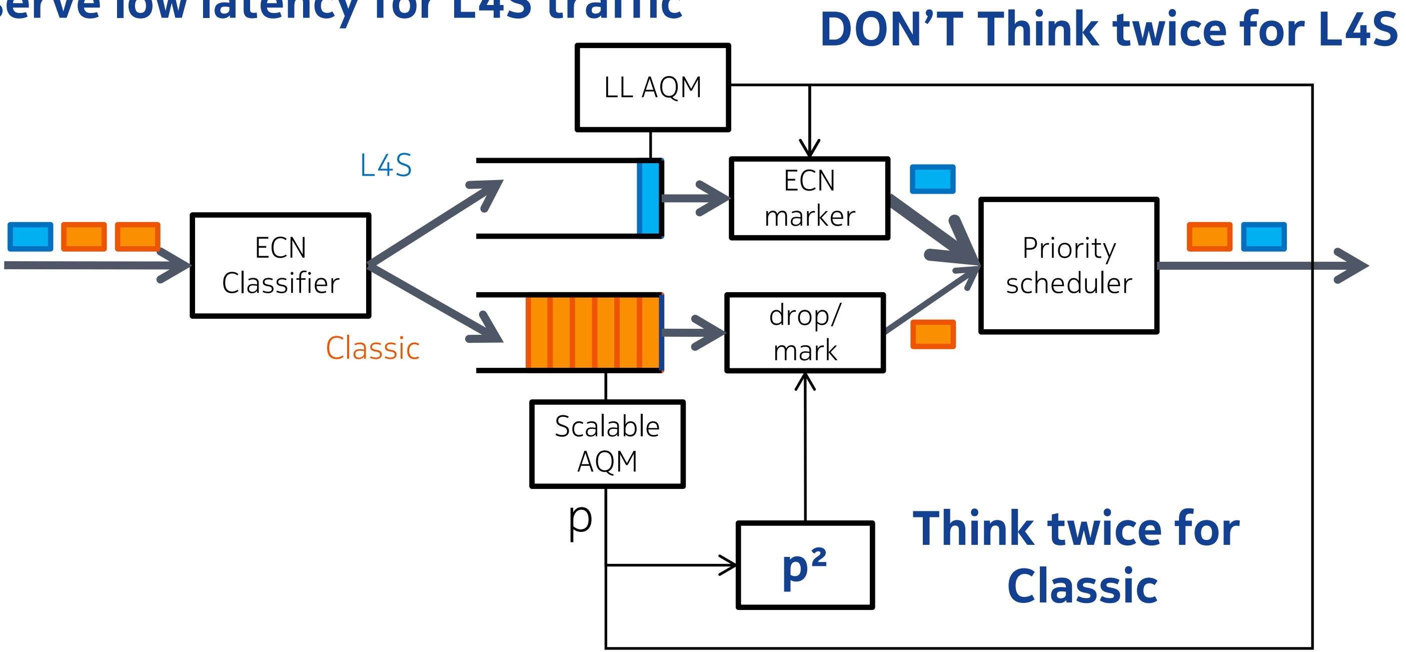

Coupled AQM for equal rate

Demo with Linux DualPI2

L4S (DCTCP)

$$r \approx 1/p$$

Classic (Cubic)

$$r \approx 1/\sqrt{p}$$

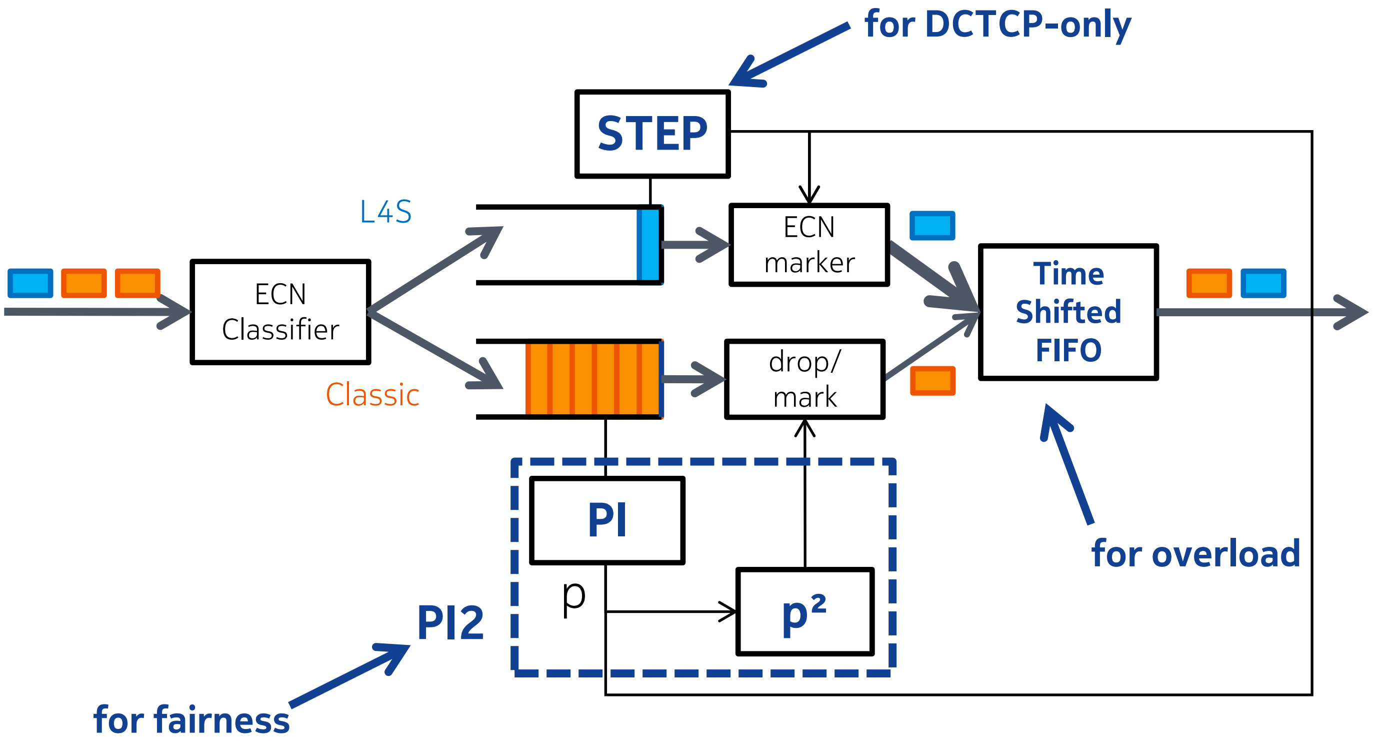

Demo experiments: coexistence

DCTCP $p = p'$

Cubic $p = \left(\frac{p'}{2}\right)^2$

- PI2 DCTCP and Cubic
  - Fairness: same throughput:
  - different (equal) RTTs and link speeds

1-1 flows

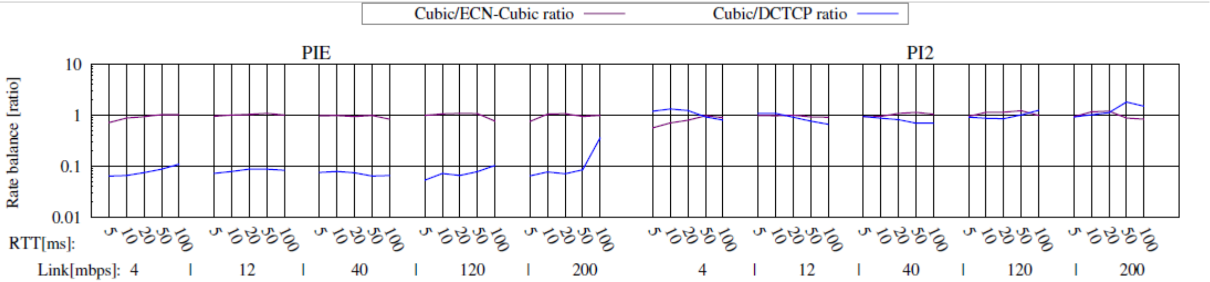

  - different numbers of flows

40Mbps

10ms

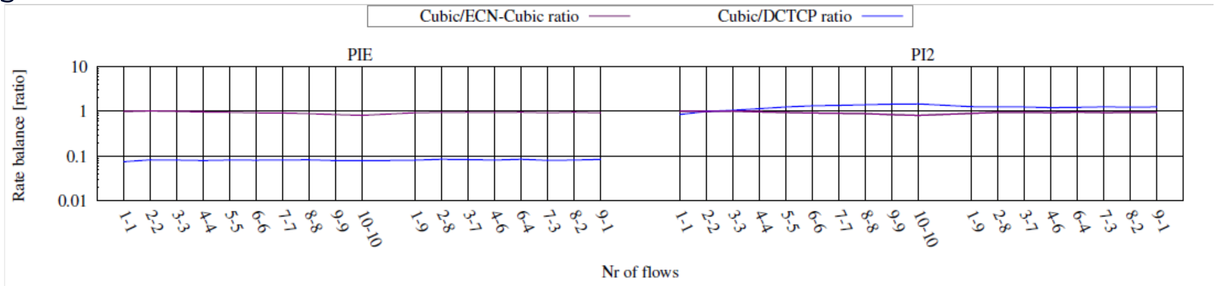

TCP-Prague:
Compensate for the advantages of big queue targets

Big buffers are today a network solution for Classic TCP limitations

L4S allows TCP-Prague to solve problems in the end-point

→ TCP-Prague should remove as much as possible the compromises of defining shallow thresholds

Demo experiments: ECN, DCTCP

- Effect of classic ECN and L4S ECN on all packets
- PIE CUBIC-drop ->

PIE CUBIC-ECN ->

PI2 DCTCP

- Low latency queue
- STEP DCTCP