L4S

A gradually deployable simplifying clean-slate opportunity

Bob Briscoe
[simula . research laboratory]

Koen De Schepper
NOKIA Bell Labs

April 6th, 2016

Know the competition

The benefits of big standing queues:

- Keeps links fully utilized, by keeping backup packets for the dip in the big saw-teeth

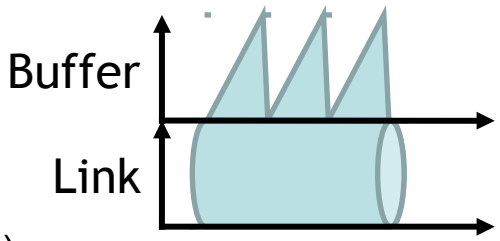

- Allows bursts by delaying congestion response (AQM smoothing)
- Reduces the loss rate, by increasing RTT

$$r \approx \frac{1.22}{\sqrt{p}.RTT}$$

- Reduces RTT dependency on throughput, by equalizing RTT
- Absorbs minimum window sizes

  Every flows emits at least 2 packets:

  if N*2p > BDP then queuing is unavoidable

$$\frac{100ms + 0}{1ms + 0} = 100\times$$ No queue

$$\frac{100ms + 100}{1ms + 100} = 2\times$$ Big queue

Alternative solutions ?

Removing queues for ultra low latency needs alternative solutions

End-to-End principle: Try in end system first (a big buffer was the network solution)

- Reduce saw-teeth
- Remove NW smoothing, smooth in the TCP sender
- Use ECN, support lots of signaling AND scalability
- Reduce RTT dependency on throughput in the end-system
- Remove minimum window sizes

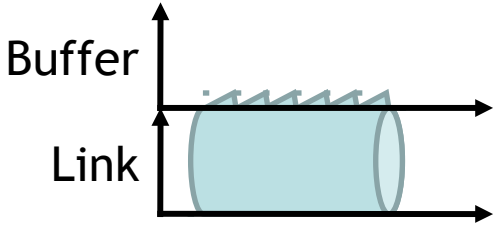

$$p.r \approx C$$

$$r \approx \frac{x}{p.\cancel{RTT}}$$

L4S and DualQ

a gradual deployable

simplifying

clean-slate

opportunity!

ect(1) reuse of ECN

think twice: coupling rule for marking

simplifies classic AQM control