

L4S: Low Latency, Low Loss, Scalable Throughput Internet Service Architecture and Identifier

draft-ietf-tsvwg-l4s-arch-01
draft-ietf-tsvwg-ecn-l4s-id-01

Bob Briscoe, CableLabs
Koen De Schepper, Nokia Bell Labs
Marcelo Bagnulo, UC3M

IETF-100 Nov 2017


The authors were part-funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700). The views expressed here are solely those of the authors.

Recap

- Motivation

- Extremely low queuing delay for all Internet traffic
- already 1-2 orders better than state of the art
- 500 μ s vs 5-15 ms (fq-CoDel or PIE)


- Architecture



tsvwg-l4s-arch

tsvwg-ecn-l4s-id

status




- Now in a holding pattern
 - pending possible changes to main assumptions about the “TCP Prague Requirements”
 - Exit holding pattern when reference implementation of TCP Prague complete
- Reviews, comments from implementers etc, obviously welcome

L4S Status Update (1/3)


- Source Code
 - Dual Queue Coupled AQM, DualPI2 for Linux [UPDATED internally, release shortly]
 - Data Centre TCP (DCTCP) for Linux (in the mainline kernel), FreeBSD patch, ns2 patch.
 - Accurate ECN TCP Feedback for Linux [UPDATED, but still not fully tested]
- Implementations
 - DualQ Coupled AQM: in at least one chipset aimed at the data centre environment [**NEW – availability TBA**]
 - L4S Scalable congestion control: rmcats SCReAM [**NEW**]
 - Whole L4S system: ns3 [**to be released early 2018**]

SCReAM, 100Mbps , CoDel ECN


qdel[s]




RTP queue delay



CWND & in flight [byte]




Bitrate [kbps]




SCReAM, 100Mbps , L4S


qdel[s]



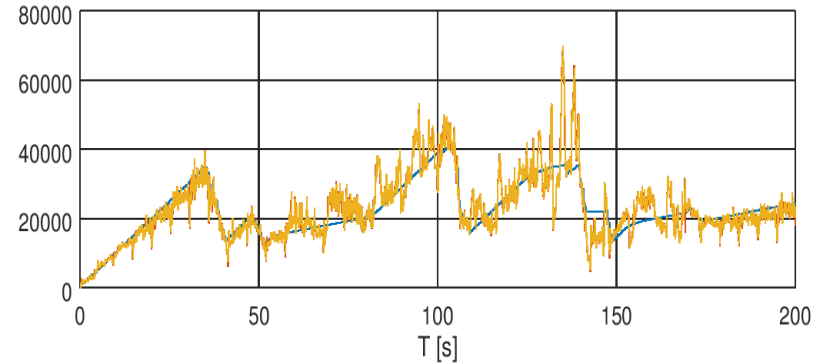
RTP queue delay



CWND & in flight [byte]



Bitrate [kbps]



L4S Status Update (2/3)

- IETF specs
 - Low Latency, Low Loss, Scalable Throughput (L4S) Internet Service: Architecture <draft-ietf-tsvwg-l4s-arch-01> [MINOR UPDATE]
 - A proposed new identifier for Low Latency, Low Loss, Scalable throughput (L4S) packets <draft-ietf-tsvwg-ecn-l4s-id-01> [MINOR UPDATE]
 - Dual-queue AQM: : <draft-ietf-tsvwg-aqm-dualq-coupled-02> [UPDATED]
 - enabled by <draft-ietf-tsvwg-ecn-experimentation> [Passed IESG; UPDATED]
 - scalable TCP algorithms, e.g. Data Centre TCP (DCTCP) <RFC8257>, TCP Prague [RFC PUBLISHED]
 - Accurate ECN: <draft-ietf-tcpm-accurate-ecn-05> [UPDATED – WGLC LIKELY]
 - ECN++ Adding ECN to TCP control packets: <draft-ietf-tcpm-generalized-ecn-02> [UPDATED]
 - ECN support in trill <draft-ietf-trill-ecn-support-03>, motivated by L4S [Completed WGLC]
 - ECN in QUIC <draft-johansson-quic-ecn-03>, motivated by L4S [DES TEAM FORMED]
- 3GPP Proposal
 - ECN visibility to Radio Link Control (RLC) layer, motivated by L4S [Rejected for R15; Retry for R16]

L4S Status Update (3/3)

- Papers
 - Article in the IETF Journal describing the Demo in Bits-N-Bites at the IETF in Prague, July 2015. “Ultra-Low Delay for All” IETF Journal, Nov 2015.
 - “Ultra-Low Delay for All: Live Experience, Live Analysis“, Proc. ACM Multimedia Systems; Demo Session (May 2016).
 - “PI2: A Linearized AQM for both Classic and Scalable TCP,” Proc. ACM CONEXT 2016 (To appear Dec 2016).
 - “Data Centre to the Home: Deployable Ultra-Low Queuing Delay for All” [Journal submission]
- Landing page for code, specs, papers
<https://riteproject.eu/dctth/>

DualQ Coupled AQM

for Low Latency Low Loss Scalable throughput (L4S)

draft-ietf-tsvwg-aqm-dualq-coupled-02

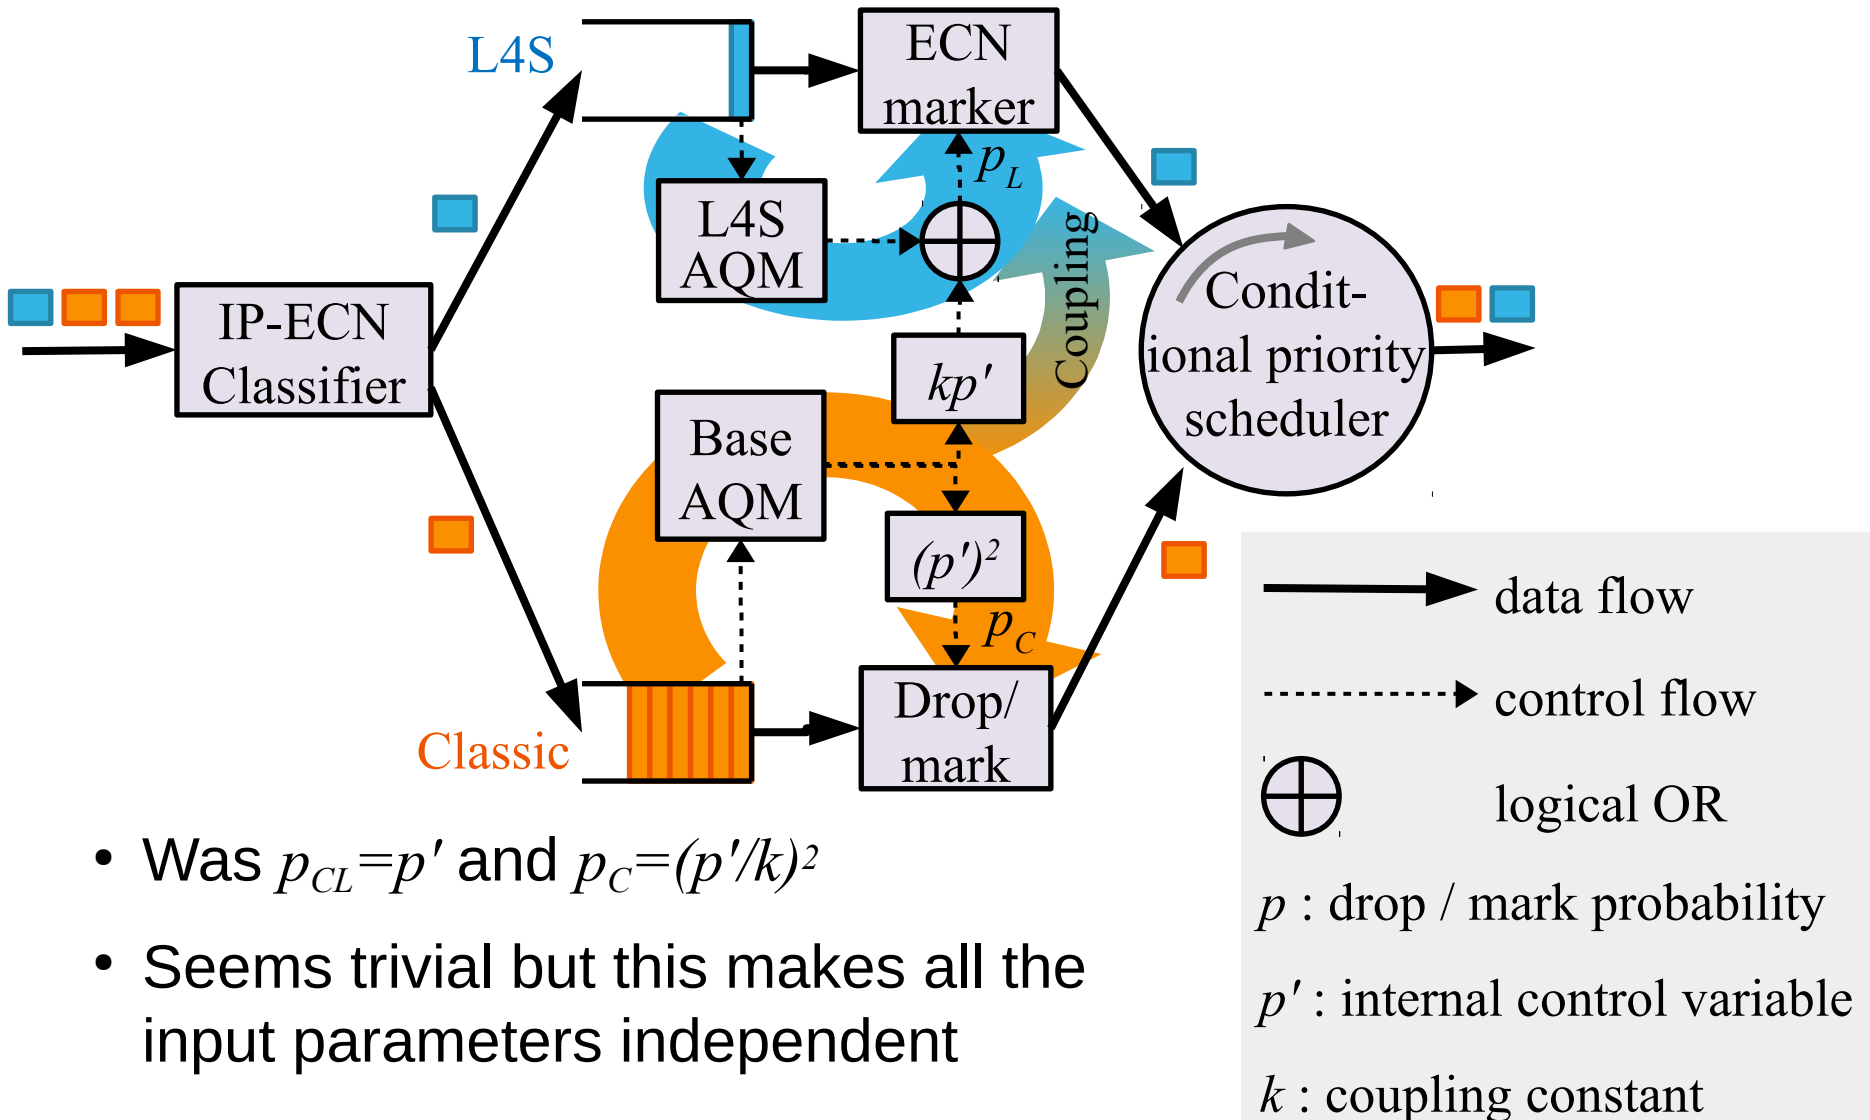
Bob Briscoe, CableLabs

Koen De Schepper, Nokia Bell Labs

IETF-100 Nov 2017

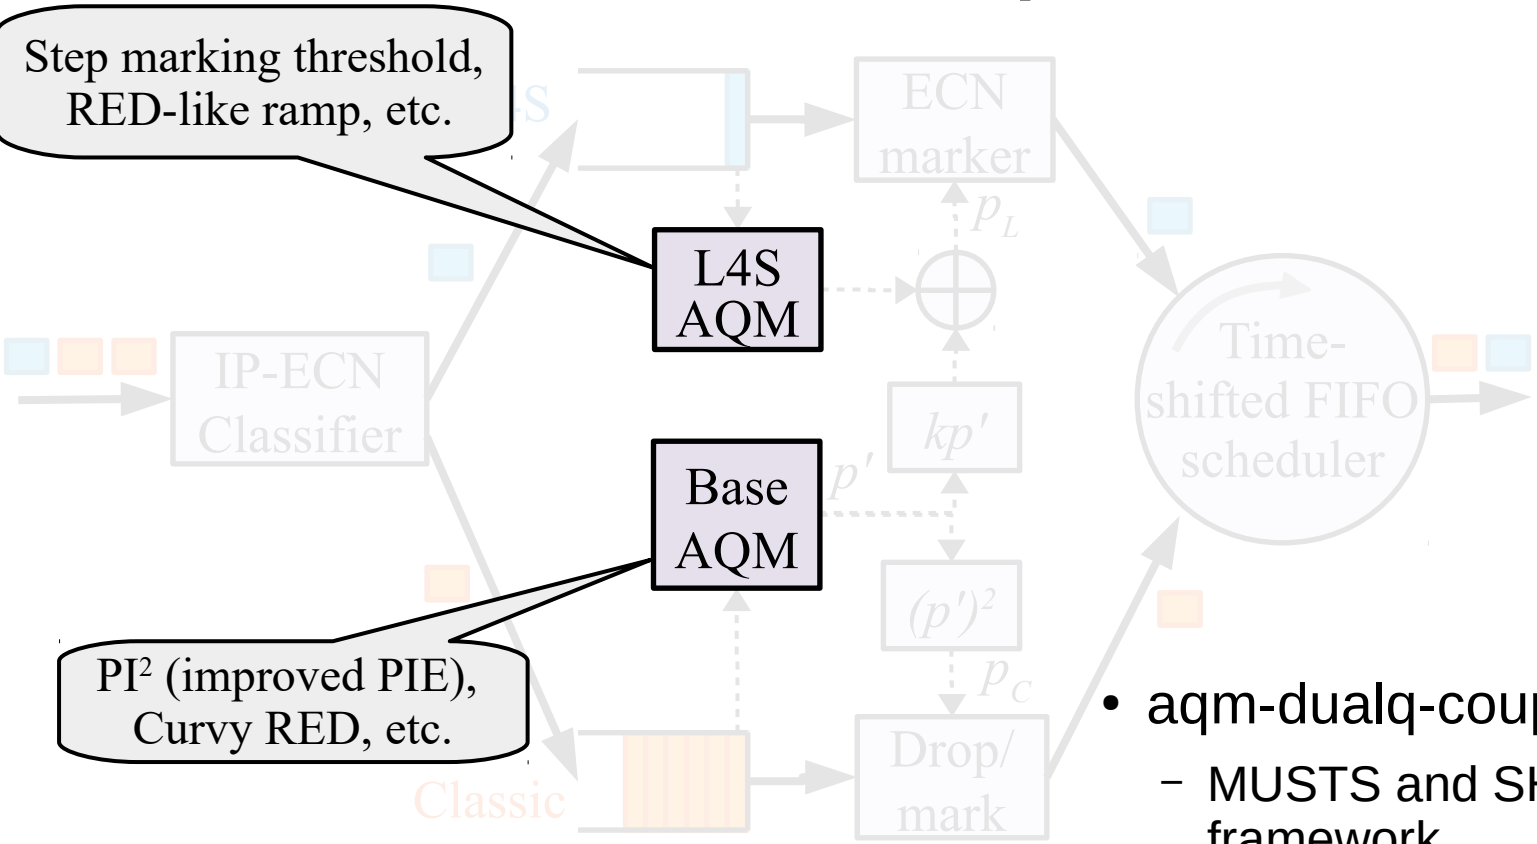
The authors were part-funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700). The views expressed here are solely those of the authors.

DualQ Coupled AQM: Structure



- Was $p_{CL} = p'$ and $p_C = (p'/k)^2$
- Seems trivial but this makes all the input parameters independent

DualQ Coupled AQM: modular separation



- not actually *an* AQM
 - a framework for coupling AQMs
 - drop in your chosen 'native AQM' for each queue

- aqm-dualq-coupled defines:
 - MUSTS and SHOULDs for framework
 - informational appendices explaining 2 example AQMs with pseudocode

DualQ Coupled AQM: Management Requirements

- Configuration
 - 3 categories:
 - 1) Framework
 - 2) Base AQM
 - 3) Native L4S AQM
- Monitoring
 - per queue per sample interval
 - Bits forwarded
 - utilization
 - Packets arriving; enqueued; dequeued
 - packets marked; dropped (ECT & not-ECT)

- Linux ref. implementation
 - Fixed the API so each config parameter is independent
 - Added external config of classifier (Linux classifier arch)
 - Made classful: visibility of each queue (e.g. stats) consistent externally

Status & Next Steps

- Draft is now understandable to implementers
 - thanks to Sowmini Varadhan & Gabi Bracha

Next:

- Management requirements for experiments
 - prompted by IESG review of tsvwg-ecn-experimentation
- Release changes to Linux implementation
- Relationship to Diffserv
- Policing / queue protection

draft-ietf-tsvwg-l4s-arch-01

draft-ietf-tsvwg-ecn-l4s-id-01

draft-ietf-tsvwg-aqm-dualq-coupled-02

Q&A