L4S: Low Latency, Low Loss, Scalable Throughput Internet Service

draft-ietf-tsvwg-l4s-arch-02
draft-ietf-tsvwg-ecn-l4s-id-02
draft-ietf-tsvwg-aqm-dualq-coupled-04

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

IETF-101 Mar 2018

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 µs vs 5-15 ms (fq-CoDel or PIE)
- Architecture

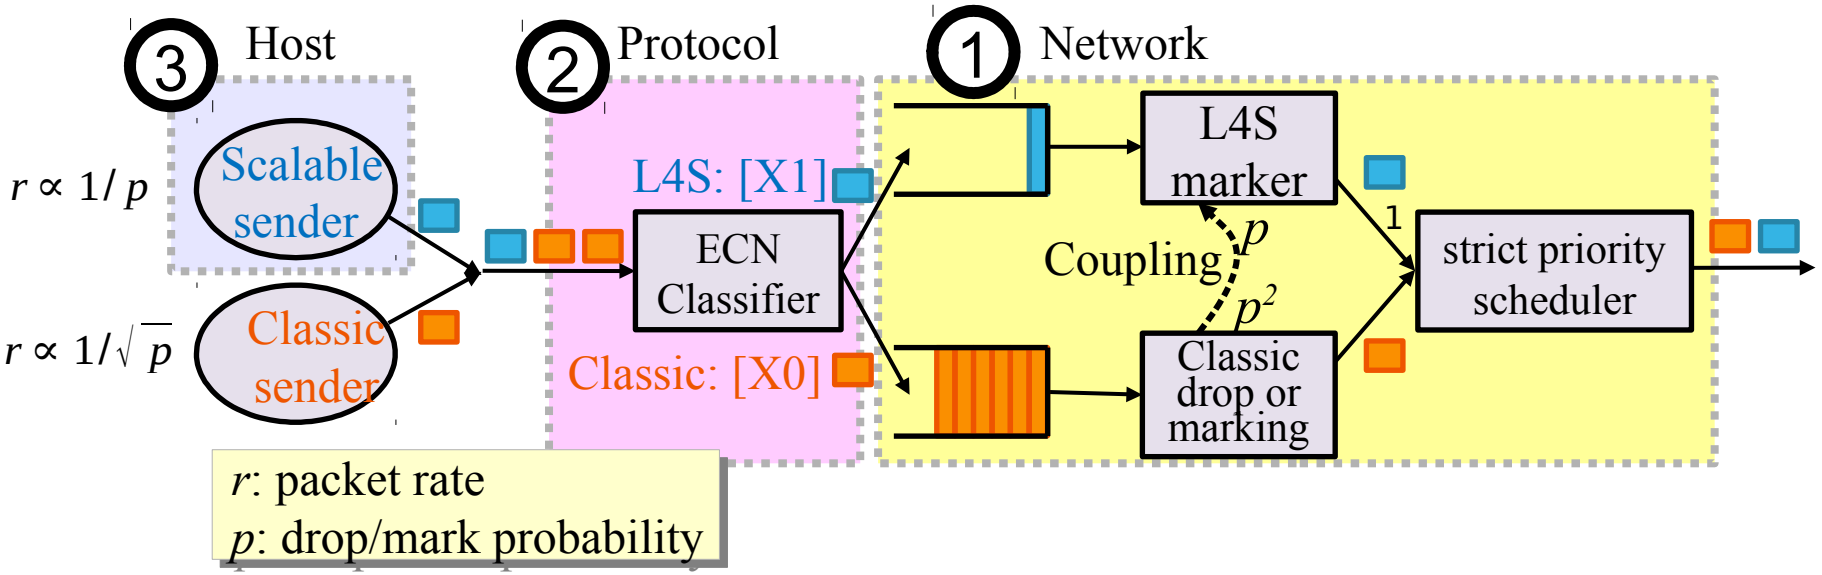

L4S Status Update (1/2)

- Landing page for code, specs, papers

  https://riteproject.eu/dctth/

- Source Code
  - Dual Queue Coupled AQM, DualPI2 for Linux [UPDATE in progress]
  - Data Centre TCP (DCTCP) for Linux (in the mainline kernel), FreeBSD patch, ns2 patch.
  - Accurate ECN TCP Feedback for Linux [testing needed]
- Implementations
  - DualQ Coupled AQM: in at least one chipset aimed at the data centre environment [availability TBA]
  - L4S Scalable congestion control: rmcat SCReAM
  - BBRevo, evolution of BBR with L4S support [NEW, see iccrg Fri]
  - Whole L4S system in ns3 [complete but evolving, first release Jun'18 timeframe]

L4S Status Update (2/2)

- IETF specs
  - Low Latency, Low Loss, Scalable Throughput (L4S) Internet Service: Architecture <draft-ietf-tsvwg-l4s-arch-02> [MINOR UPDATE]
  - A proposed new identifier for Low Latency, Low Loss, Scalable throughput (L4S) packets <draft-ietf-tsvwg-ecn-l4s-id-02> [MINOR UPDATE]
  - Dual-queue AQM: : <draft-ietf-tsvwg-aqm-dualq-coupled-04> [2 UPDATES] see later
  - Interactions of L4S with Diffserv <draft-briscoe-tsvwg-l4s-diffserv-00> [NEW] see later
  - enabled by <RFC8311> [RFC published]
  - scalable TCP algorithms, e.g. Data Centre TCP (DCTCP) <RFC8257>, TCP Prague
  - Accurate ECN: <draft-ietf-tcpm-accurate-ecn-06> [UPDATED – WGLC pending rvw(s)]
  - ECN++ Adding ECN to TCP control packets: <draft-ietf-tcpm-generalized-ecn-02> [Supporting measurement paper published in IEEECommMag]
  - ECN support in trill <draft-ietf-trill-ecn-support-07>, motivated by L4S [4 updates, RFC Ed Q]
  - 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]

DualQ Coupled AQMs for L4S
draft-ietf-tsvwg-aqm-dualq-coupled-04

- Two updates in this IETF cycle
- Overload handling
  - explained under security considerations
    (not just pseudocode in appendix A.2)
- Additional terminology for the control variables
- Un-deprecated WRR for inter-Q scheduler
- Added classifier flexibility, not only ECT(1):
  - addressing
  - protocols (e.g. DNS, LDAP, ARP)
  - or DSCP (see new draft – next slides)

Interactions between
Low Latency Low Loss Scalable throughput (L4S) and Diffserv

draft-briscoe-tsvwg-l4s-diffserv-00

Bob Briscoe, CableLabs

IETF-101 Mar 2018

Why is DualQ Coupled AQM different?

- Diffserv controls bandwidth
  - controls queue latency by allocating bandwidth
- Coupled DualQ: semi-permeable membrane:
  - latency: delay of L queue isolated[1] from C
  - bandwidth: behaves as 1 pool of capacity
- Can add bandwidth allocation to DualQ
  - but typically unnecessary

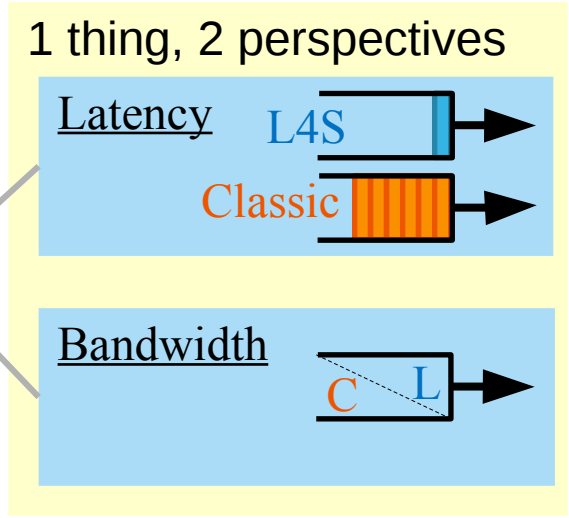

[1]: Lower latency is not at the expense of the C queue, C is for legacy

L4S-Diffserv: 4 types of interaction

1) None
   - typical
2) Codepoint mapping (next slide)
   - in absence of additional Diffserv queues
3) Operator configures bandwidth alloc'n around DualQ
   - no different from any AQM + Diffserv
4) Operator configures bandwidth alloc'n within DualQ

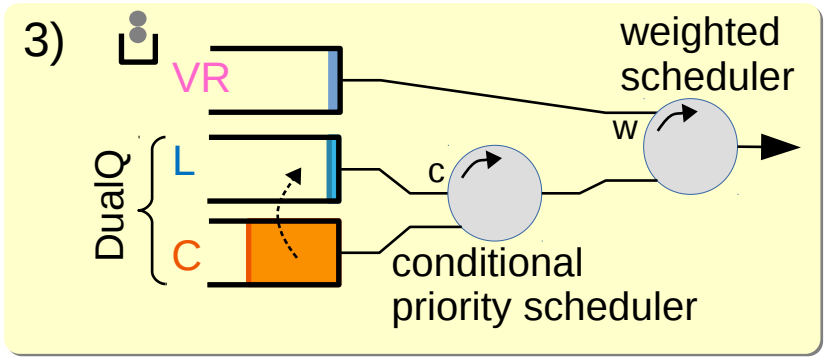

Assured b/w for demanding app, e.g. VR

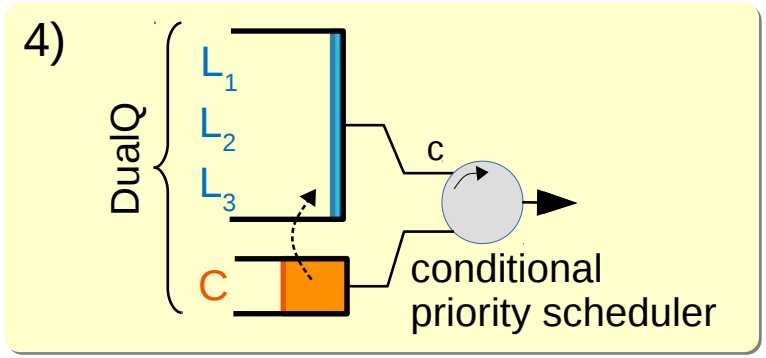

Low latency for L_1-L_3
More b/w per flow for L_1 than L_3
(like Diffserv Assured Forwarding)

Mapping Diffserv Service Classes [RFC4594] to L4S (if operator solely offers Latency & Classic queues)

Service Class Name	DSCP Name	DSCP Value	App example	AQM	LLD
Network Control	CS7	111000	(Resv'd for) Network routing	Y & N	L if ECT1
Network Control	CS6	110000	Internetwork routing	Y & N	L if ECT1
OAM	CS2	010000	Ops, admin, mgmt & provis'ng	Y & N	L if ECT1
Signalling	CS5	101000	IP telephony signalling	N	L
Telephony	EF	101110	IP telephony bearer	N	L
	Voice Admit[1]	101100	Admission-control'd IP telephony	N	L[1]
Real-Time Interactive	CS4	100000	Video conf & interactive gaming	N	L if ECT1
MM Conferencing	AF4x; x=1,2,3	100{01,10,11}0	H.323/V2 video conf. (adaptive)	Y	L if ECT1
Broadcast Video	CS3	011000	Broadcast TV & live events	N	L if ECT1
Multimedia Streaming	AF3x; x=1,2,3	011{01,10,11}0	Streamed video & audio	Y	L if ECT1
Low Latency Data	AF2x; x=1,2,3	010{01,10,11}0	Client-server transactions, Web	Y	L if ECT1
High Thrughput Data	AF1x; x=1,2,3	001{01,10,11}0	Store and forward applications	Y	L if ECT1[2]
Standard	DF (CS0)	000000	Undifferentiated applications	Y	L if ECT1
Low Priority Data	LE[3]	000001[3]	Any flow with no b/w assurance	Y	L if ECT1[4]

- “L if ECT1” is not classifed into L 'cos of its DSCP
- Need to consider NTP (advice in RFC4594 n/a for LLD)

1. RFC5865 gives Voice Admit priority over EF
2. To take advantage of scalable congestion control
3. Less Effort [draft-ietf-tsvwg-le-phb] update to RFC4594
4. Flows using LE SHOULD also use LE congestion ctrl

Status & Next Steps

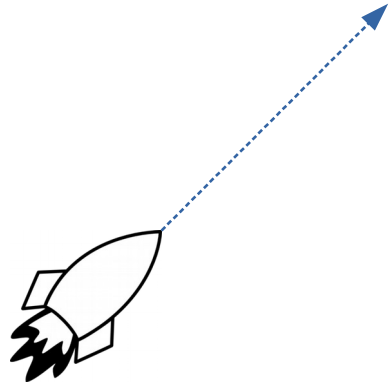

- AQM: **aqm-dualq-coupled**: continual improvement
  - feedback from implementers (all offlist)
  - evaluation over other specific links besides DSL (DOCSIS, LTE/5G, DC) using ns3, Linux and prototype h/w implementations
- Architecture & Identifier: in holding pattern
  - **l4s-arch, ecn-l4s-id**

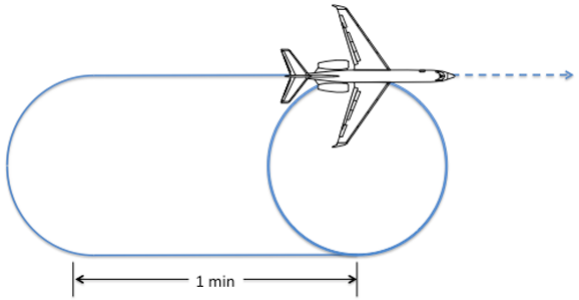

Next:

- TCP Prague: Pulling parts together
- **aqm-dualq-coupled**: more on policing / queue protection

- Review of relationship with Diffserv
- **draft-briscoe-tsvwg-l4s-diffserv** adoption?
  - loosely coupled to rest of L4S process?

draft-ietf-tsvwg-l4s-arch
draft-ietf-tsvwg-ecn-l4s-id
draft-ietf-tsvwg-aqm-dualq-coupled
draft-briscoe-tsvwg-l4s-diffserv

Q&A