To adopt or not to adopt L4S-compatible congestion control?

NYU TANDON SCHOOL OF ENGINEERING · NYU WIRELESS

Fatih Berkay Sarpkaya, Ashutosh Srivastava, Fraida Fund, Shivendra Panwar

What is L4S?

Low Latency, Low Loss, Scalable Throughput (L4S) [3] architecture is designed to enable **both high throughput and low latency** while coexisting with classic flows.

Classic congestion control cannot achieve extremely low delay with high throughput.

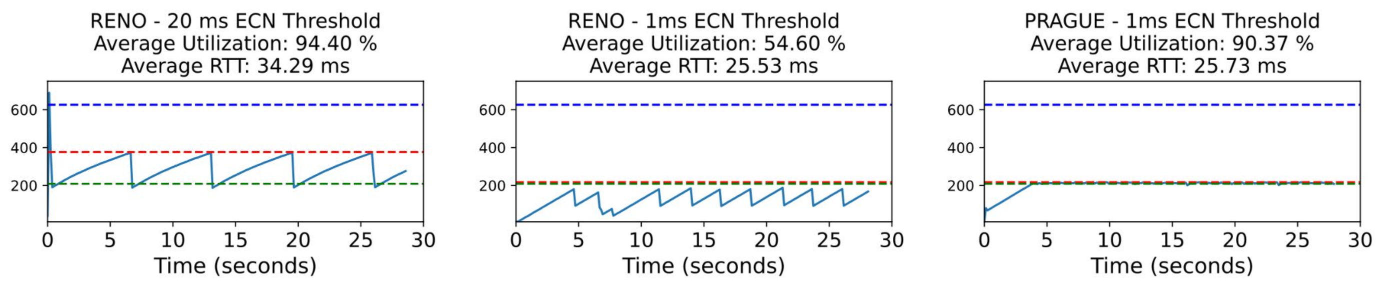

Classic architecture: High ECN threshold, high delay... | ...OR small ECN threshold, low throughput | **L4S**: Small ECN threshold, low delay, high throughput

But: Unlike other mechanisms for low latency, L4S involves endpoints *and* middleboxes.

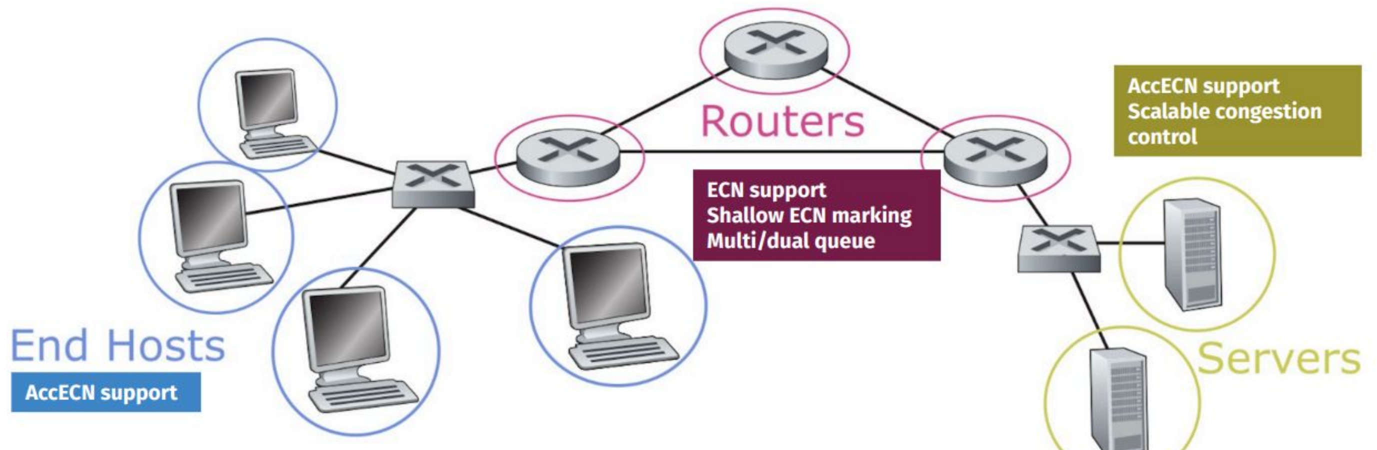

When all of these components are in place, an L4S flow can achieve high throughput with very low latency.

➔ **However**, like any new Internet technology, the deployment of L4S will be **incremental**.

➔ In the initial stages of deployment, L4S flows will coexist with classic flows at L4S or non-L4S bottleneck routers.

L4S Gaining Industry Traction

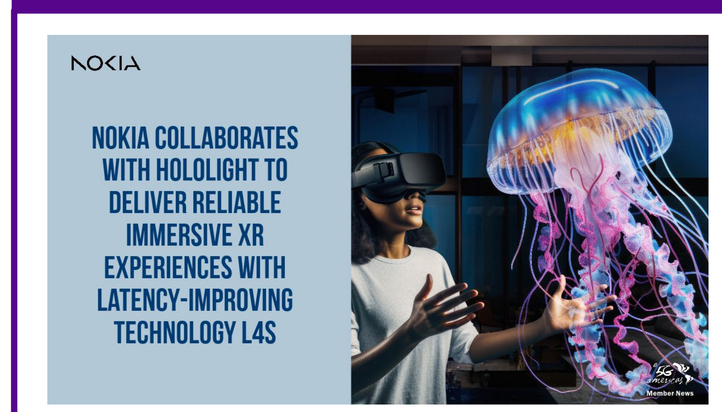

Nokia

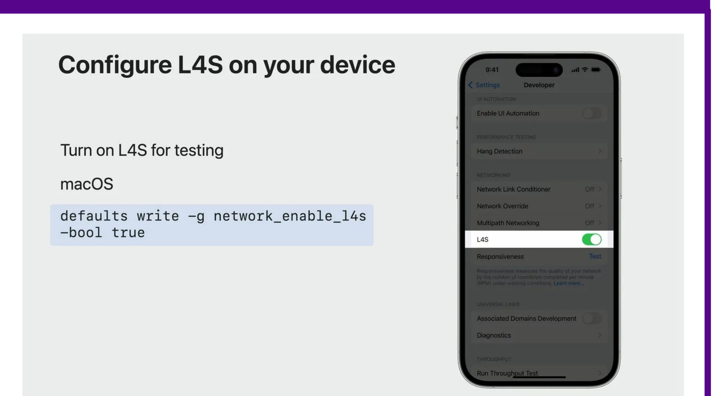

Apple

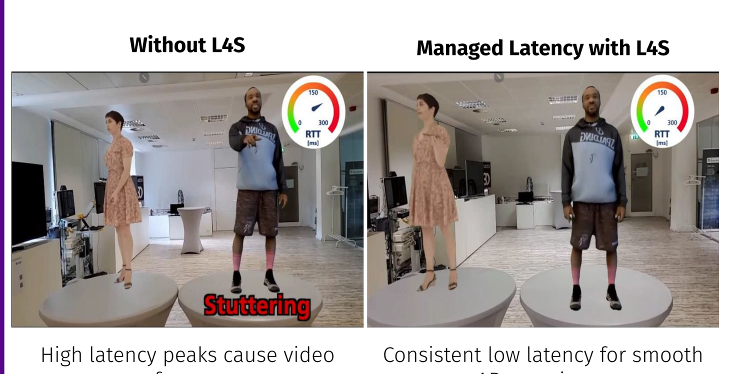

Ericsson

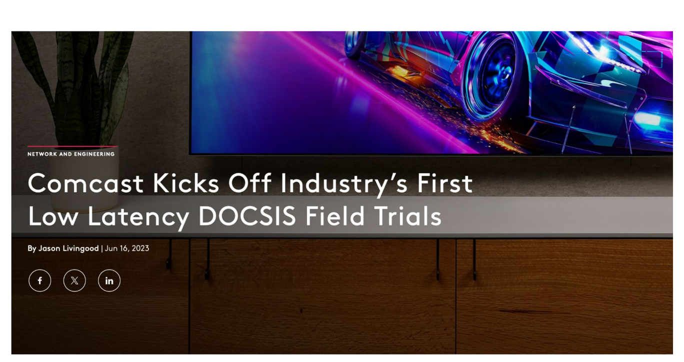

Comcast

Research Questions

1. Can senders be assured that **TCP Prague** (L4S flow) will not cause **harm to, or be harmed by**, another flow at a shared bottleneck link? [1], **we partially analyzed in [2]**
2. Does **L4S-compatible BBRv2** have **more favorable properties** for adoption than TCP Prague? [1]
3. Is the harm caused by or to the L4S-compatible flow mitigated when the bottleneck is shared by a **large number of flows**? [1]
4. **BBRv3** enables ECN based on a path delay threshold. Is this approach better suited for L4S-compatible congestion control deployment? [1]

Experiment Methodology

Line topology on **FABRIC** [4] testbed comprising two sending and two receiving hosts.

Residential home broadband scenario

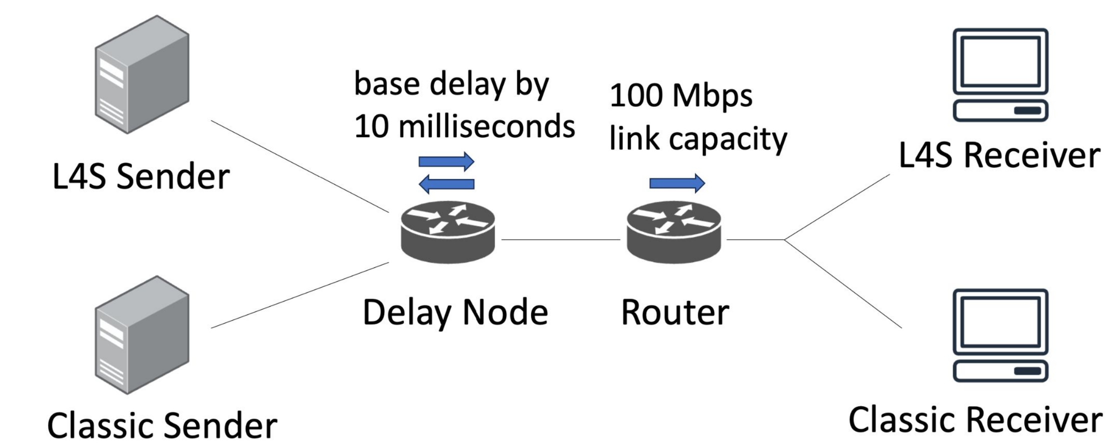

- TCP flow(s) from the L4S sender and from the classic sender, each for a duration of 60 seconds.
- A **wide variety of queue types** that may be encountered at the bottleneck router.
- FIFO, CoDel, PIE, FQ, L4S-aware FQ-CoDel, DualPI2
- **L4S flows**: TCP Prague and L4S-compatible BBRv2
- **non-L4S flows**: CUBIC, BBRv1/v2/v3

Findings - L4S Flows Throughput

*with competing flows in parentheses

RQ-1

➔ **Senders cannot be fully assured**, as TCP Prague may cause harm or be harmed in various bottlenecks, though it coexists well in drop-based and fairness-enforcing queues.

Prague vs CUBIC - ECN (Buffer Size (BDP))

	0.5	1.0	2.0	4.0	8.0
FIFO	40.6 (53.8)	44.6 (50.8)	52.3 (43.2)	51.5 (44.1)	41.9 (53.9)
FIFO (+ECN)	52.1 (41.4)	89.4 (5.3)	90.3 (4.6)	90.5 (4.3)	90.0 (5.0)
CoDel (+ECN+AQM)	42.9 (51.5)	74.9 (19.9)	88.4 (6.5)	88.6 (6.4)	88.6 (6.4)
DualPI2 (+ECN+AQM +DualQueue)	42.9 (52.1)	42.8 (52.2)	42.7 (52.3)	43.1 (52.0)	43.5 (51.6)

Prague vs BBRv2 - ECN (+Rx DCTCP marking) (Buffer Size (BDP))

	0.5	1.0	2.0	4.0	8.0
FIFO	42.6 (50.3)	55.1 (39.4)	77.8 (17.2)	75.2 (20.0)	70.5 (24.7)
FIFO (+ECN)	38.9 (53.3)	62.9 (31.0)	76.3 (18.1)	77.3 (17.0)	77.9 (16.3)
CoDel (+ECN+AQM)	41.4 (51.1)	59.4 (34.8)	64.3 (30.9)	61.7 (33.5)	62.8 (32.2)
DualPI2 (+ECN+AQM +DualQueue)	3.5 (90.4)	4.1 (90.0)	3.8 (90.3)	3.5 (90.4)	3.8 (90.3)

RQ-2

➔ While L4S-compatible BBRv2 has **more favorable properties** for adoption than TCP Prague, **senders still cannot be fully assured** that it will not cause harm to, or be harmed by, other flows at a shared bottleneck link.

L4S-BBRv2 vs Cubic - ECN (Buffer Size (BDP))

	0.5	1.0	2.0	4.0	8.0
FIFO	44.3 (49.9)	40.1 (55.0)	38.0 (57.6)	38.5 (57.0)	32.9 (62.6)
FIFO (+ECN)	47.5 (45.9)	56.1 (37.9)	64.6 (29.4)	62.3 (31.7)	60.8 (33.3)
CoDel (+ECN+AQM)	43.9 (50.3)	57.4 (37.2)	70.5 (23.9)	83.4 (10.1)	83.3 (10.4)
DualPI2 (+ECN+AQM +DualQueue)	56.2 (31.9)	56.2 (31.9)	56.0 (32.2)	56.2 (32.0)	56.2 (31.8)

L4S-BBRv2 vs BBRv2 - ECN (DCTCP Style marking on Receiver) (Buffer Size (BDP))

	0.5	1.0	2.0	4.0	8.0
FIFO	50.6 (43.5)	52.1 (42.1)	60.9 (32.7)	54.5 (39.7)	57.9 (35.9)
FIFO (+ECN)	50.4 (43.6)	56.9 (36.8)	48.3 (45.7)	51.0 (43.1)	49.5 (44.4)
CoDel (+ECN+AQM)	50.6 (43.1)	47.5 (46.4)	40.8 (53.3)	55.6 (38.3)	57.5 (36.4)
DualPI2 (+ECN+AQM +DualQueue)	4.9 (87.9)	4.6 (88.4)	4.8 (88.4)	4.7 (88.2)	4.7 (88.3)

RQ-3

➔ **The harm is not necessarily mitigated** when a scalable flow shares a bottleneck with multiple classic flows.

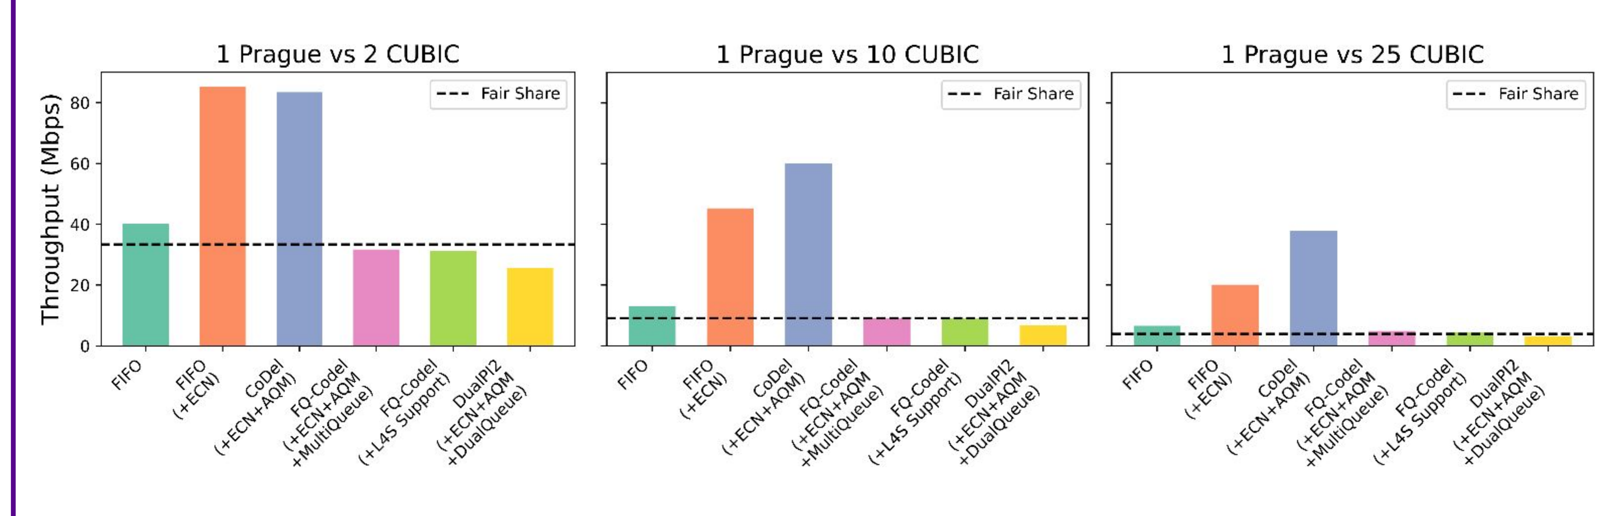

RQ-4

➔ **BBRv3 approach is not more favorable**, as it presents coexistence challenges with both TCP Prague and L4S-compatible BBRv2 across various bottlenecks.

Prague vs BBRv3-Long path delay (Buffer Size (BDP))

	0.5	1.0	2.0	4.0	8.0
FIFO	23.7 (68.9)	25.2 (69.0)	43.1 (52.3)	56.5 (38.7)	55.6 (39.7)
FIFO (+ECN)	22.7 (69.6)	3.6 (89.7)	1.7 (92.7)	1.6 (92.8)	1.5 (92.9)
CoDel (+ECN+AQM)	23.6 (69.3)	24.2 (69.9)	39.0 (56.2)	45.6 (49.9)	41.7 (53.8)
DualPI2 (+ECN+AQM +DualQueue)	3.9 (90.2)	4.1 (90.2)	4.3 (89.9)	4.6 (89.6)	3.6 (90.7)

Prague vs BBRv3-Short path delay (Buffer Size (BDP))

	0.5	1.0	2.0	4.0	8.0
FIFO	12.4 (75.8)	14.2 (78.7)	22.4 (71.9)	61.0 (33.8)	79.9 (14.8)
FIFO (+ECN)	13.3 (74.2)	10.1 (80.3)	11.4 (80.5)	11.2 (80.7)	11.6 (80.4)
CoDel (+ECN+AQM)	13.3 (77.5)	14.5 (78.8)	25.8 (68.8)	56.0 (39.1)	64.0 (30.6)
DualPI2 (+ECN+AQM +DualQueue)	5.0 (87.3)	5.2 (87.1)	5.0 (87.3)	5.2 (87.2)	7.0 (87.3)

Takeaway

➔ Given that the sender of an L4S flow cannot be sure what type of queue is at the bottleneck router or what other flows will share it, **safe coexistence cannot be guaranteed right now**.

Future Work

➔ Investigating the impact on L4S adoption and strategies for mitigation more directly.

➔ Extending the evaluation to more realistic traffic patterns and other network settings.

References

[1] **F. B. Sarpkaya**, F. Fund, and S. Panwar, "To adopt or not to adopt L4S-compatible congestion control? Understanding performance in a partial L4S deployment," *Under review*, Oct. 2024.
[2] **F. B. Sarpkaya**, A. Srivastava, F. Fund, and S. Panwar, "To switch or not to switch to TCP Prague? Incentives for adoption in a partial L4S deployment," ANRW '24, Vancouver, Canada, Jul. 2024. [**QR**]
[3] B. Briscoe, K. D. Schepper, M. Bagnulo, and G. White, "Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture," RFC 9330, Jan. 2023. [Online].
[4] Baldin, Ilya, et al. "*FABRIC: A national-scale programmable experimental network infrastructure.*" IEEE Internet Computing 23.6 (2019): 38-47.