

# (L4S XOR RFC3168) ECN Marking for improved detection of Classic ECN AQMs?

Bob Briscoe, Independent

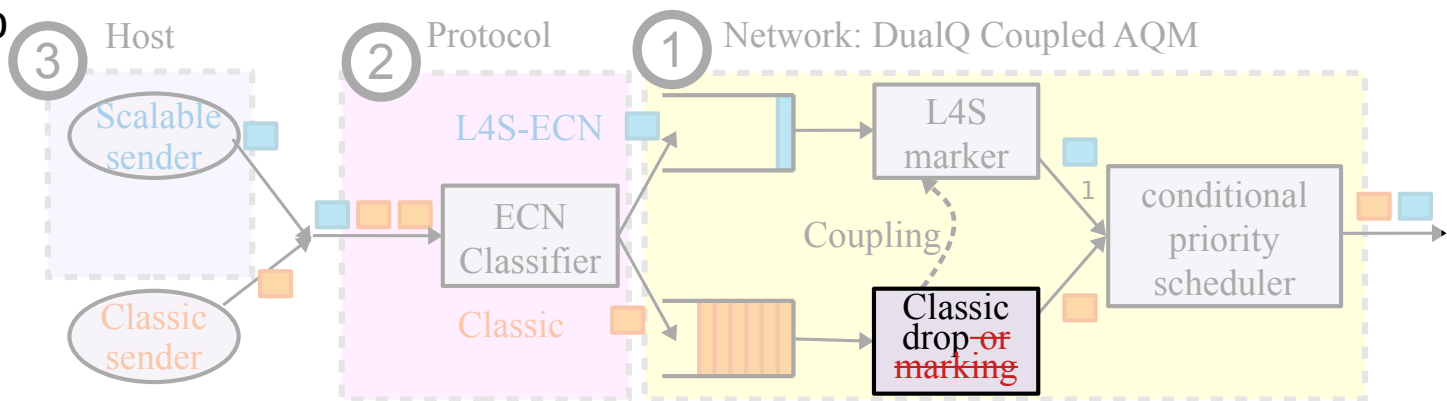
<ietf@bobbriscoe.net>



Mar 2021

# Exclusive ECN Marking – the Base Proposal

- An L4S AQM node that marks ECT1 packets **MUST NOT** also mark ECT0 packets
- Rationale: Would make the presence of an RFC3168 AQM more clear-cut
- Recap of Problem
  - L4S would outcompete Classic in an RFC3168 AQM
  - L4S sources are meant to detect an RFC3168 AQM
  - Certainty that it's not L4S has proved challenging




# In-Band Active Detection Ex.#1

## ECT0 probes

- L4S source

- minimise extra load
- 8 ECT1 data packets : 1 ECT0 probe ( $P = 8$ )
- data 1500B, probes 75B ( $r = 1500/75 = 20$ )




ECT0 drop mark		ECT1 drop mark		Inferred bottleneck type
>0	0	>0	0	
	>0		>0	Classic ECN AQM (FQ or FIFO)
>0	0		>0	L4S AQM

- How long to decide no ECT0 have been marked?

- after  $N$  ECT1 packets marked, where  $N = r * P * 5$  (say)
- and at least one ECT0 drop

# FQ-Exclusive ECN Marking

- An L4S <sup>per-flow-queue</sup> AQM ~~node~~ that marks ECT1 packets MUST NOT also mark ECT0 packets
- if ECT1 seen 'recently', disable marking ECT0
  - just in that flow-queue
  - 'recently' either requires a timer, or for the life of the queue
- Not essential to disable ECT0 marking
  - An L4S source ought to keep queue below ECT0 target anyway




# In-Band Active Detection Ex.#1

## ECT0 probes – problems

- L4S source
  - minimise extra load
  - 8 ECT1 data packets : 1 ECT0 probe ( $P = 8$ )
  - data 1500B, probes 75B ( $r = 1500/75 = 20$ )

ECT0 drop mark		ECT1 drop mark		Inferred bottleneck type
>0	0	>0	0	
	>0		>0	Classic ECN AQM (FQ or FIFO)
>0	0		>0	L4S AQM



- If ECT0 marked, proves RFC3168
  - but no ECT0 marked, doesn't disprove
- Other reasons for no ECT0 markings
  - variable congestion didn't coincide with probes
  - size-based marking, e.g. DOCSIS PIE
  - didn't test for long enough

- How long to decide no ECT0 have been marked?
  - after  $N$  ECT1 packets marked, where  $N = r * P * 5$  (say)
  - and at least one ECT0 drop

- Challenges
  - delayed ACKs – which packet was marked?
    - if TCP,  $seqno = snd\_next - 1$
  - 800 marks is too long for in-band detection

# Exclusive ECN marking


## What does it give us?

- Ideal would be an in-band passive test
  - But exclusive ECN marking is inherently for active testing
- In-band active test (ECT0 probes) have to minimize extra load
  - Then, too slow to catch unfairness in time
- Can use exclusive ECN marking for a fast out-of-band test (→ spare slide)
  - but once we've resorted to out-of-band, no longer constrained to minimize extra load
  - then, we already have good tests **without** exclusive ECN marking...

Out-of-Band test without exclusive marking  
parallel L4S (L) & Classic (C) test flows  
can distinguish everything

Rate	RTT	Inferred AQM
$L > C$	$L = C$	Classic ECN AQM (FIFO)
$L = C$	$L = C$	Classic ECN AQM (FQ)
$L = C$	$L < C$	FQ Classic+L4S AQM
$L \approx C$	$L < C$	DualQ Coupled AQM

# Exclusive ECN marking Summary



## Cons

- Rapid in-band detection strategy not possible (yet?)
- Useful out-of-band, but we already have good out-of-band tests without it
- Only useful if near-universal compliance

## Pros

- Seems promising, but not a silver bullet
- Can be withdrawn later (but can't be introduced later)

Tech report: <https://arxiv.org/pdf/1911.00710.pdf#subsection.5.3>

(L4S XOR RFC3168) ECN Marking  
for improved detection of Classic ECN AQMs?

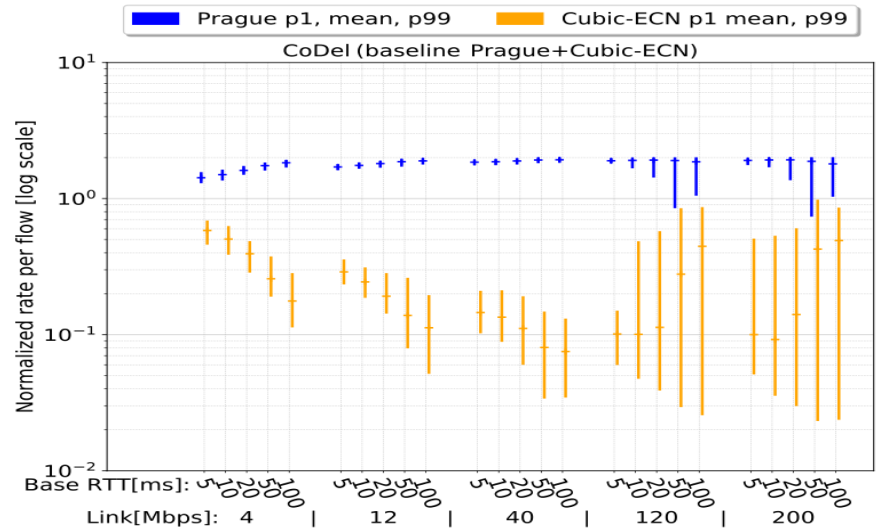
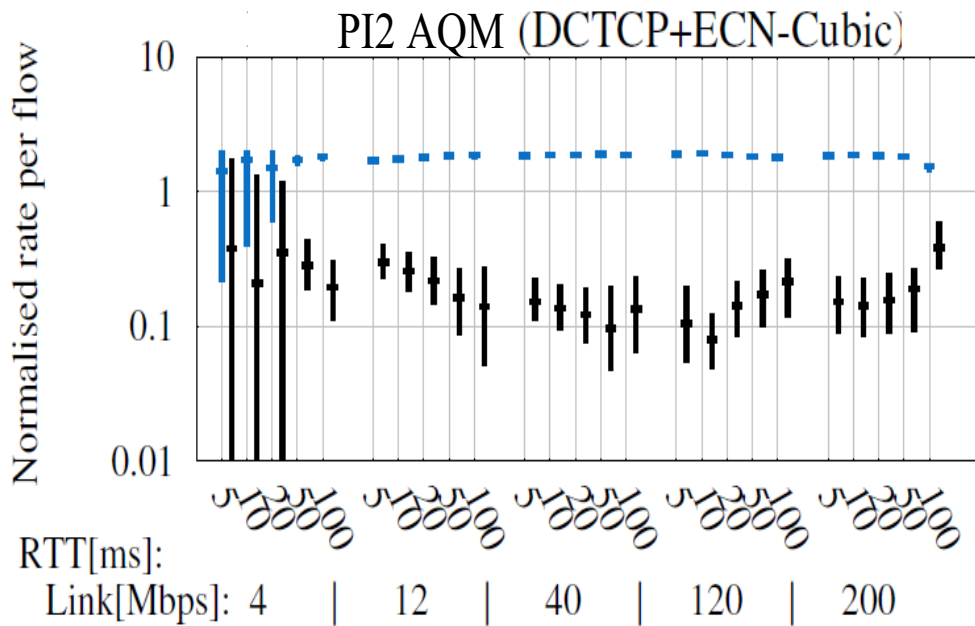
Q&A

spare slides



# Problem: Coexistence between L4S and Classic flows in a FIFO RFC3168 ECN AQM

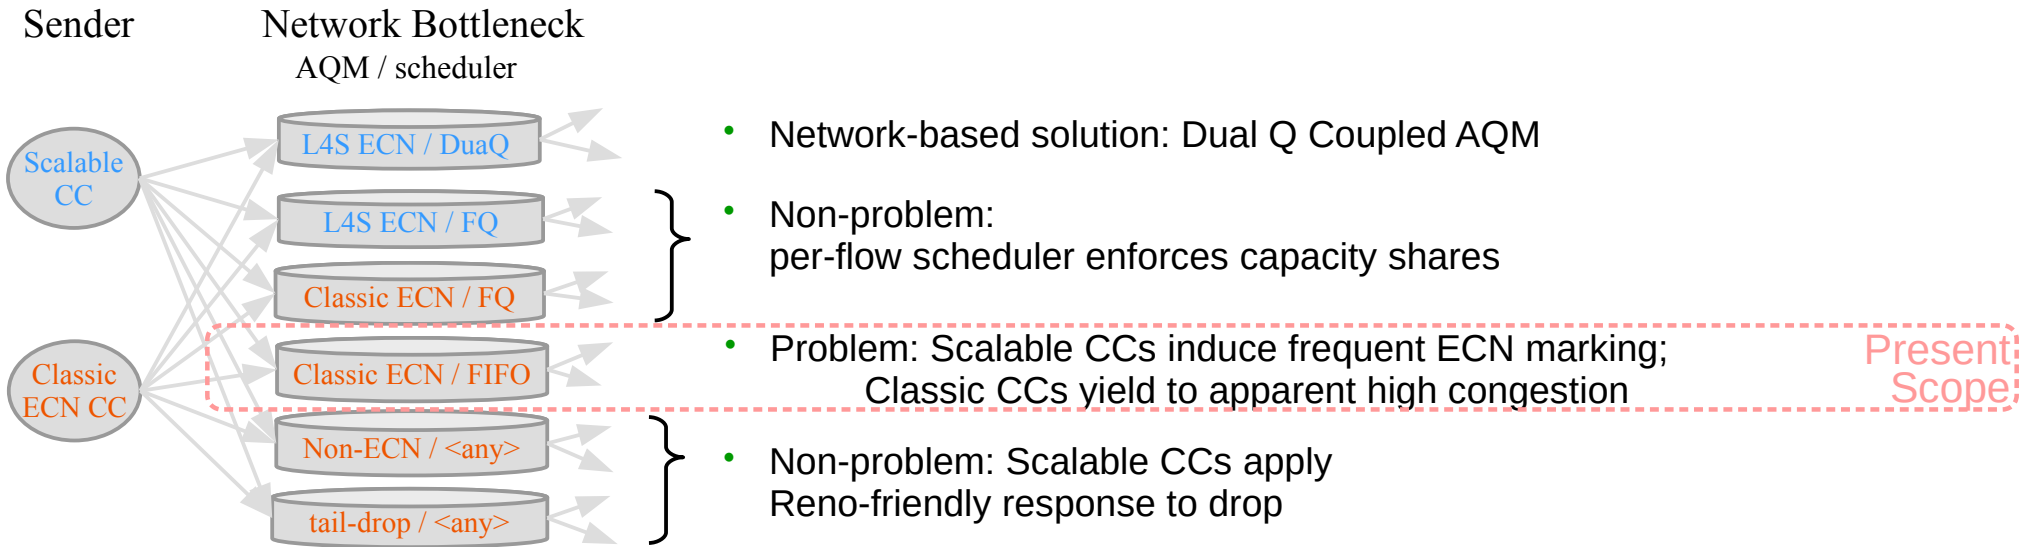
DCTCP P1, mean, P99 ■ ECN-Cubic P1, mean, P99 ■



- Normalized rate per flow = flow rate after convergence / (capacity / no. of flows)
- 1 v 1 long-running flows
- Default config. for all CCs and AQMs

# The Full Coexistence Scope

- Across all combinations of congestion control, AQM & scheduler




- Classic ECN: RFC3168 Explicit Congestion Notification
- CC: Congestion Control
- Scalable CC:  $1/p$  response to congestion ( $p$ )
- Classic CC: Reno-Friendly CC
- AQM: Active Queue Management
- FIFO: First-In First-Out
- FQ: Per-Flow Queuing
- L4S: Low Latency Low Loss Scalable throughput

# Exclusive ECN marking Implications if adopted by WG

- Would need to update:
  - 3 main L4S drafts
  - Linux DualPI2 code & other implementations
  - (Low Latency DOCSIS already doesn't support ECT0 marking, for hardware backward compatibility)
- Not onerous

# Out-of-Band Active Detection Ex.#2

## Late onset ECT1 samples

- L4S source 
  - ECT0 until CE mark
  - then 1 ECT1 : 20 ECT0 (all full-sized data packets)

 ECT(1)  
 ECT(0)

All ECT0		95% ECT0		5% ECT1		Implied bottleneck type
drop	mark	drop	mark	drop	mark	
>0	0	>0	0	>0	0	Tail drop or non-ECN AQM
	>0		>0		>0	RFC3168 ECN AQM (FQ or FIFO)
	>0		>0		most	FQ Classic+L4S AQM (non-exclusive)
	>0	>0	0		>0	FQ Classic+L4S AQM (exclusive)
>0	0	>0	0		>0	DualQ L4S AQM (exclusive)