



Fighting the bufferbloat: on the coexistence of **AQM** and **LPCC**

Y. Gong, D. Rossi, C. Testa, S. Valenti
Télécom ParisTech, France
first.last@telecom-paristech.fr

M. D. Täht
bufferbloat.net
dave.taht@gmail.com

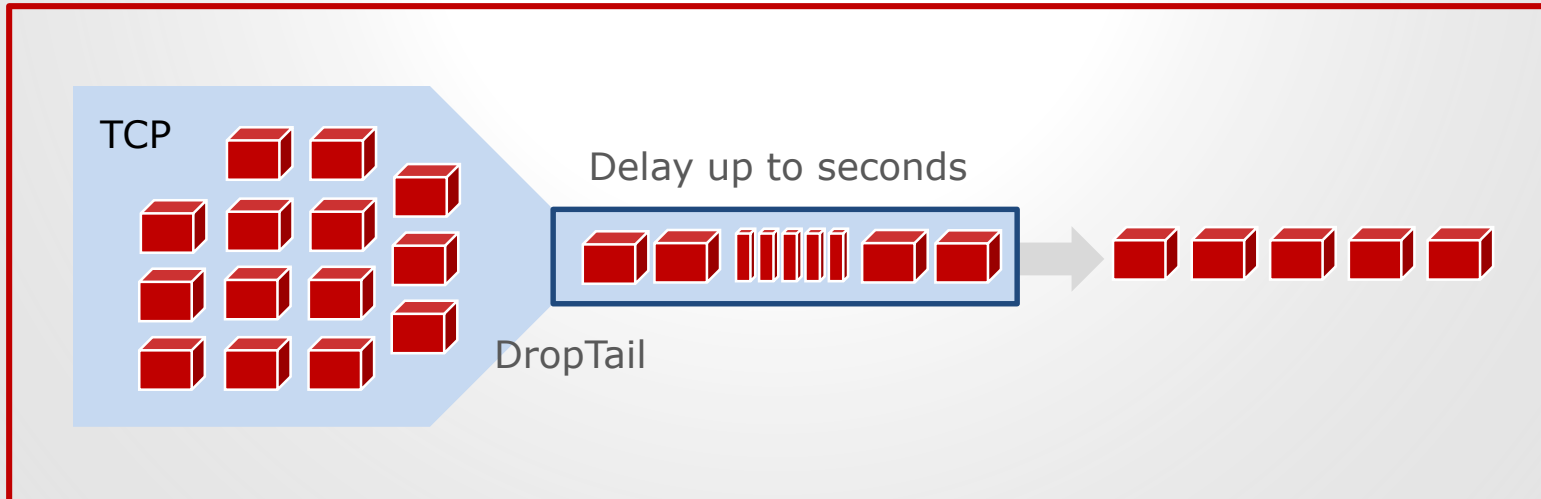
AQM: Active Queue Management
LPCC: Low Priority Congestion Control

IEEE INFOCOM TMA 2013
April 19, 2013
Turin, Italy

Motivation

Bufferbloat

excessively large buffers in network communication systems.




Motivation

Active Queue Management

SFQ, RED, DRR, Choke, CoDel ...

Bufferbloat

excessively large buffers in network communication systems.



Motivation

Active Queue Management

SFQ, RED, DRR, Choke, CoDel ...

RED: Random Early Drop

Bufferbloat

excessive buffers in network communication systems.

TCP



Motivation

Active Queue Management

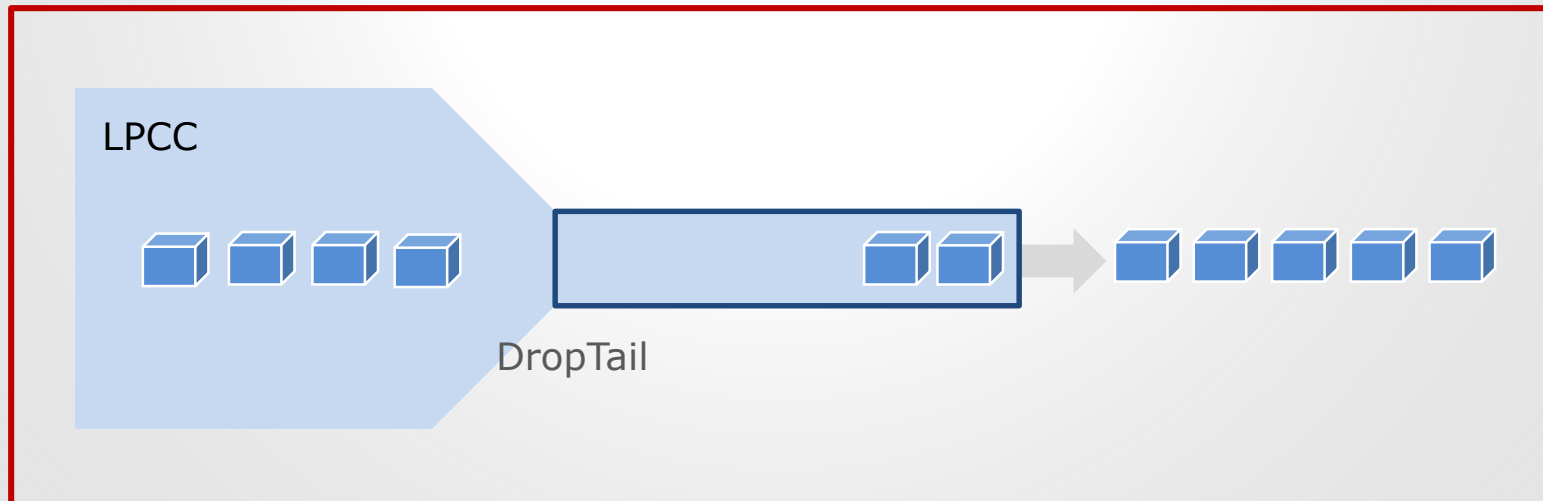
SFQ, RED, DRR, Choke, CoDel ...

Bufferbloat

excessively large buffers in network communication systems.

Low Priority Congestion Control

NICE, TCP-LP, LEDBAT ...



Motivation

Active Queue Management

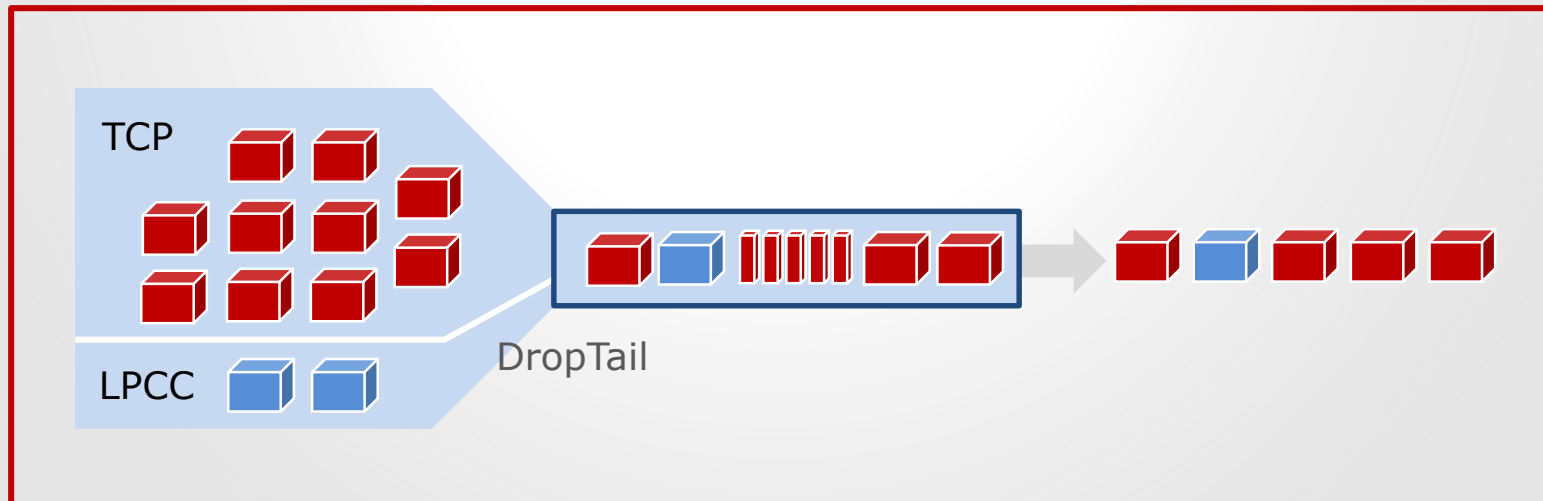
SFQ, RED, DRR, Choke, CoDel ...

Bufferbloat

excessively large buffers in network communication systems.

Low Priority Congestion Control

NICE, TCP-LP, LEDBAT ...

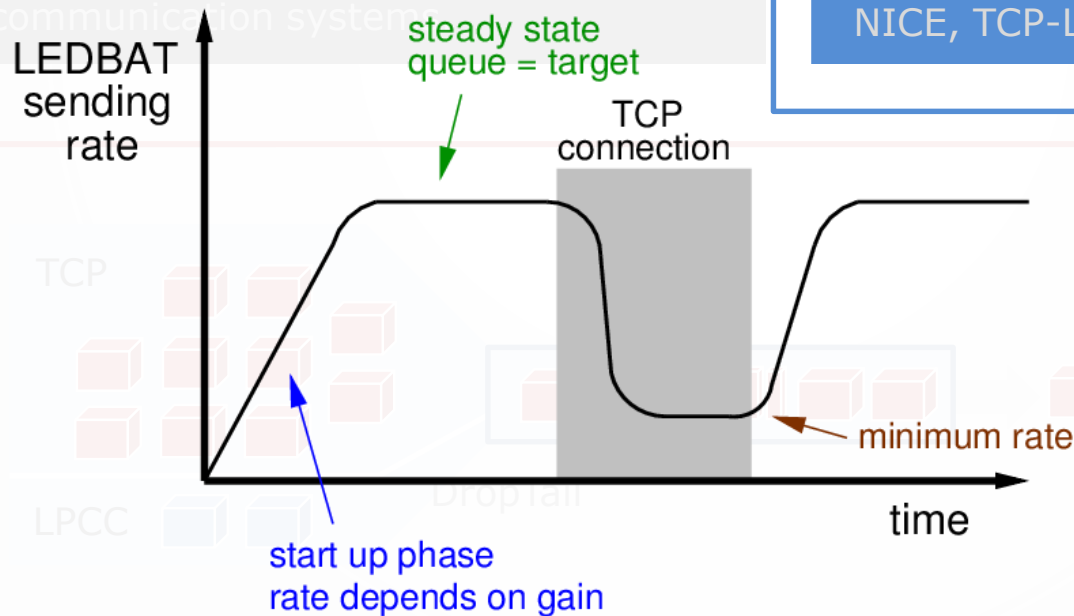


Motivation

LEDBAT: Low Extra Delay Background Transport

... calculate on-way-delay ...
 $\text{off_target} = \text{TARGET} - \text{queuing_delay}$
 $\text{cwnd} += \text{GAIN} * \text{off_target} / \text{cwnd}$

excessively large buffers in network
communication systems



Low Priority Congestion Control

NICE, TCP-LP, LEDBAT ...

Motivation

Active Queue Management

SFQ, RED, DRR, Choke, CoDel ...



Reprioritization



Bufferbloat

excessively large buffers in network communication systems.

Low Priority Congestion Control

NICE, TCP-LP, LEDBAT ...

Why

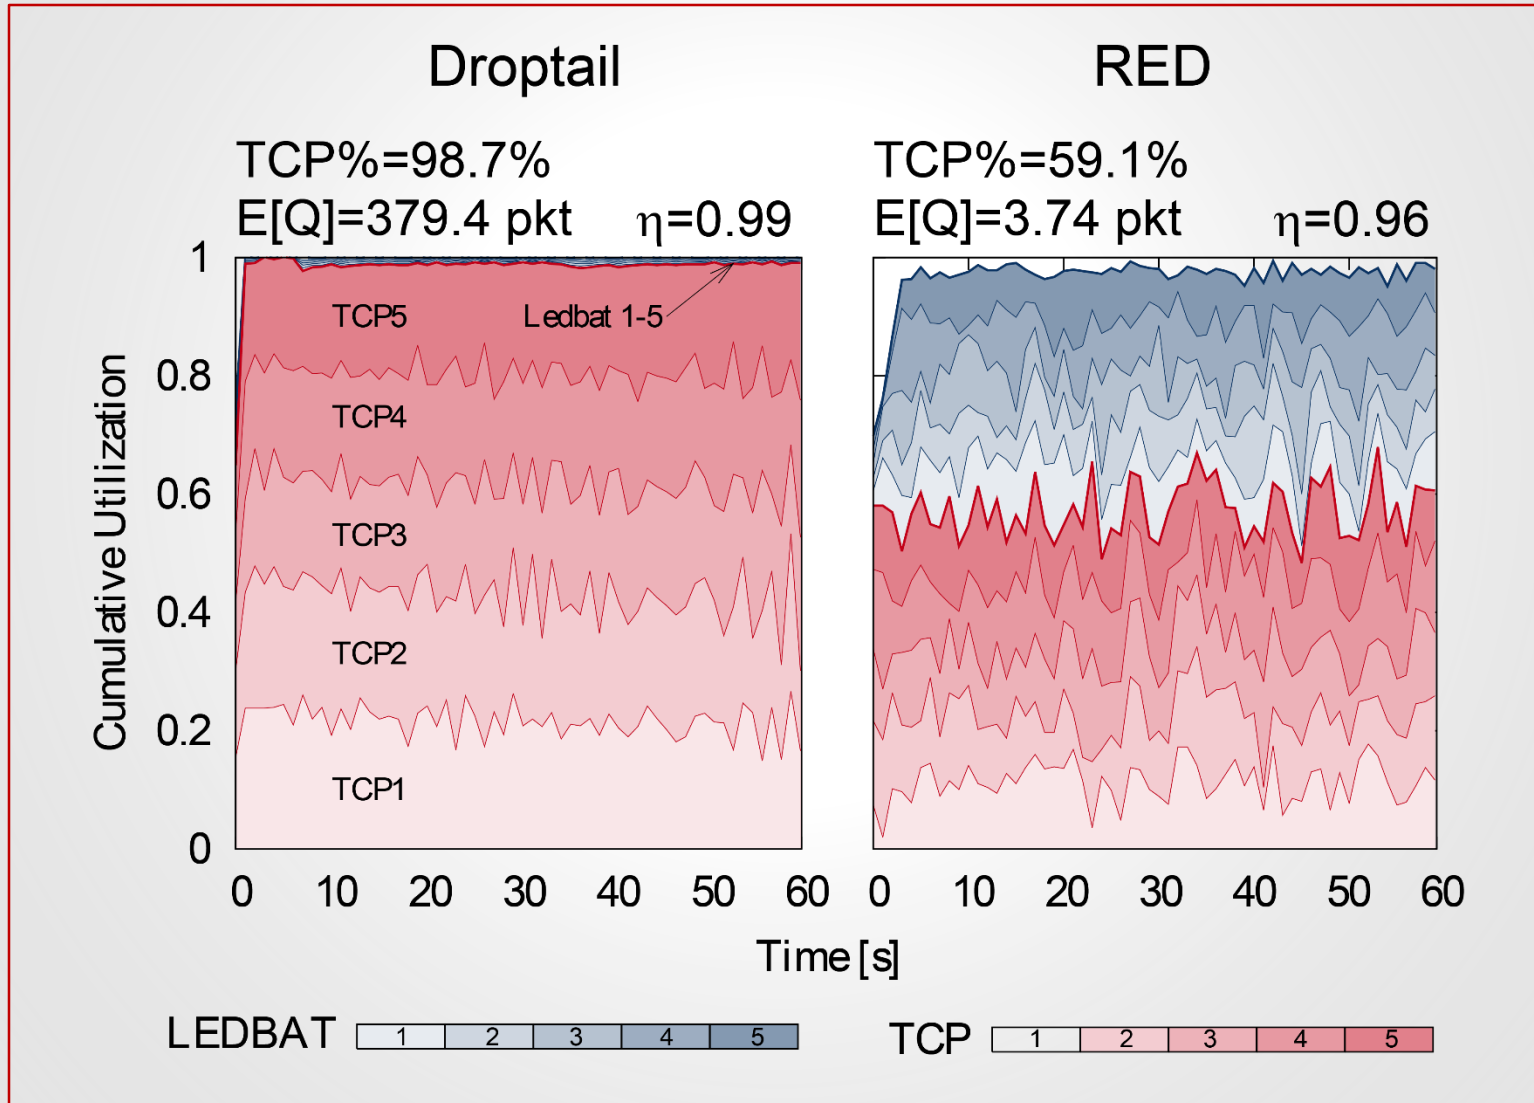
- Long history and active development
- Deployment in real world
- Independent progress

TCP

PCC

AQM

Reprioritization



Methodology

Methodology

- Experiment



- Simulation



Result



Results

- Validity (experiment)
 - LPCC \in [LEDBAT, TCP-LP]
 - AQM/Schedule \in [RED, SFQ]
 - $N \in [1, 5]$
 - 24 scenarios
- Generality (simulation)
 - LPCC \in [LEDBAT, NICE, TCP-LP]
 - AQM/Schedule \in [RED, CHOKe, DRR, SFQ, CoDel]
 - 5000+ simulations

- Scripts

<http://goo.gl/wvG0D>

Testbed



LPCCC	AQM	5+5 flows		1+1 flows	
		TCP%	E[Q] ms	TCP%	E[Q] ms
 LEDBAT	FIFO	94.9	6437.3	99.5	5304.3
	RED	49.3	11.3	1.7	9.5
	SFQ	76.1	106.4	57.7	15.1
 TCP-LP	FIFO	50.8	7879.6	65.8	7471.9
	RED	57.0	21.0	97.5	2.7
	SFQ	49.8	144.2	50.0	25.2

LPCCC dynamics may introduce unexpected behavior

- TCP starvation under LEDBAT + RED (1+1 flows)
- Queuing delay under SFQ approaches limit* (100-150ms)

* Considered to be harmful for interactive communication by ITU G.114

Internet

		5+5 flows		1+1 flows	
AQM		TCP%	E[Q] ms	TCP%	E[Q] ms
 Testbed	FIFO	94.9	6437.3	99.5	5304.3
	RED	49.3	11.3	1.7	9.5
	RED#	71.9	763.9	92.2	333.8
 Internet	FIFO	97.0	6551.7	98.9	916.0
	RED	2.6	45.0	5.1	29.1
	RED#	63.8	745.2	86.7	305.4

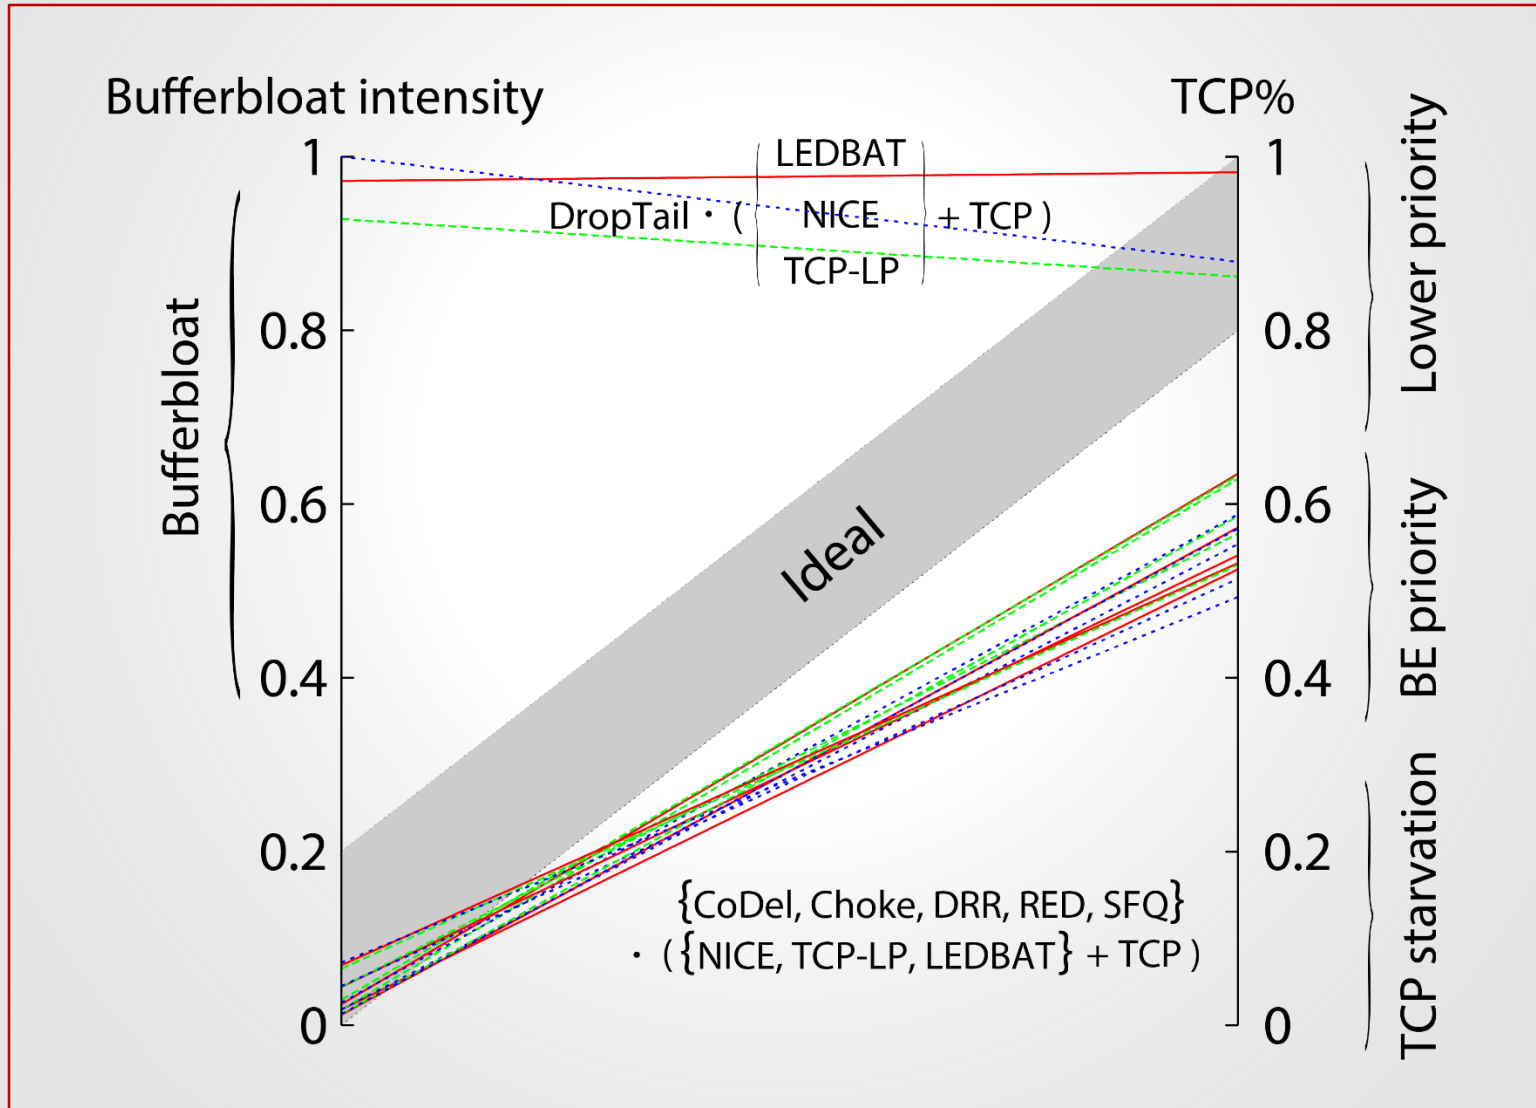
* LPCC = LEDBAT

Internet experiment confirms previous observation

- Confirm testbed result (reprioritization, tcp starvation)
- RED# reinstates LPCC priority but still has bufferbloat

* RED# ad hoc setting with `min_threshold > LEDBAT target (100ma)`, not practical

Simulation



Conclusion and Future Work

Conclusion

- Reprioritization is general and unavoidable under present AQM&LPCC design

Future work

- Refining results w/ recent updates on kernel
 - TCP small queue (3.6)
 - CoDel (3.5)
 - Byte Queue Limit (3.3)
- Analytical modeling
 - Fluid model of TCP + LEDBAT under AQM
 - Equilibrium TCP% is related to LEDBAT delay target and equilibrium queue length
 - <http://arxiv.org/abs/1303.6817>

