THE GAME THEORY BEHIND RUNNING CUBIC AND BBR ON THE INTERNET

Ayush Mishra[1], Jingzhi Zhang[2], Melodies Sim[1], Sean Ng[1], Raj Joshi[1], Ben Leong[1]

[1] National University of Singapore
[2] Wuhan University

IETF 112

A lot of websites have made the performance driven decision to switch to BBR.

It has been reported that switching to BBR has improved throughput and reduced delay

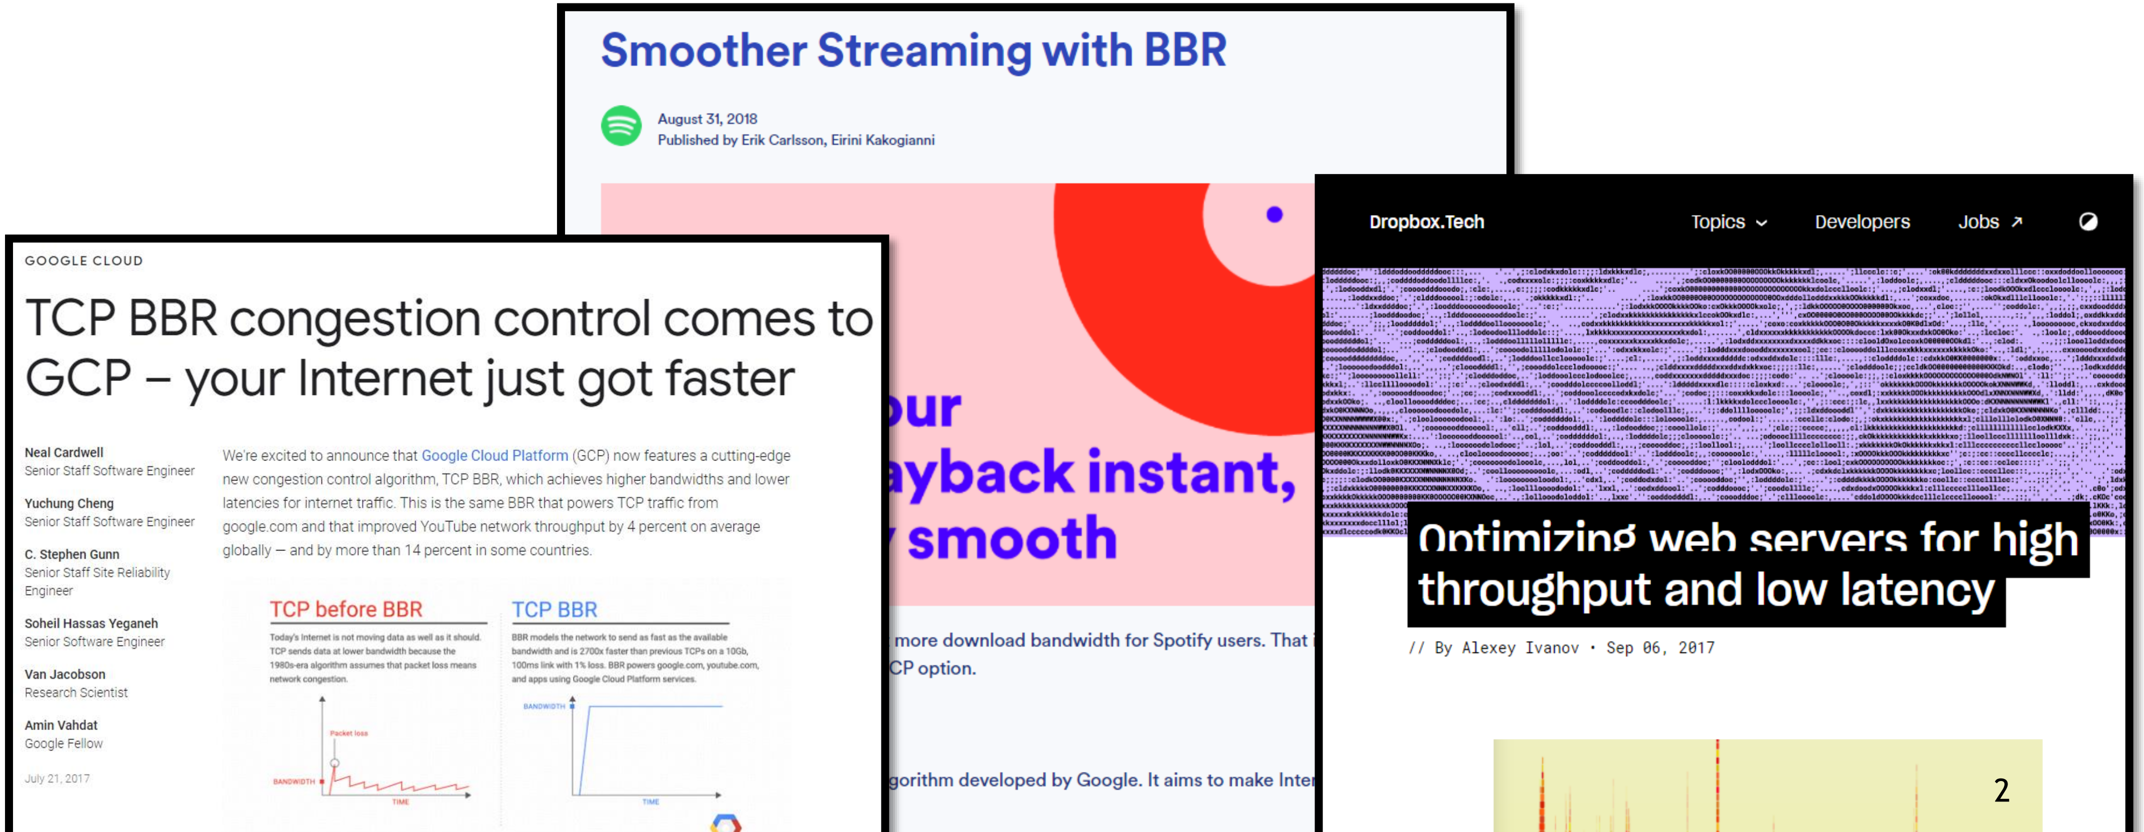

Close to **18%** of the Alexa Top 20,000 websites run BBR

This share is even larger among websites that are more popular or have a higher share of the downstream traffic

Variant	Websites	Proportion
CUBIC [15]	6,139	30.70%
BBR [4]	3,550	17.75%
BBR G1.1	167	0.84%
YeAH [2]	1,162	5.81%
CTCP [34]/Illinois[22]	1,148	5.74%
Vegas [3]/Veno [13]	564	2.82%
HTCP [21]	560	2.80%
BIC [37]	181	0.90%
New Reno [28]/HSTCP [12]	160	0.80%
Scalable [20]	39	0.20%
Westwood [7]	0	0.00%
Unknown	3,535	17.67%
Short flows	1,493	7.46%
Unresponsive websites	1,302	6.51%
Total	20,000	100%

The Great Internet TCP Congestion Control Census

AYUSH MISHRA, National University of Singapore, Singapore
XIANGPENG SUN, National University of Singapore, Singapore
ATISHYA JAIN, Indian Institute of Technology, Delhi, India
SAMEER PANDE, Indian Institute of Technology, Delhi, India
RAJ JOSHI, National University of Singapore, Singapore
BEN LEONG, National University of Singapore, Singapore

In 2016, Google proposed and deployed a new TCP variant called BBR. BBR represents a major departure from traditional congestion-window-based congestion control. Instead of using loss as a congestion signal, BBR

Where is this evolution headed?

What is the next paradigm shift in the Internet's congestion control landscape going to look like?

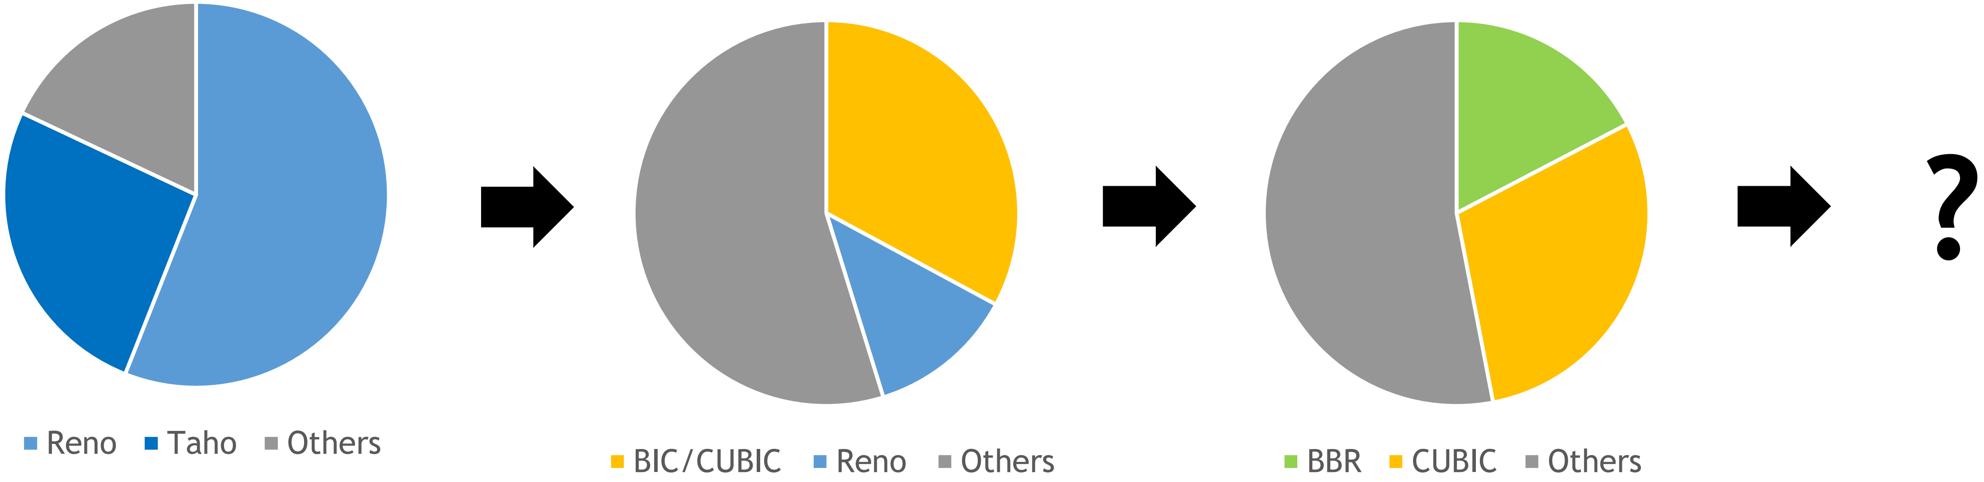

[2001] Padhya et al.

[2011] Yang et al.

[2019] Mishra et al.

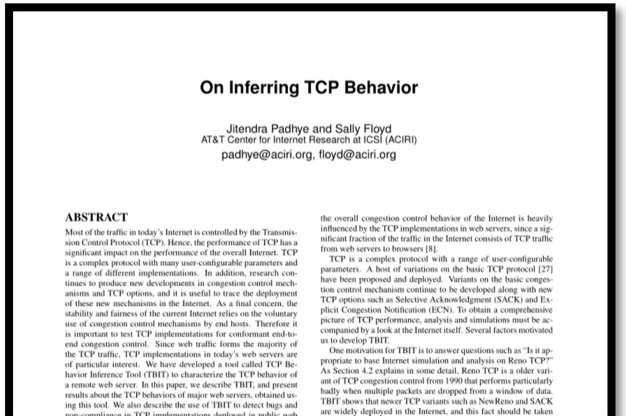

On Inferring TCP Behavior

Jitendra Padhye and Sally Floyd
AT&T Center for Internet Research at ICSI (ACIRI)
padhye@aciri.org, floyd@aciri.org

ABSTRACT

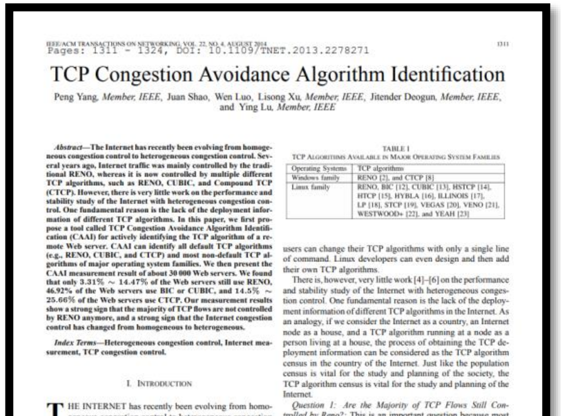

TCP Congestion Avoidance Algorithm Identification

Peng Yang, Member, IEEE, Juan Shao, Wen Luo, Lisong Xu, Member, IEEE, Jitender Deogun, Member, IEEE, and Ying Lu, Member, IEEE

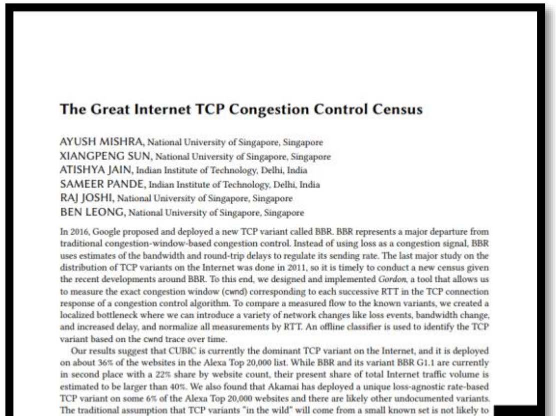

The Great Internet TCP Congestion Control Census

AYUSH MISHRA, National University of Singapore, Singapore
XIANGPENG SUN, National University of Singapore, Singapore
ATISHYA JAIN, Indian Institute of Technology, Delhi, India
SAMEER PANDE, Indian Institute of Technology, Delhi, India
RAJ JOSHI, National University of Singapore, Singapore
BEN LEONG, National University of Singapore, Singapore

Given this performance improvement, how do we expect the Internet to evolve?

Is it reasonable to expect everyone to switch from CUBIC to BBR?

APNET 2021

Conjecture: Existence of Nash Equilibria in Modern Internet Congestion Control

AYUSH MISHRA, National University of Singapore, Singapore
JINGZHI ZHANG, Wuhan University, China
MELODIES SIM, National University of Singapore, Singapore
SEAN NG, National University of Singapore, Singapore
RAJ JOSHI, National University of Singapore, Singapore
BEN LEONG, National University of Singapore, Singapore

The Internet's congestion control landscape is currently in the midst of an unprecedented paradigm shift. A recent measurement study found that BBR, a congestion control algorithm introduced by Google in 2016, has seen rapid adoption and is deployed at more

China, China. ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3469393.3469397

Given this performance improvement, how do we expect the Internet to evolve?

Is it reasonable to expect everyone to switch from CUBIC to BBR?

Key Insight:

We have some players that can maximize some **utility**
All the players have **strategies** (CUBIC/BBR) available to them to maximize their utility.

It's a Normal Form Game!

Given some network, we can calculate the Nash Equilibrium CC distribution.

CC Distribution:

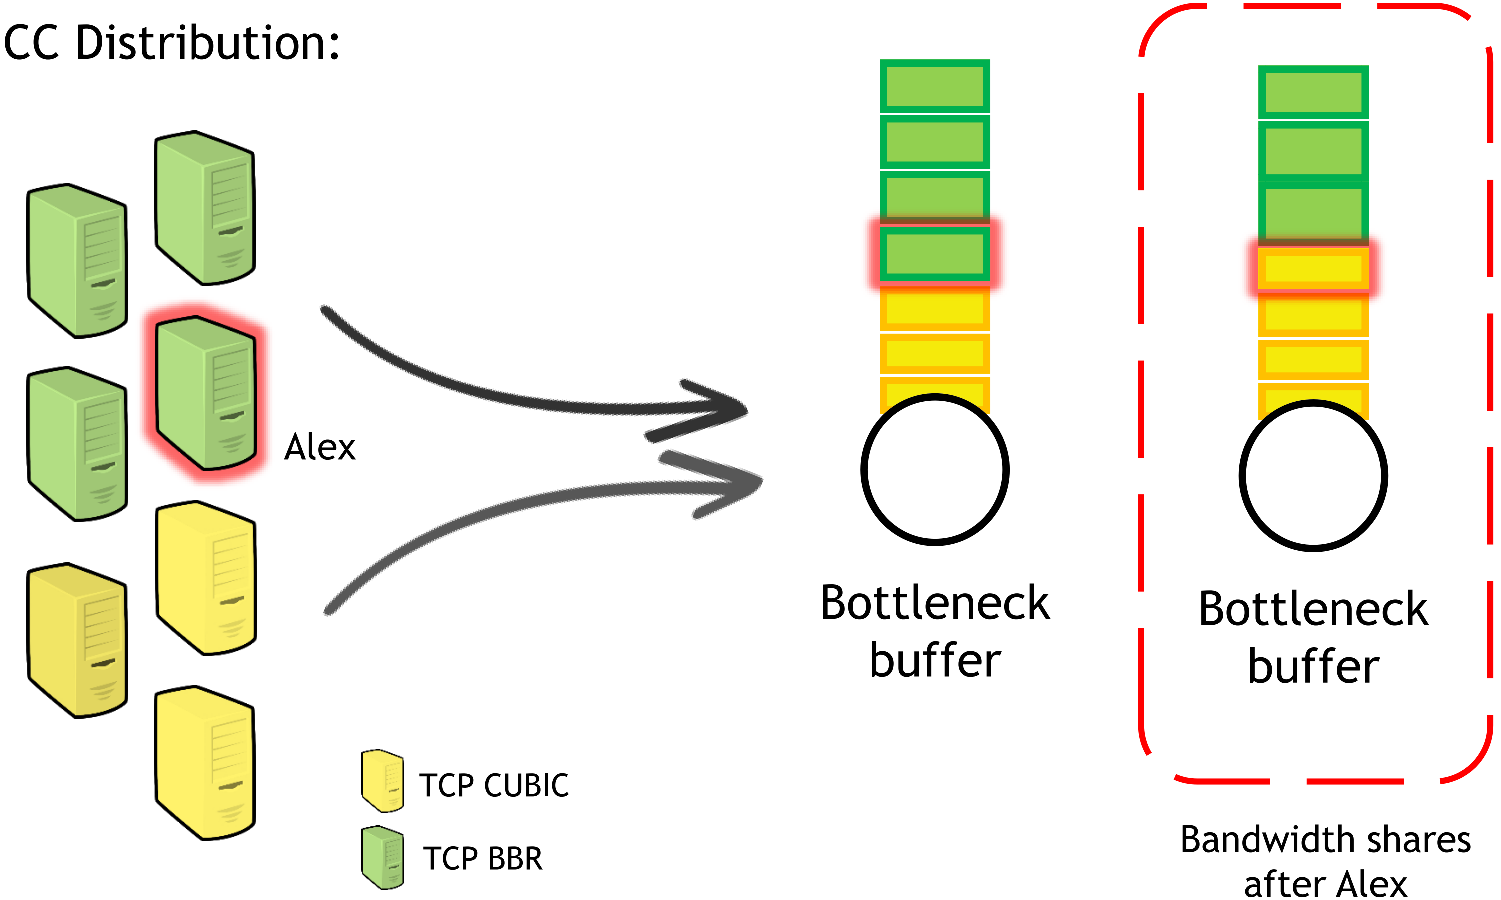

Bandwidth shares after Alex switches CCA

If everyone does worse after making a switch, then the given CC Distribution is the **Nash Equilibrium**.

If everyone on the Internet chooses between CUBIC and BBR based on throughput, this Nash Equilibrium distribution is our best *estimate* of the future of the congestion control landscape.

We think that there will always be a Nash Equilibrium in a network with CUBIC and BBR flows.

Combined throughput of all the BBR flows

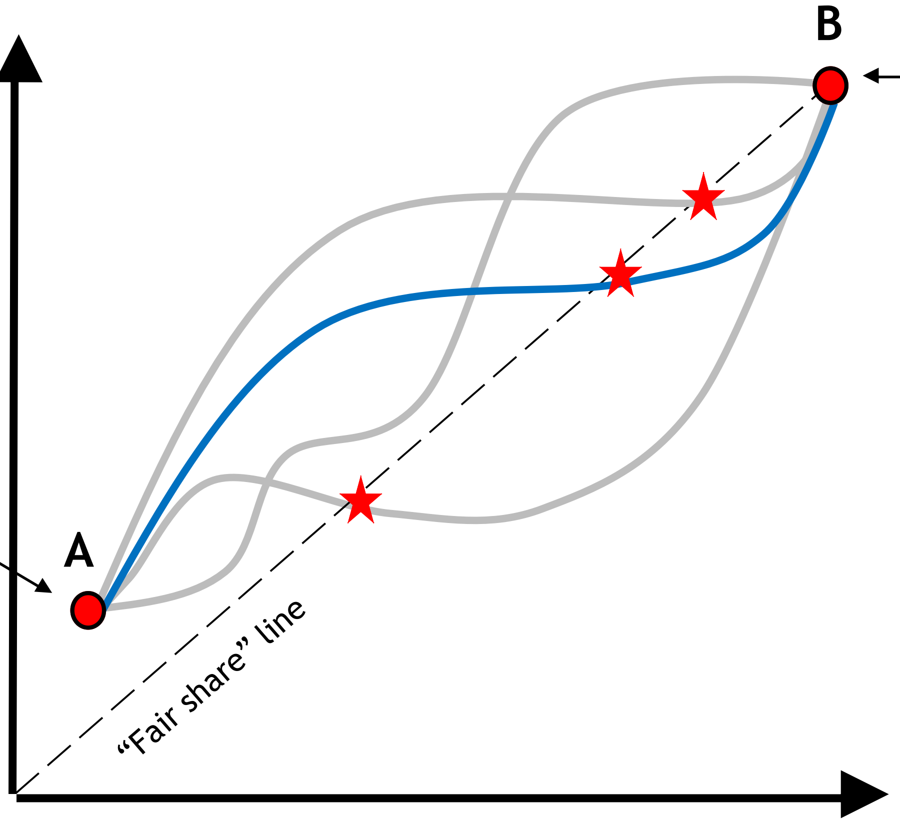

Percentage of flows running BBR at the bottleneck

When all the flows at the bottleneck are BBR flows, they will utilize the bottleneck link

A small number of BBR flows can get a disproportionately high share of the bottleneck bandwidth

Each point of intersection is a Nash Equilibrium distribution!

[2019] Ware et al.

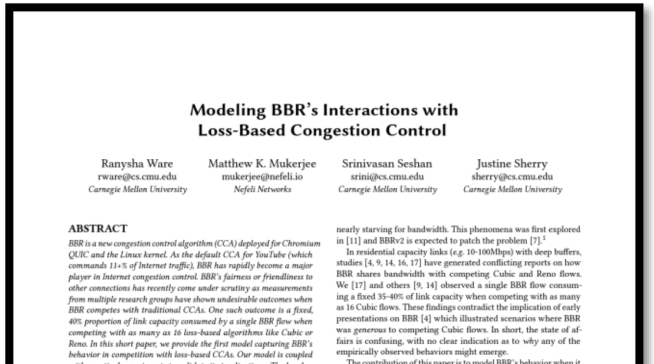

Modeling BBR's Interactions with Loss-Based Congestion Control

Ranysha Ware, rware@cs.cmu.edu, Carnegie Mellon University
Matthew K. Mukerjee, mukerjee@nefeli.io, Nefeli Networks
Srinivasan Seshan, srini@cs.cmu.edu, Carnegie Mellon University
Justine Sherry, sherry@cs.cmu.edu, Carnegie Mellon University

ABSTRACT

BBR is a new congestion control algorithm (CCA) deployed for Chromium QUIC and the Linux kernel. As the default CCA for YouTube (which commands 11+% of Internet traffic), BBR has rapidly become a major player in Internet congestion control. BBR's fairness or friendliness to other connections has recently come under scrutiny as measurements from multiple research groups have shown undesirable outcomes when BBR competes with traditional CCAs. One such outcome is a fixed, 40% proportion of link capacity consumed by a single BBR flow when competing with as many as 16 loss-based algorithms like Cubic or Reno. In this short paper, we provide the first model capturing BBR's behavior in competition with loss-based CCAs. Our model is coupled

nearly starving for bandwidth. This phenomena was first explored in [11] and BBRv2 is expected to patch the problem [7].[1]

In residential capacity links (*e.g.* 10-100Mbps) with deep buffers, studies [4, 9, 14, 16, 17] have generated conflicting reports on how BBR shares bandwidth with competing Cubic and Reno flows. We [17] and others [9, 14] observed a single BBR flow consuming a fixed 35-40% of link capacity when competing with as many as 16 Cubic flows. These findings contradict the implication of early presentations on BBR [4] which illustrated scenarios where BBR was *generous* to competing Cubic flows. In short, the state of affairs is confusing, with no clear indication as to *why* any of the empirically observed behaviors might emerge.

We think that there will always be a Nash Equilibrium in a network with CUBIC and BBR flows.

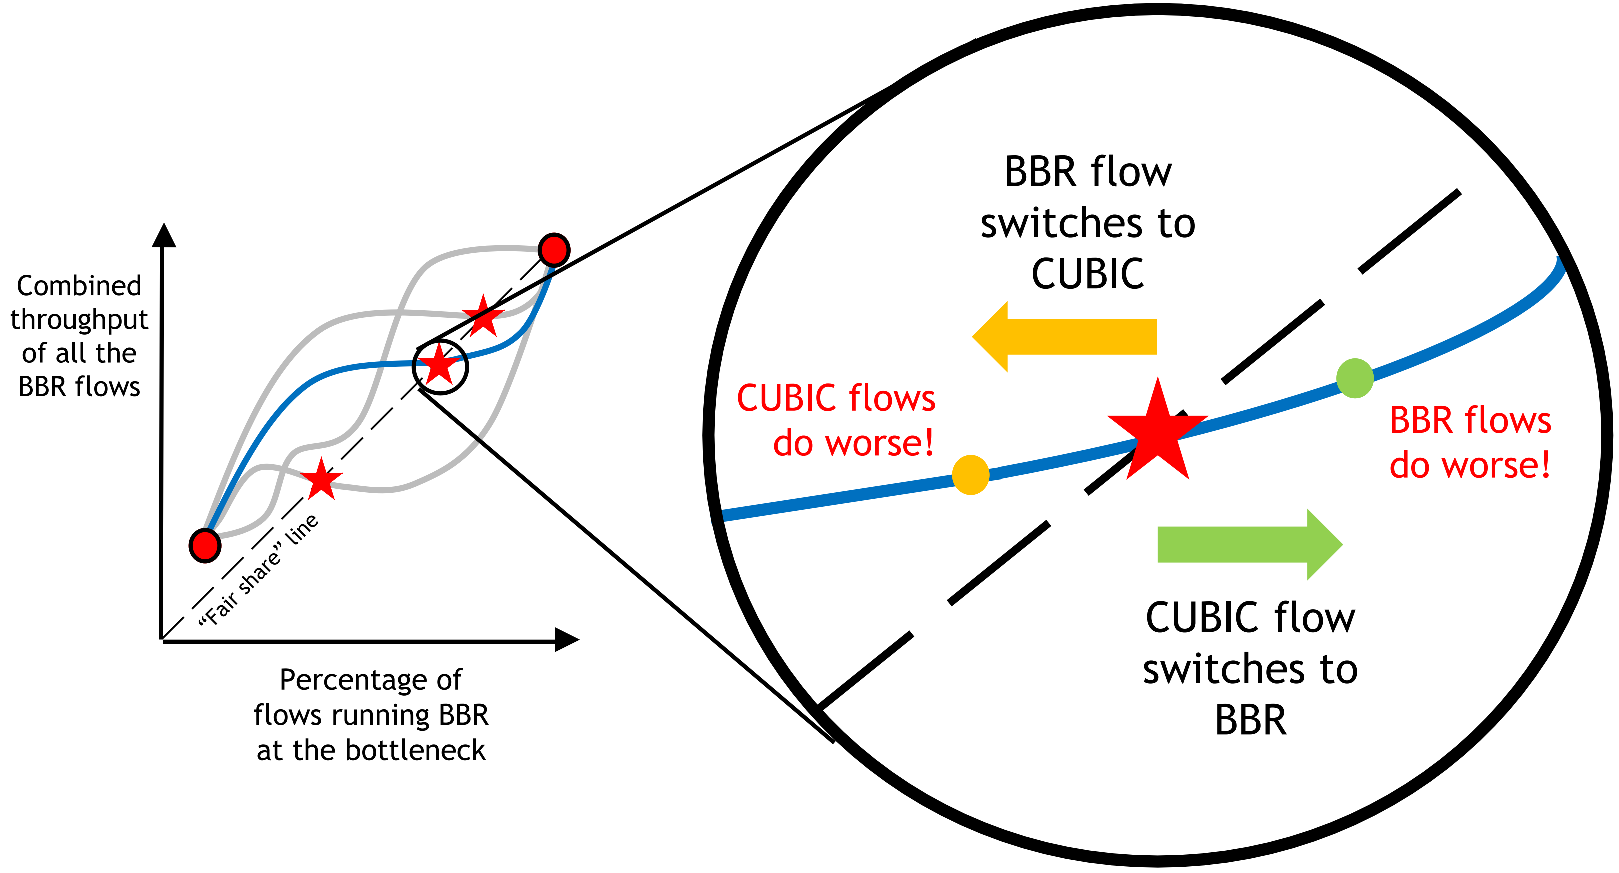

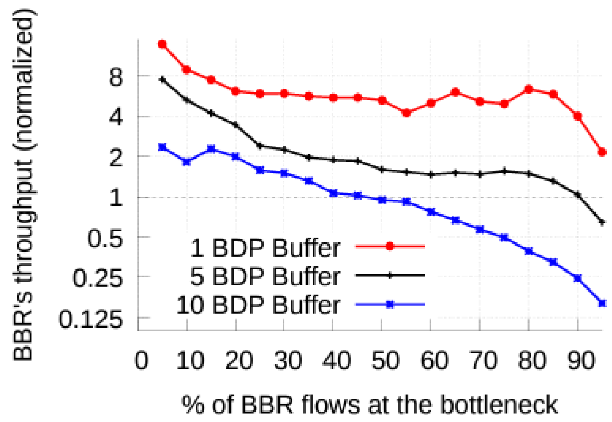

Fig. 2. BBR's throughput vs. % of BBR flows.

We can exhaustively prove a NE will always exist when 2 flows compete.

The proof is based on simple observations made by other measurement papers on how CUBIC and BBR compete.

Table 2. Outcomes in a two-flow game. ($RTT_1 > RTT_2$, winning strategies are **highlighted**)

	$Buff < T_2$				$T_2 < Buff < T_3$				$T_3 < Buff$			
Scenario	Strategies		Outcome		Strategies		Outcome		Strategies		Outcome	
	S_1	S_2	S_1	S_2	S_1	S_2	S_1	S_2	S_1	S_2	S_1	S_2
1	C	**C**	L	W	C	**C**	L	W	C	**C**	L	W
2	C	**B**	L	W	**C**	B	W	L	**C**	B	W	L
3	**B**	C	W	L	**B**	C	W	L	B	**C**	L	W
4	**B**	B	W	L	**B**	B	W	L	**B**	B	W	L
Nash Equilibrium	$(B,*)$				$(B,*)$				$(*,C)$			

Empirical validation

Checking the claims of this conjecture in a limited state space

- NE was computed in 6, 9 and 12 flow systems with each third of the flows having RTTs 20, 50 and 80 ms respectively. All flows shared a common bottleneck with a fixed link speed.
- For a given number of flows and a network configuration we ran all the possible combinations of flows running either CUBIC or BBR.

CBC

Is the NE if:

In BBC, flow 1 gets worse throughput
In CCC, flow 2 gets worse throughput
In CBB, flow 3 gets worse throughput

Properties of observed NE

- In each case, we observed exactly one Nash Equilibrium
- In each Nash Equilibrium, when the flows were sorted by the RTT, CUBIC was always picked by the smallest RTT flows

(CCC…)(…BBB)

m flows (n-m) flows

- That is, if the NE for 6 flows is when 50% of the flows are running CUBIC, the NE is at:

C C C B B B

RTTs: 20 ms 50 ms 80 ms

Effect of buffer size and link speed on the NE

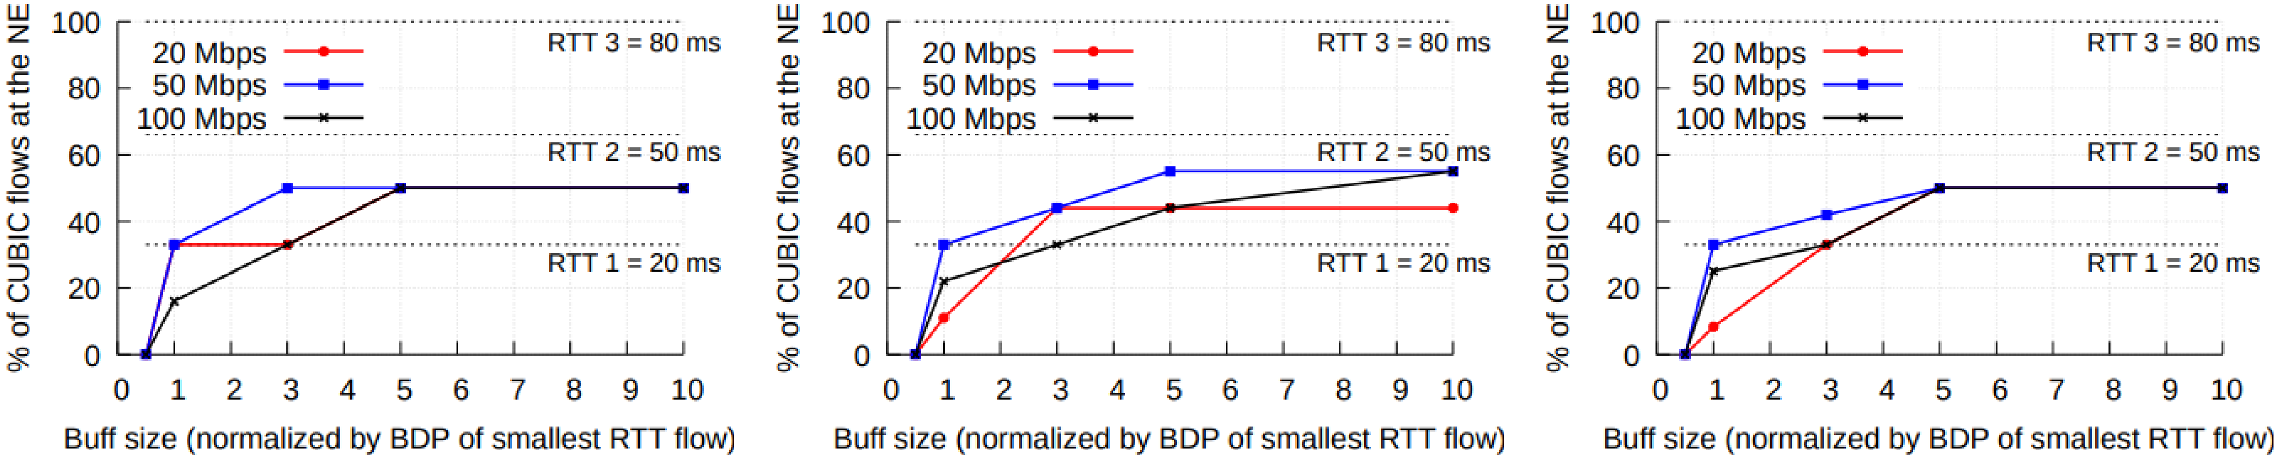

(a) 6 flows (b) 9 flows (c) 12 flows

Fig. 4. The effect of link capacity and number of flows on the Nash Equilibrium. RTTs 20 ms, 50 ms and 80 ms.

Predictably, buffer size had the biggest impact on the CC distribution at the NE

Effect of RTT distribution on the NE

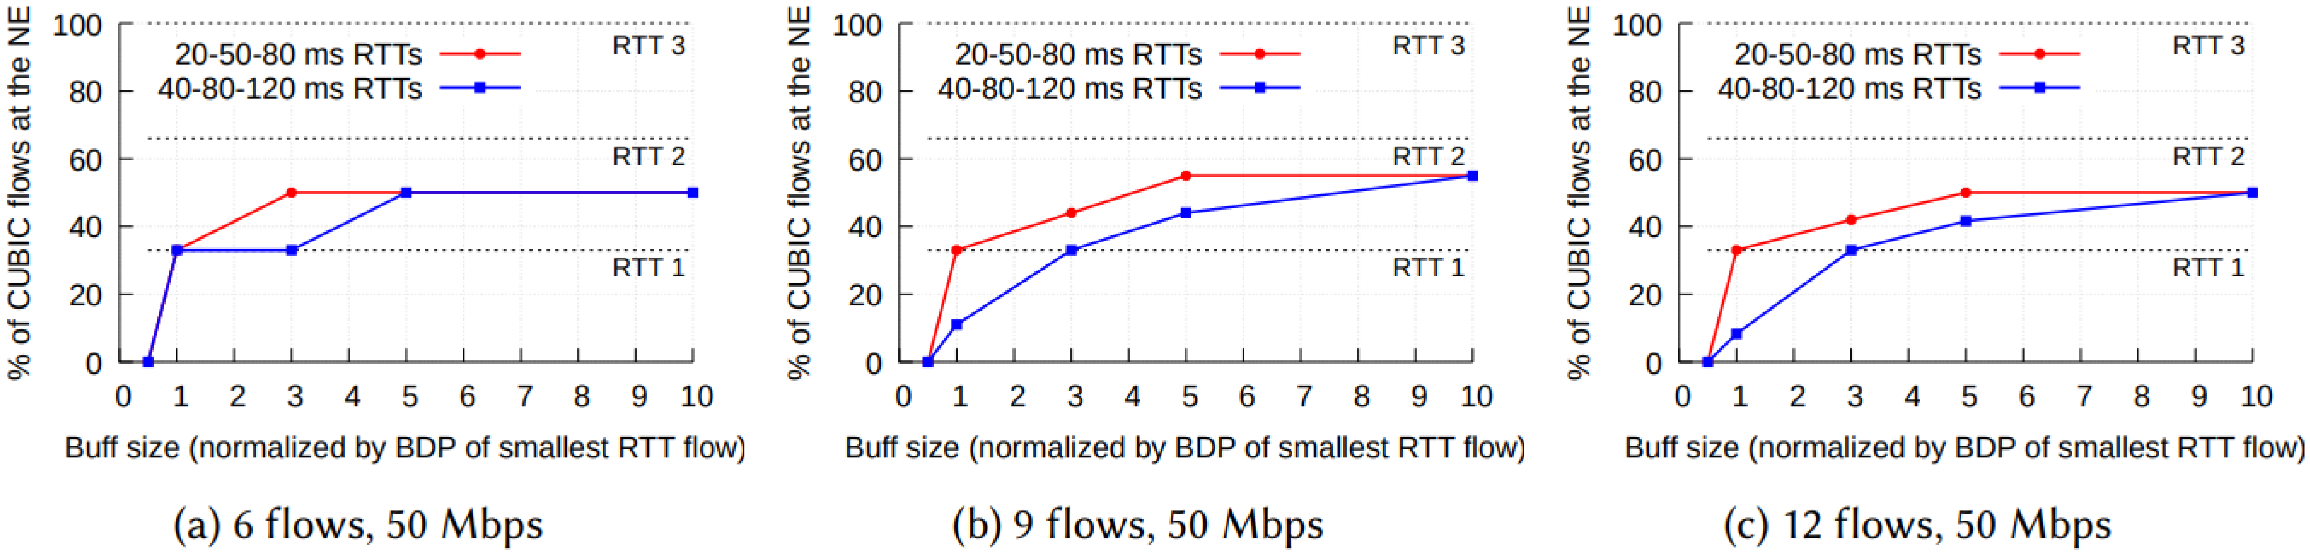

Fig. 5. The effect of the RTT distribution on the Nash Equilibrium.

RTT distribution had little effect of the where the NE was

Conclusion

- Despite BBR's current throughput benefits, CUBIC is unlikely to disappear soon from the Internet
- The Internet is likely to remain a heterogeneous mix of congestion control algorithms
- TCP performance is highly contextual
- However, the Internet does not follow economic game theory – it is not a given than the Internet will move towards the Nash Equilibrium.

Future work

- Formal proof for NE is a general n-flow game
- The effect of more complex network utility functions (delay, jitter)
- Effect on the NE in the presence of BBRv2, multi-hop paths, and AQMs
- NE in very deep buffers and a large number of flows:

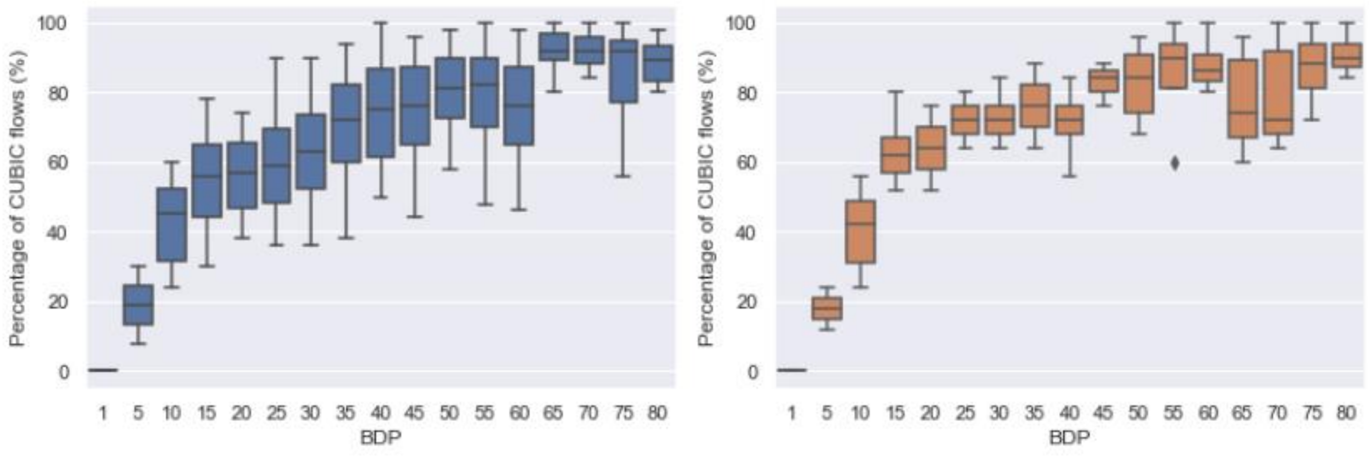

(a) 100Mbps, 50 flows (b) 50Mbps, 25 flows

Thank you for your time!

Please feel free to get in touch with me at

ayush@comp.nus.edu.sg

Extra slides

Observation 1. When competing at the bottleneck where the buffer is deep, CUBIC tends to have higher throughput than BBR; the converse is true when the buffer is shallow.

Observation 2. When a single BBR flow competes with a single CUBIC flow at a bottleneck, there must exist some threshold bottleneck buffer size T_{fair} such that when the bottleneck buffer size $Buff < T_{fair}$, the BBR flow gets higher throughput than the CUBIC flow and when $Buff > T_{fair}$, the CUBIC flow gets higher throughput than BBR.

Observation 3. As the percentage of BBR flows at the bottleneck increases, the per-flow average throughput of BBR flows at that bottleneck decreases.

Observation 4. When two BBR flows compete at a bottleneck, the flow with a longer RTT will get higher bandwidth than the flow with a shorter RTT. When two CUBIC flows compete at a bottleneck, the flow with a shorter RTT will get higher bandwidth than the flow with a longer RTT.

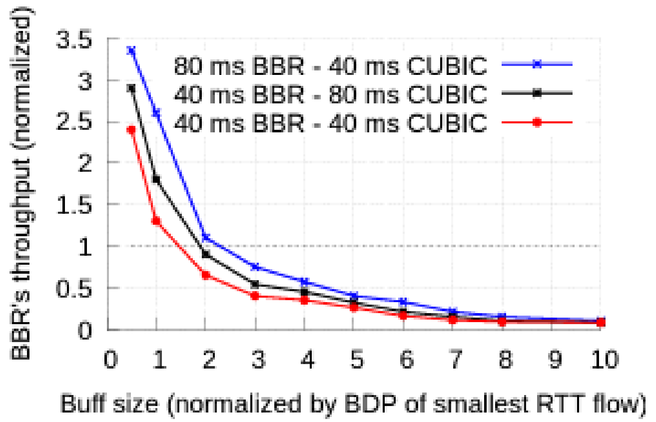

Fig. 1. BBR throughput vs. bottleneck buffer size.

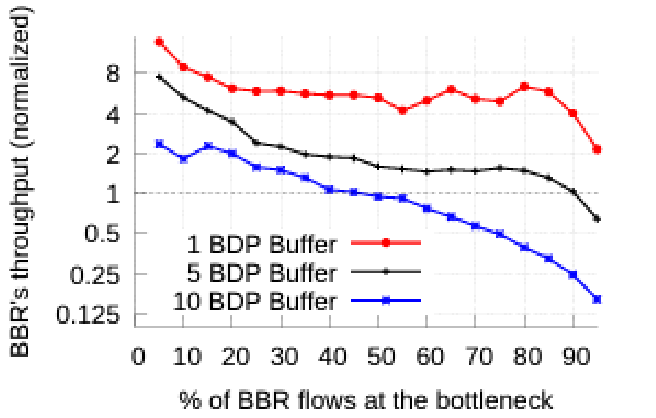

Fig. 2. BBR's throughput vs. % of BBR flows.