The Great Internet TCP Congestion Control Census

SIGMETRICS 2020

Ayush Mishra[1], Sun Xiangpeng[1], Atishya Jain[2],
Sameer Pande[2], Raj Joshi[1], Ben Leong[1]

[1] National University of Singapore
[2] Indian Institute of Technology, Delhi

The Internet is the world's largest distributed system.

More than 30 years of Congestion Control on the Internet

The Internet has been evolving both in terms of provisioning and the nature of the applications it supports.

And so has congestion control!

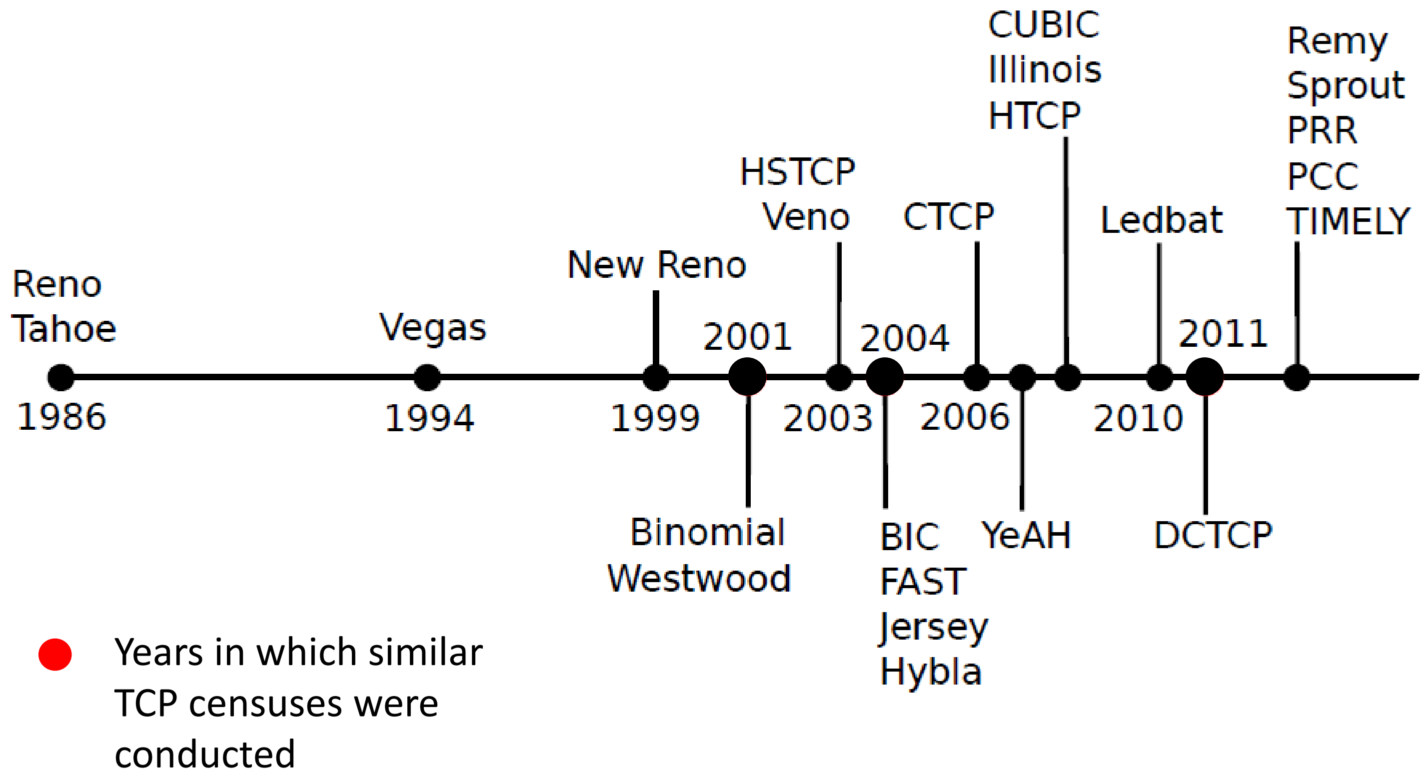

More than 30 years of Congestion Control on the Internet

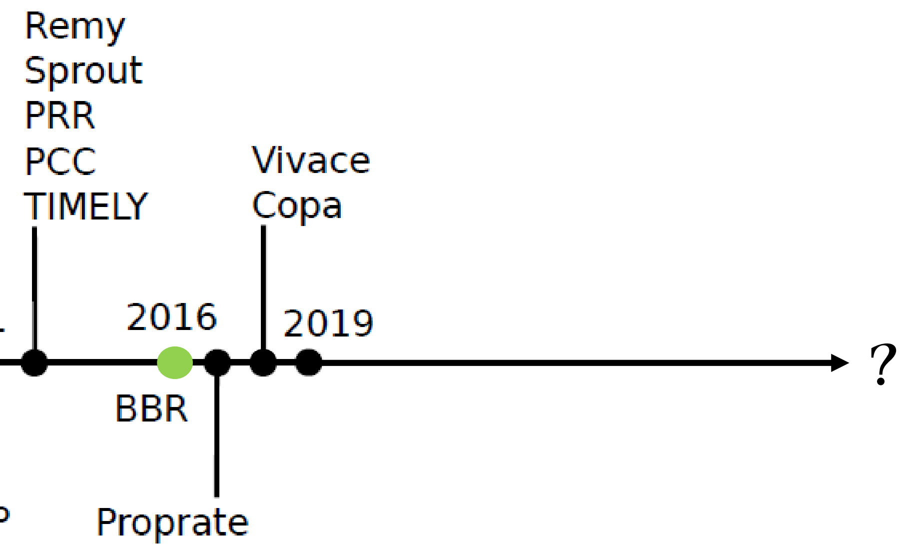

In 2016, Google introduce their own Congestion Controller, BBR. It has since been introduced in the Linux kernel and deployed by Google to support services like YouTube and GCP.

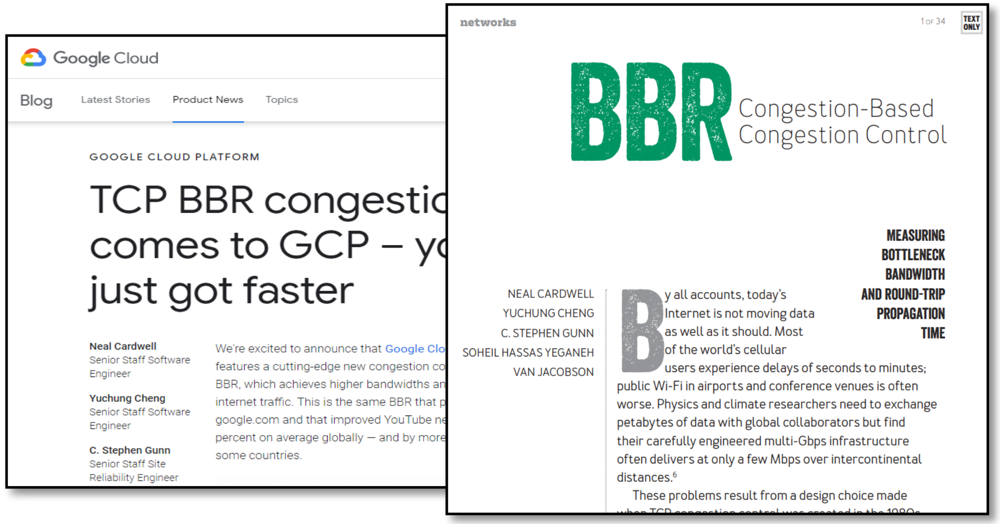

This marks a paradigm shift in Congestion Control.

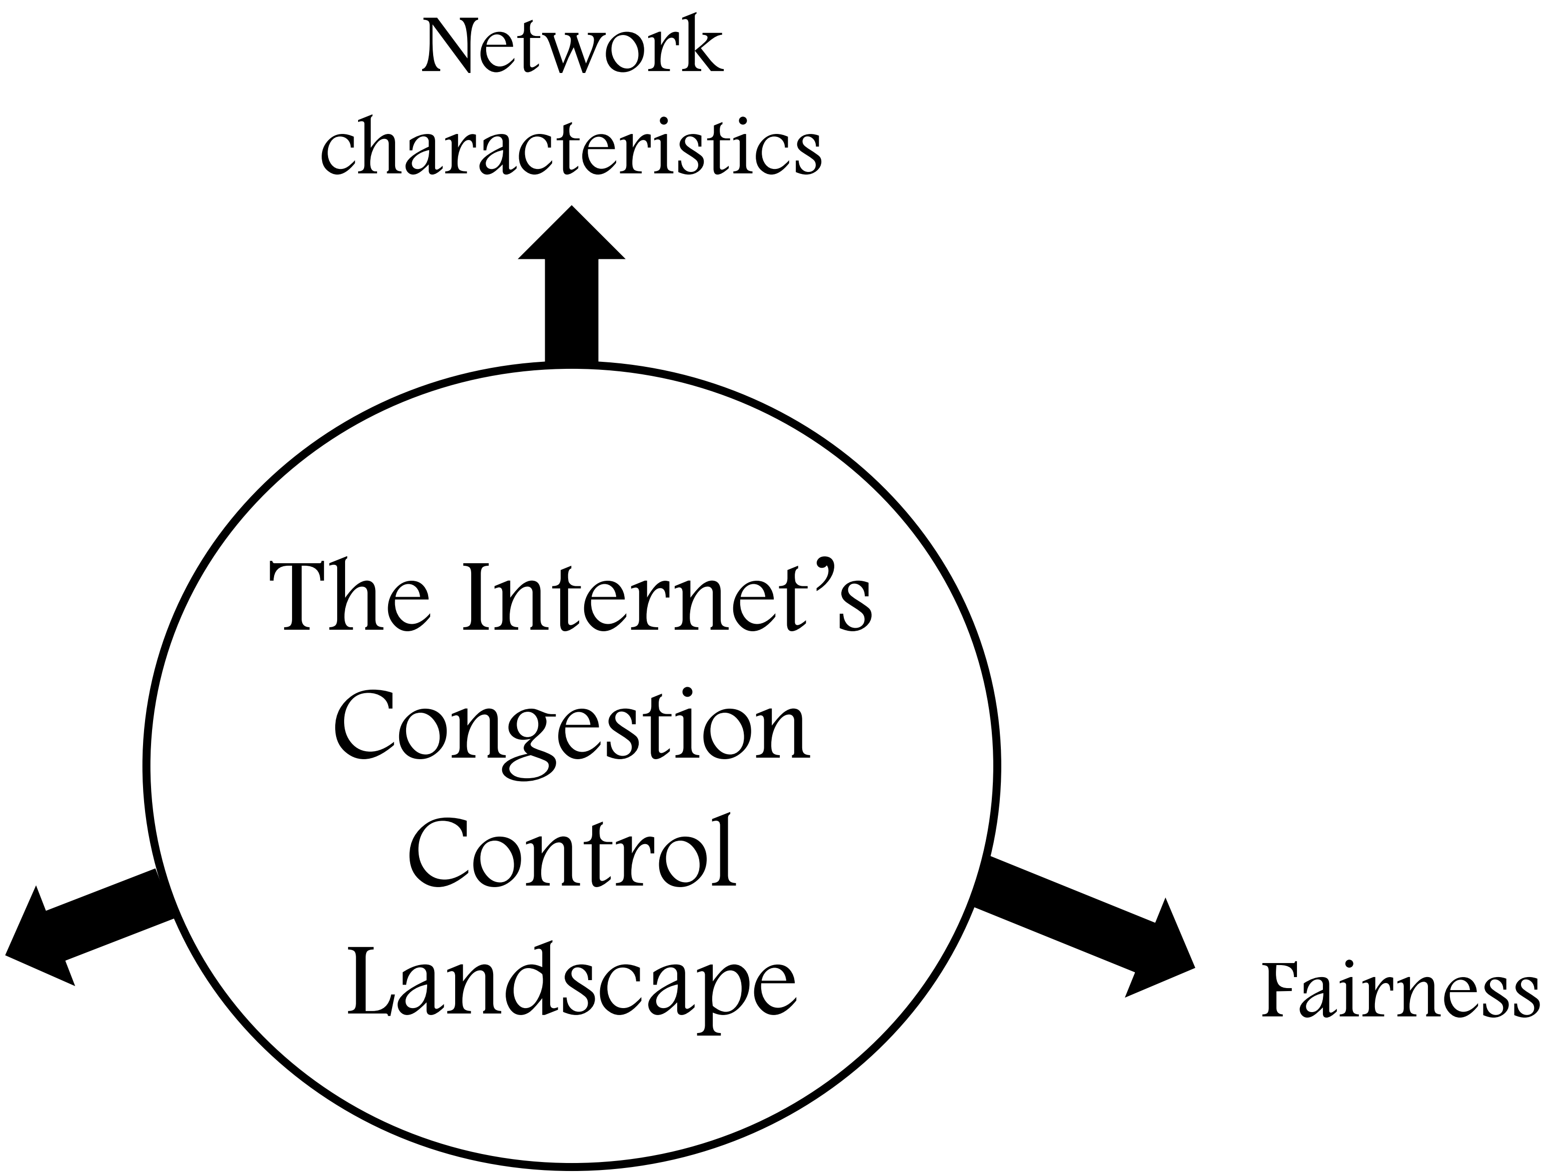
Network characteristics
The Internet's Congestion Control Landscape
Provisioning choices
Fairness

Conduct a Congestion Control Census among the 20,000 most popular websites* on the Internet.

* According to their Alexa 2019 rankings

This is a non-trivial problem.

1. Isolating the Internet's network dynamics
2. Extracting a common feature from a variety of congestion control algorithm
3. Identifying congestion control algorithm behavior within short HTTP page downloads

Our Solution: **Gordon**

1. Isolating the Internet's network dynamics

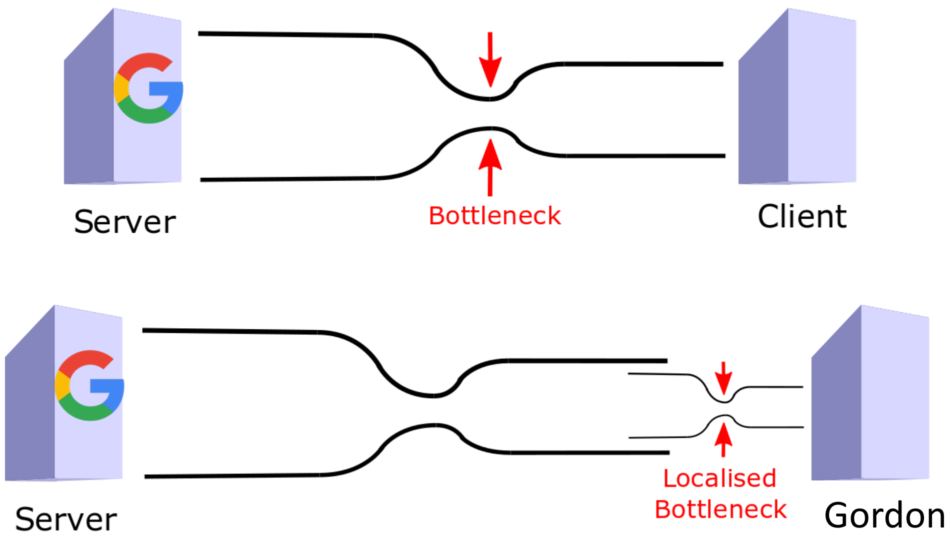

Most network dynamics, like change in bottleneck bandwidth and packet losses happen at the connection bottleneck. Localizing the bottleneck allows us to have better control over the connection.

2. Extracting a common feature from a variety of congestion control algorithm

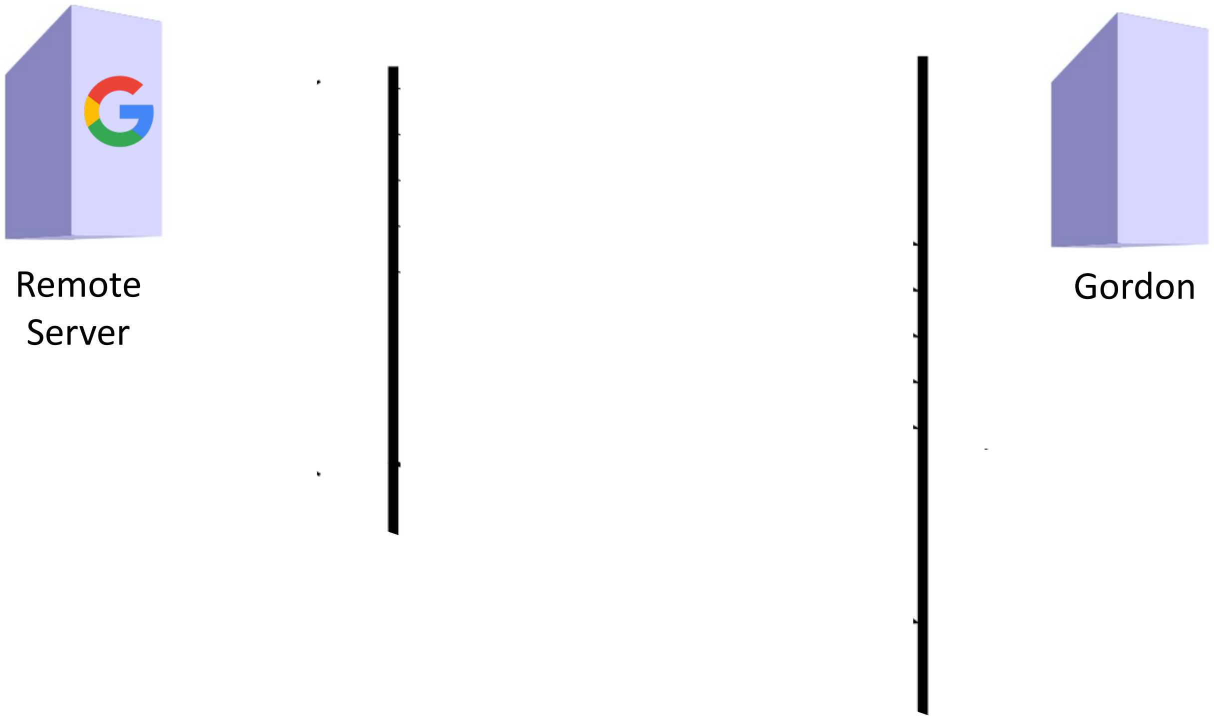

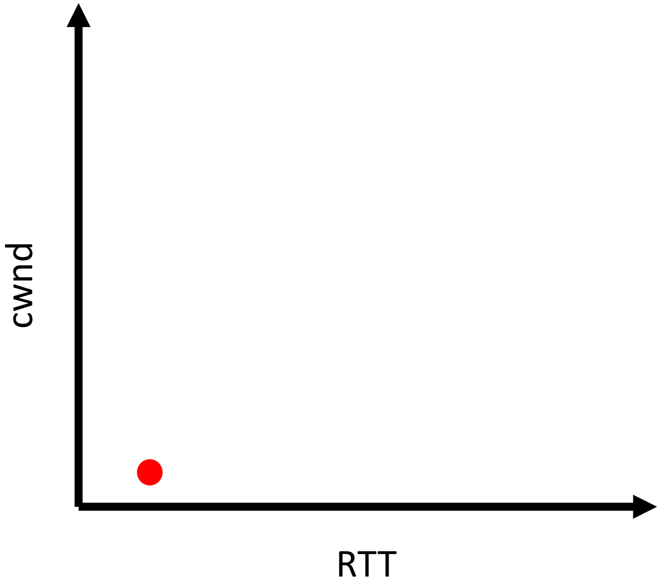

2. Extracting a common feature from a variety of congestion control algorithm

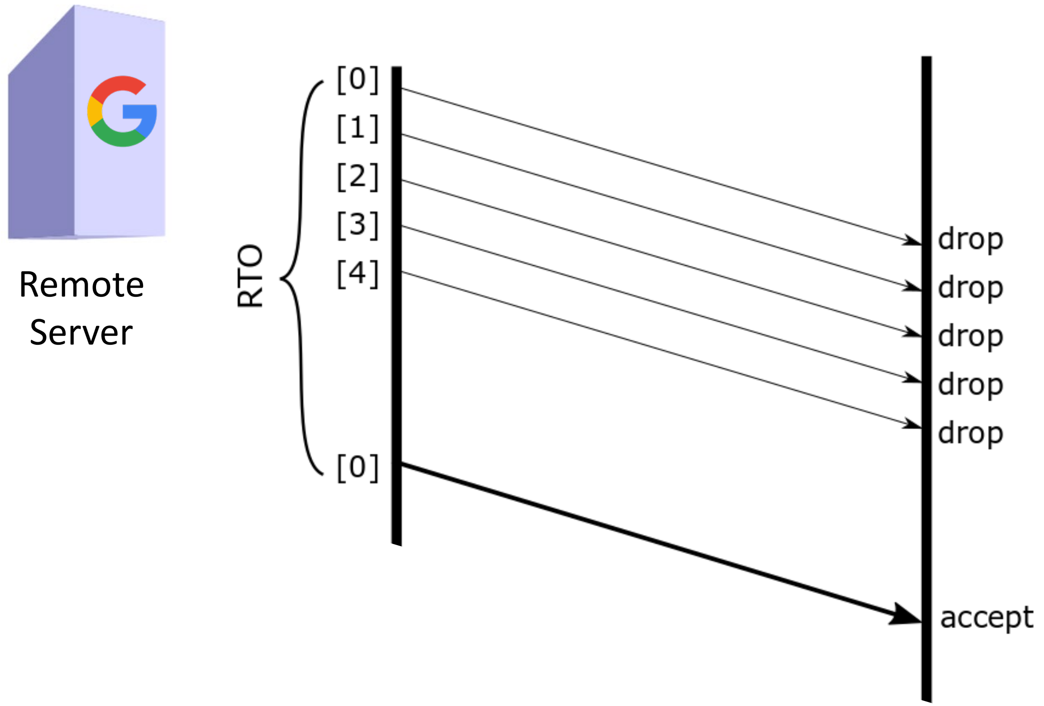

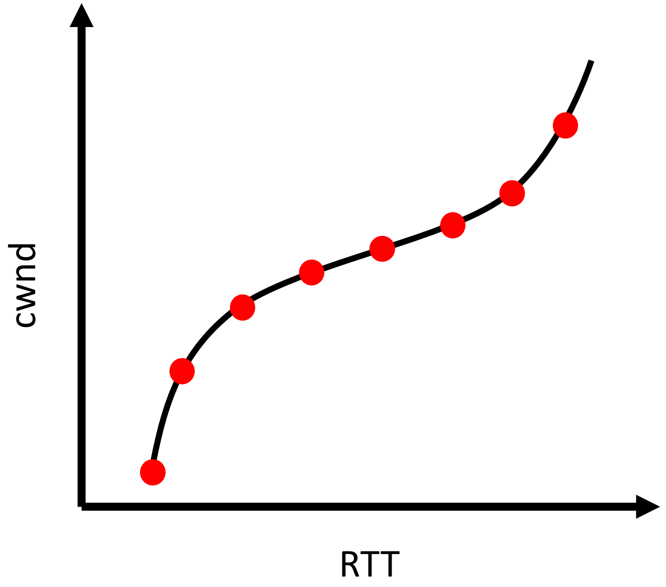

Cwnd evolution graph

3. Identifying congestion control algorithms within short HTTP page downloads

Crawling target websites for large web objects

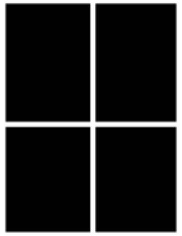

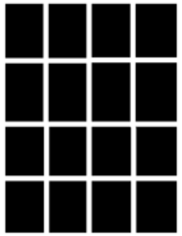

Measuring the cwnd in packets and using a smaller MTU

Final sequence of network stimuli

"Network profile":

1. Packet drop at the first cwnd that exceeds 80 packets
2. Bandwidth change after receiving 1500 packets
3. Emulating an RTT of 100 ms

Optimal page size for this network profile: **165 kb**

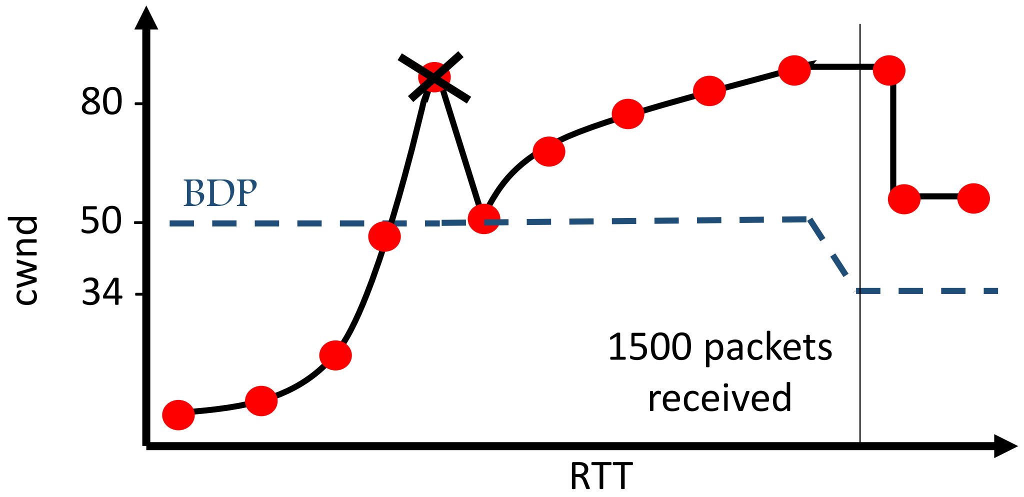

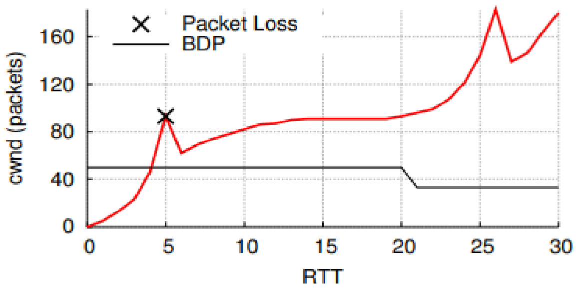

(a) CUBIC

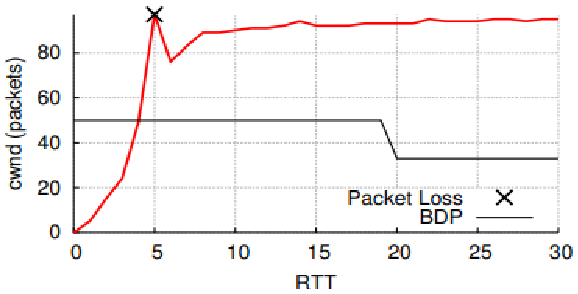

(b) BIC

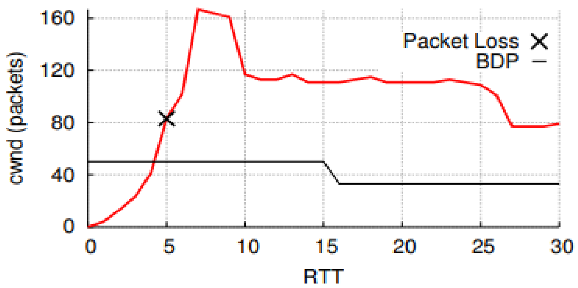

(c) BBR

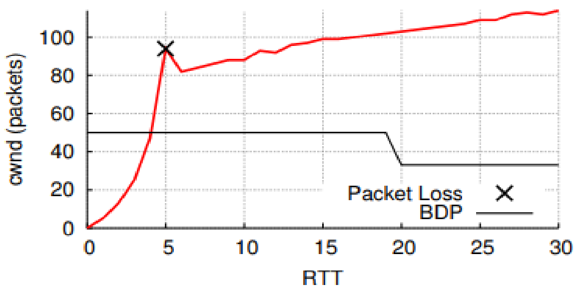

(d) Scalable

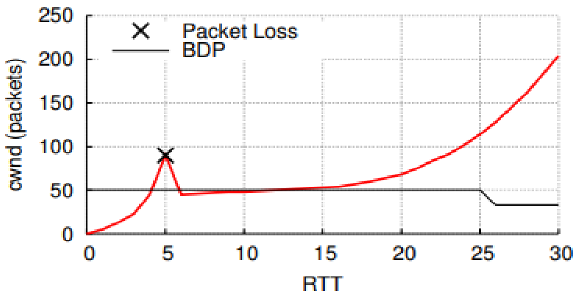

(e) HTCP

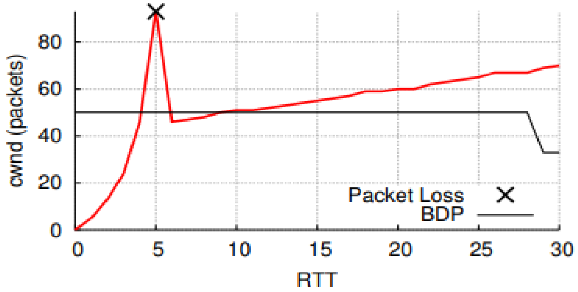

(f) New Reno

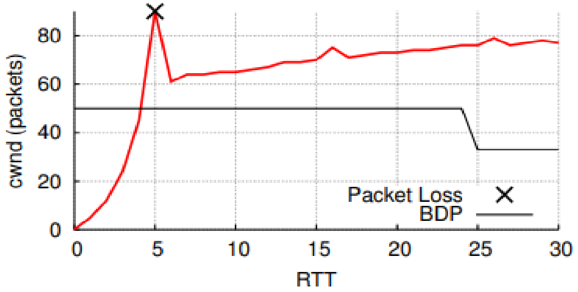

(g) Illinois

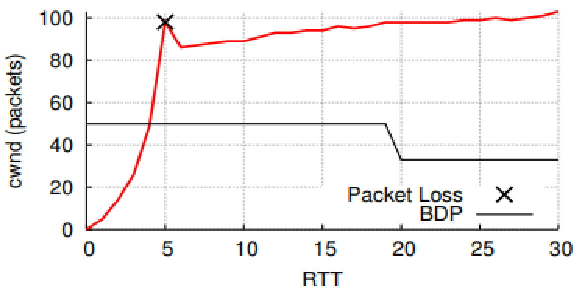

(h) YeAH

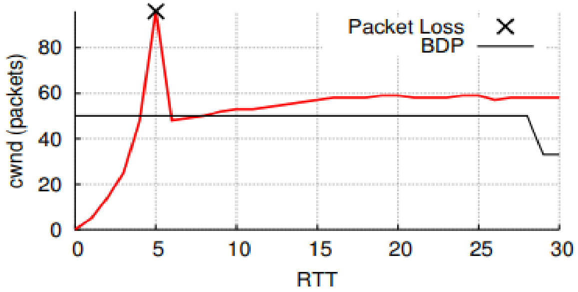

(i) Vegas

Measurement Results

Measurement and classification Accuracy

	Classified as									
	BBR	CUBIC	BIC	HTCP	Scalable	YeAH	Vegas	New Reno/Veno	CTCP/Illinois	Unknown
BBR	**98%**	0%	0%	0%	0%	0%	0%	0%	0%	2%
CUBIC	0%	**95%**	0%	0%	0%	0%	0%	0%	0%	5%
BIC	0%	9%	**91%**	0%	0%	0%	0%	0%	0%	0%
HTCP	0%	0%	0%	**95%**	0%	0%	0%	0%	0%	5%
Scalable	0%	0%	0%	0%	**98%**	0%	0%	0%	0%	2%
YeAH	0%	0%	2%	0%	0%	**98%**	0%	0%	0%	0%
Vegas	0%	0%	0%	0%	0%	0%	**94%**	6%	0%	0%
New Reno/Veno	0%	0%	0%	0%	0%	0%	0%	**96%**	0%	4%
CTCP/Illinois	0%	0%	3%	0%	0%	0%	0%	0%	**94%**	3%

Measurement details

- The measurements were done between 11 July 2019 and 17 October 2019 from servers in Singapore, Mumbai, Paris, Sao Paulo and Ohio.
- Given our network profile, 16% of pages were less than optimal in size of 165 Kb (Short flows)
- We also came across 1,302 Unresponsive websites

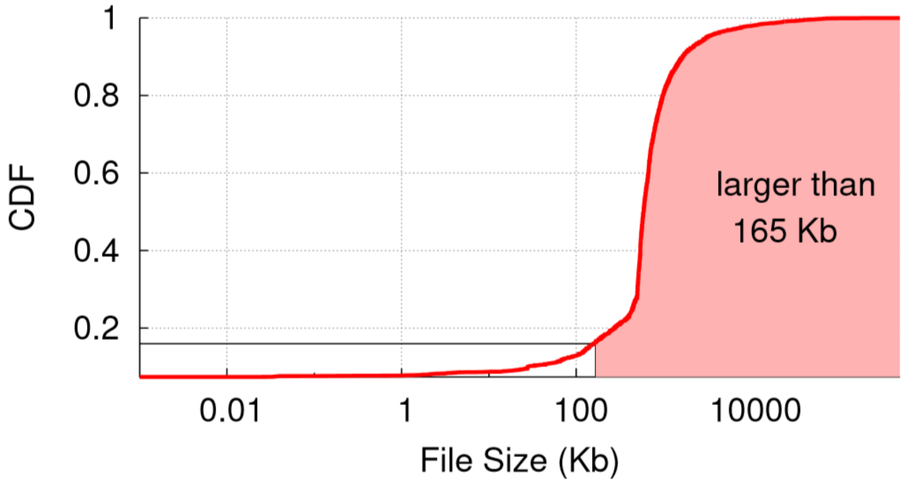

Consolidated numbers

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%

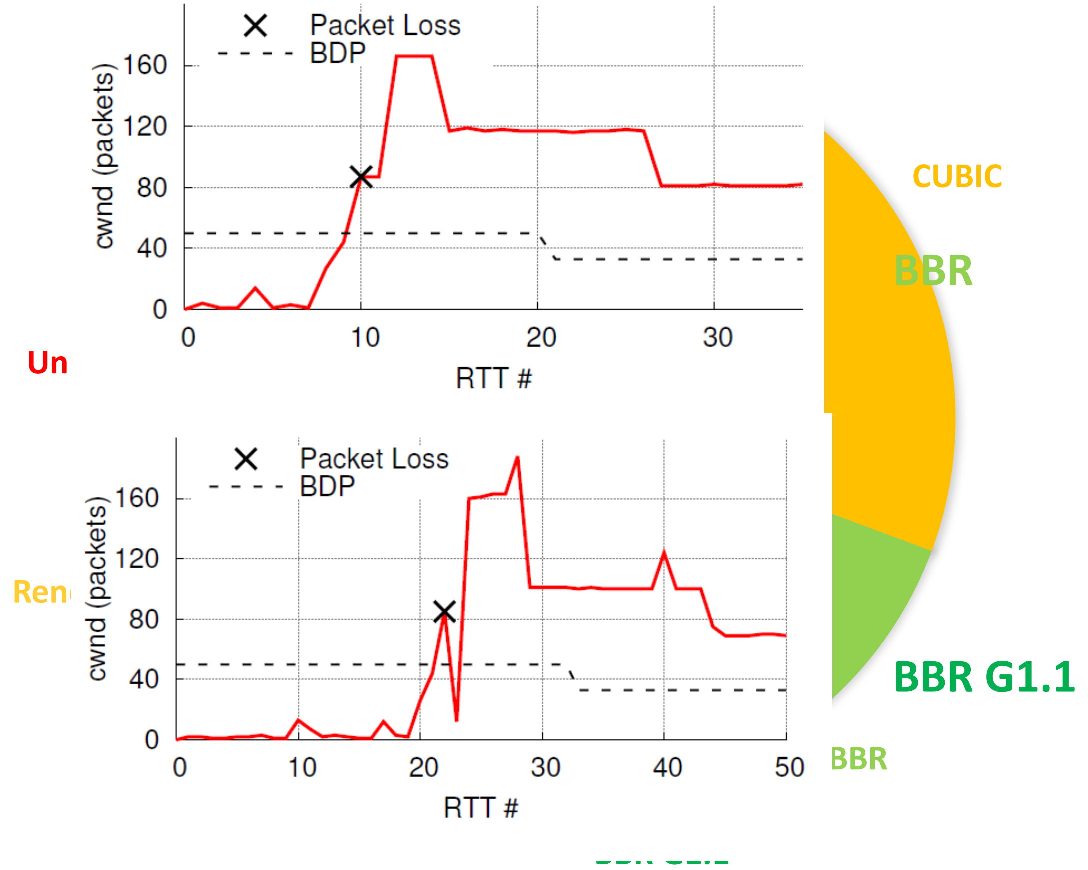

Distribution by Popularity and Traffic share

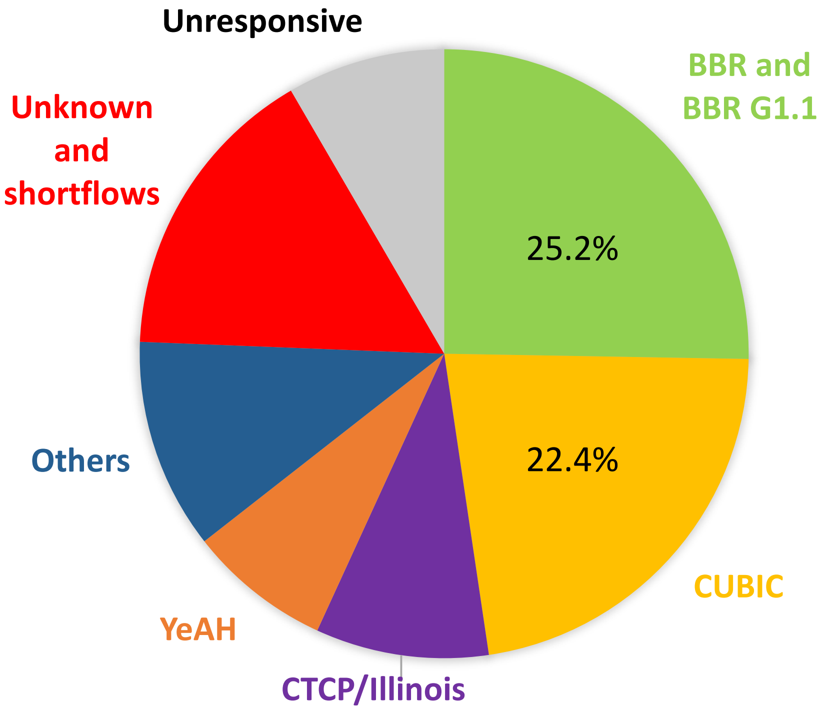

Share of congestion control algorithms deployed by website count in the Alexa Top 250 websites

- Among the top 250 Alexa websites, BBR has a larger share by website count than Cubic
- In terms of traffic share, BBR is now contributing to **more than 40%** of the downstream traffic on the Internet!

Site	Downstream traffic share	Variant
Amazon Prime	3.69%	CUBIC
Netflix	15%	CUBIC
YouTube	11.35%	BBR
Other Google sites	28%	BBR
Steam downloads	2.84%	BBR

Looking closer at the unclassified websites

- We had a total of 6,330 (31.65%) of websites that were unclassified
- We ran a variety of network profiles on these websites to infer something about their congestion control mechanism

Type	React to Packet Loss?	React to BDP?	Websites (share)
AkamaiCC	✗	✓	1,103 (5.52%)
Unknown Akamai	✗	?	157 (0.79%)
Unknown	✗	?	493 (2.47%)
	✓	?	1,782 (8.91%)
Short flows	✓	?	1,493 (7.47%)
Unresponsive	?	?	1,302 (6.51%)
Total			6,330 (31.65%)

Looking closer at the unclassified websites: AkamaiCC

- Most websites hosted by Akamai had deployed a congestion control variant that didn't react to loss. It's cwnd closely followed the BDP of the network path

- It is likely this is a variant of FAST TCP

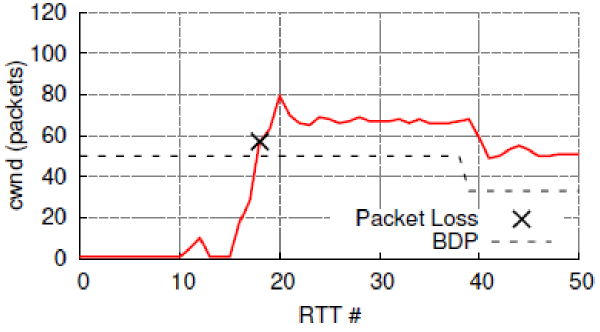

(a) Custom Network Profile 4 - Shape 1.

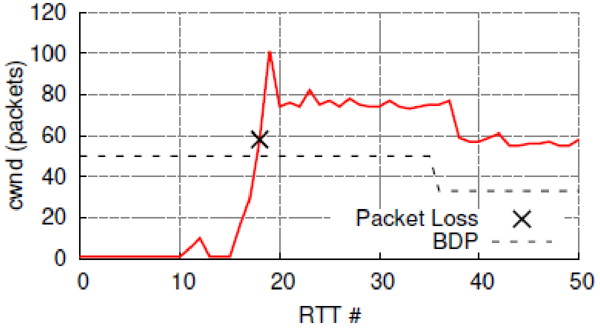

(b) Custom Network Profile 4 - Shape 2.

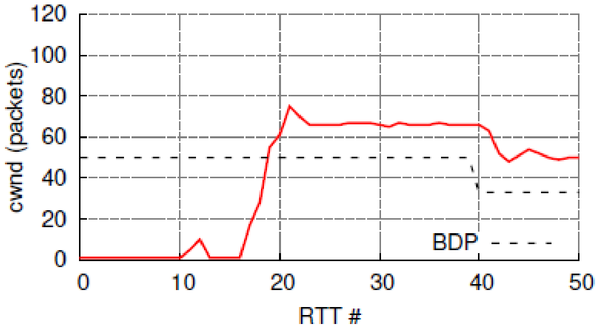

(c) Custom Network Profile 1 - Shape 1.

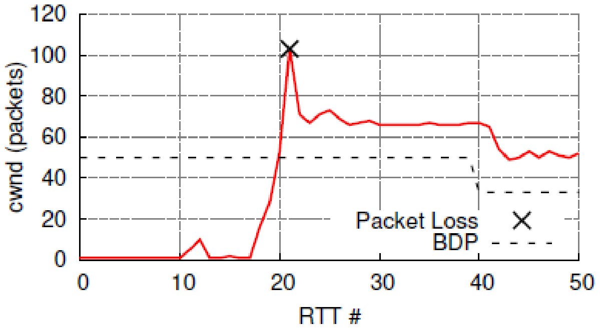

(d) Custom Network Profile 1 - Shape 2.

Looking closer at the unclassified websites

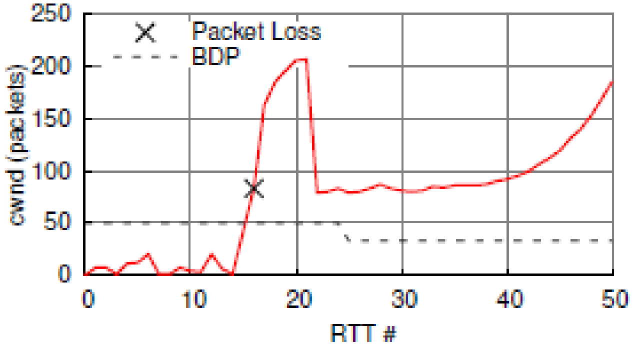

(a) amazon.com

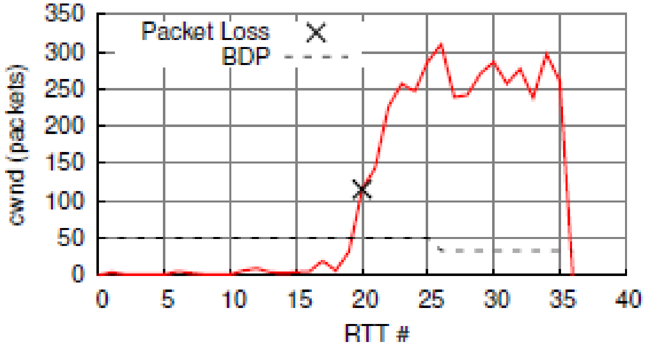

(b) zhihu.com

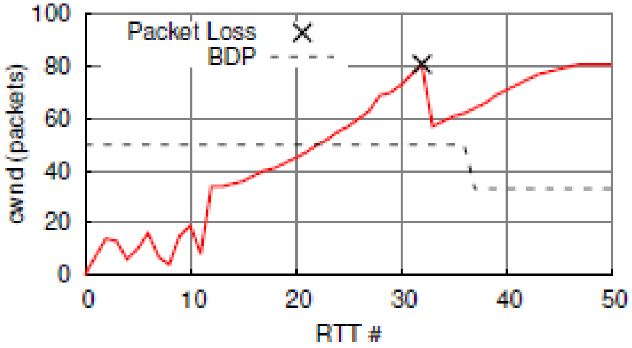

(c) yahoo.co.jp

Some other popular websites classified
as Unknown by Gordon

Evolution of the TCP Ecosystem

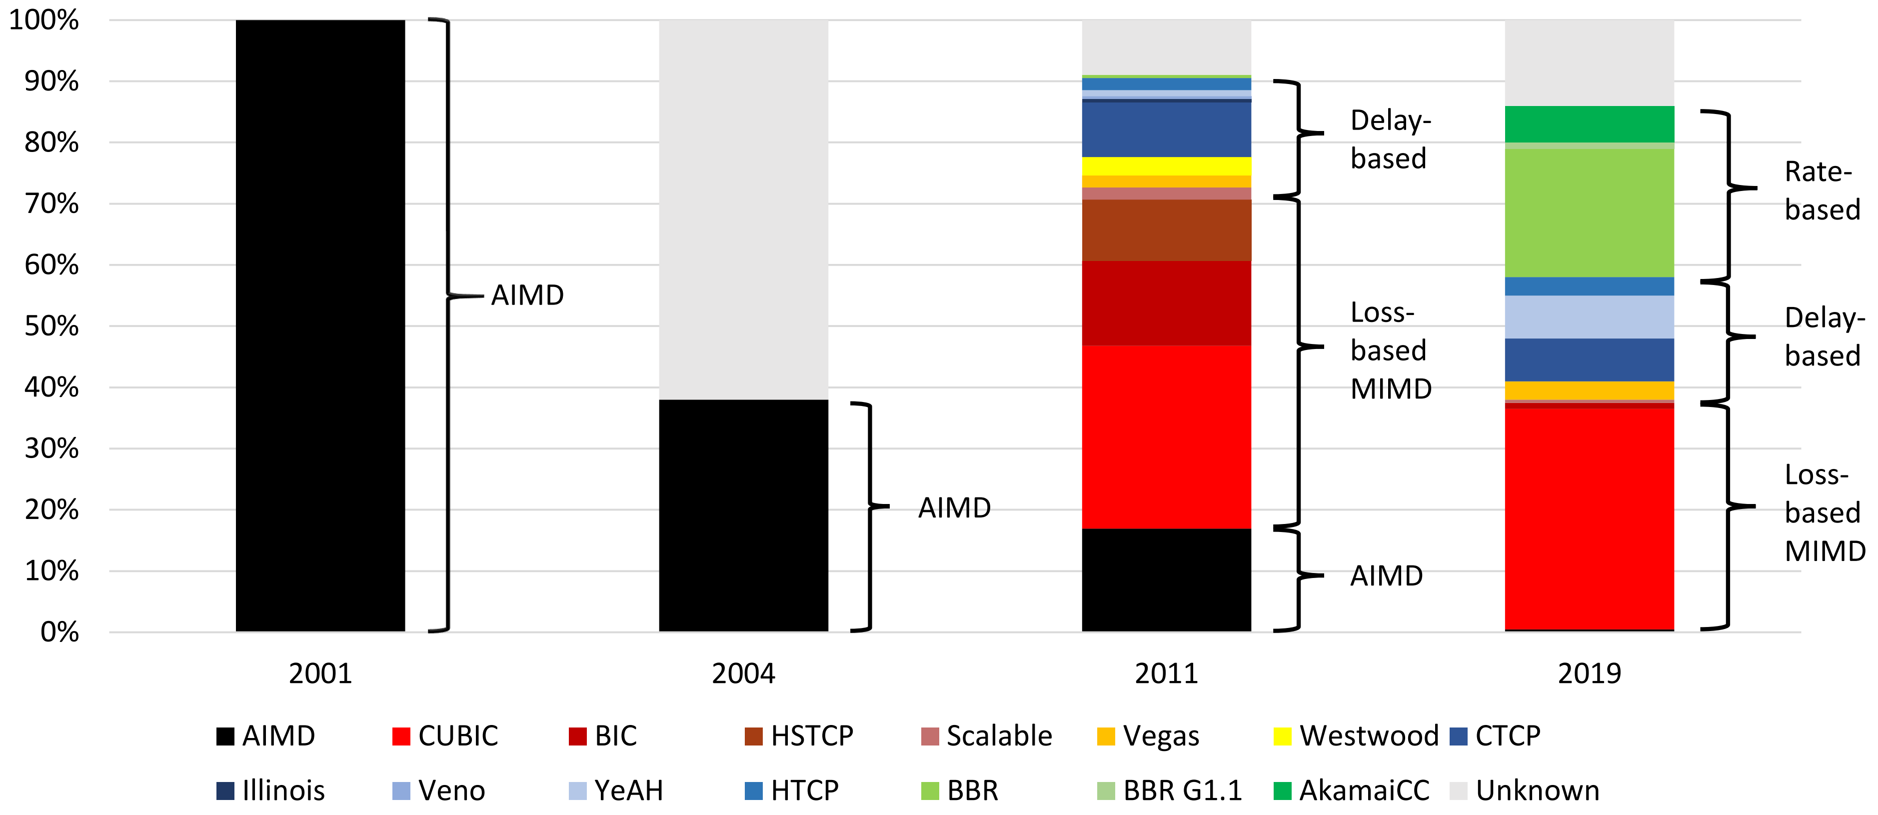

Future work

- Investigate Unresponsive websites to increase data set
- Experiment with other Network stimuli
- Extending the tool to detect sub-RTT behaviors
- Using clustering algorithms to better analyze the Unknown websites
- Identifying other rate-based algorithms, beyond BBR

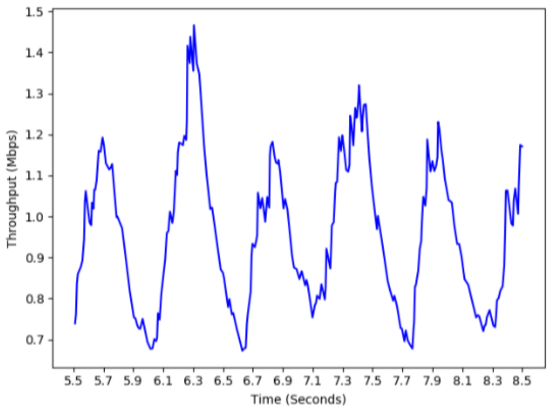

Copa

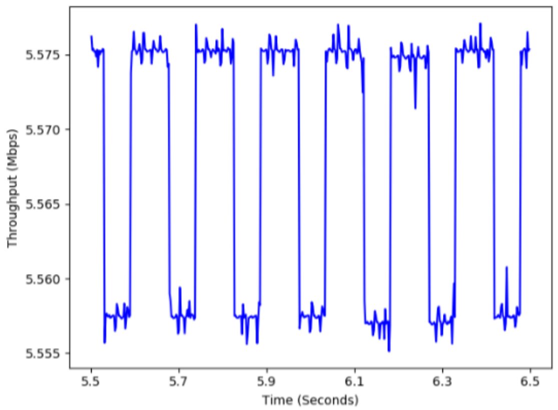

PCC-Vivace

Summary

- We measured the Alexa Top 20,000 on the Internet.
- BBR is poised to overtake CUBIC as the dominant TCP in the wild.
- The TCP Congestion control landscape is becoming increasingly heterogeneous.
- We are witnessing a paradigm shift in how Internet traffic behaves given the rise of rate-based algorithms. This represents a fundamental departure from how congestion control has traditionally been done on the Internet.

Looking forward to your questions and comments!

ayush@comp.nus.edu.sg

https://github.com/NUS-SNL/Gordon