


Keeping on Eye on Congestion Control in the Wild with *Nebby*

ACM SIGCOMM
August 4-8, 2024
Sydney, NSW, Australia


Ayush Mishra¹, Lakshay Rastogi², Raj Joshi^{1,3}, Ben Leong¹

¹National University of Singapore, ²IIT Kanpur, ³Harvard University



Background and Motivation

The makeup of the Internet's congestion control landscape influences how we think about



There is also a need to monitor the responsible deployment of congestion control algorithms (CCAs) on the Internet.

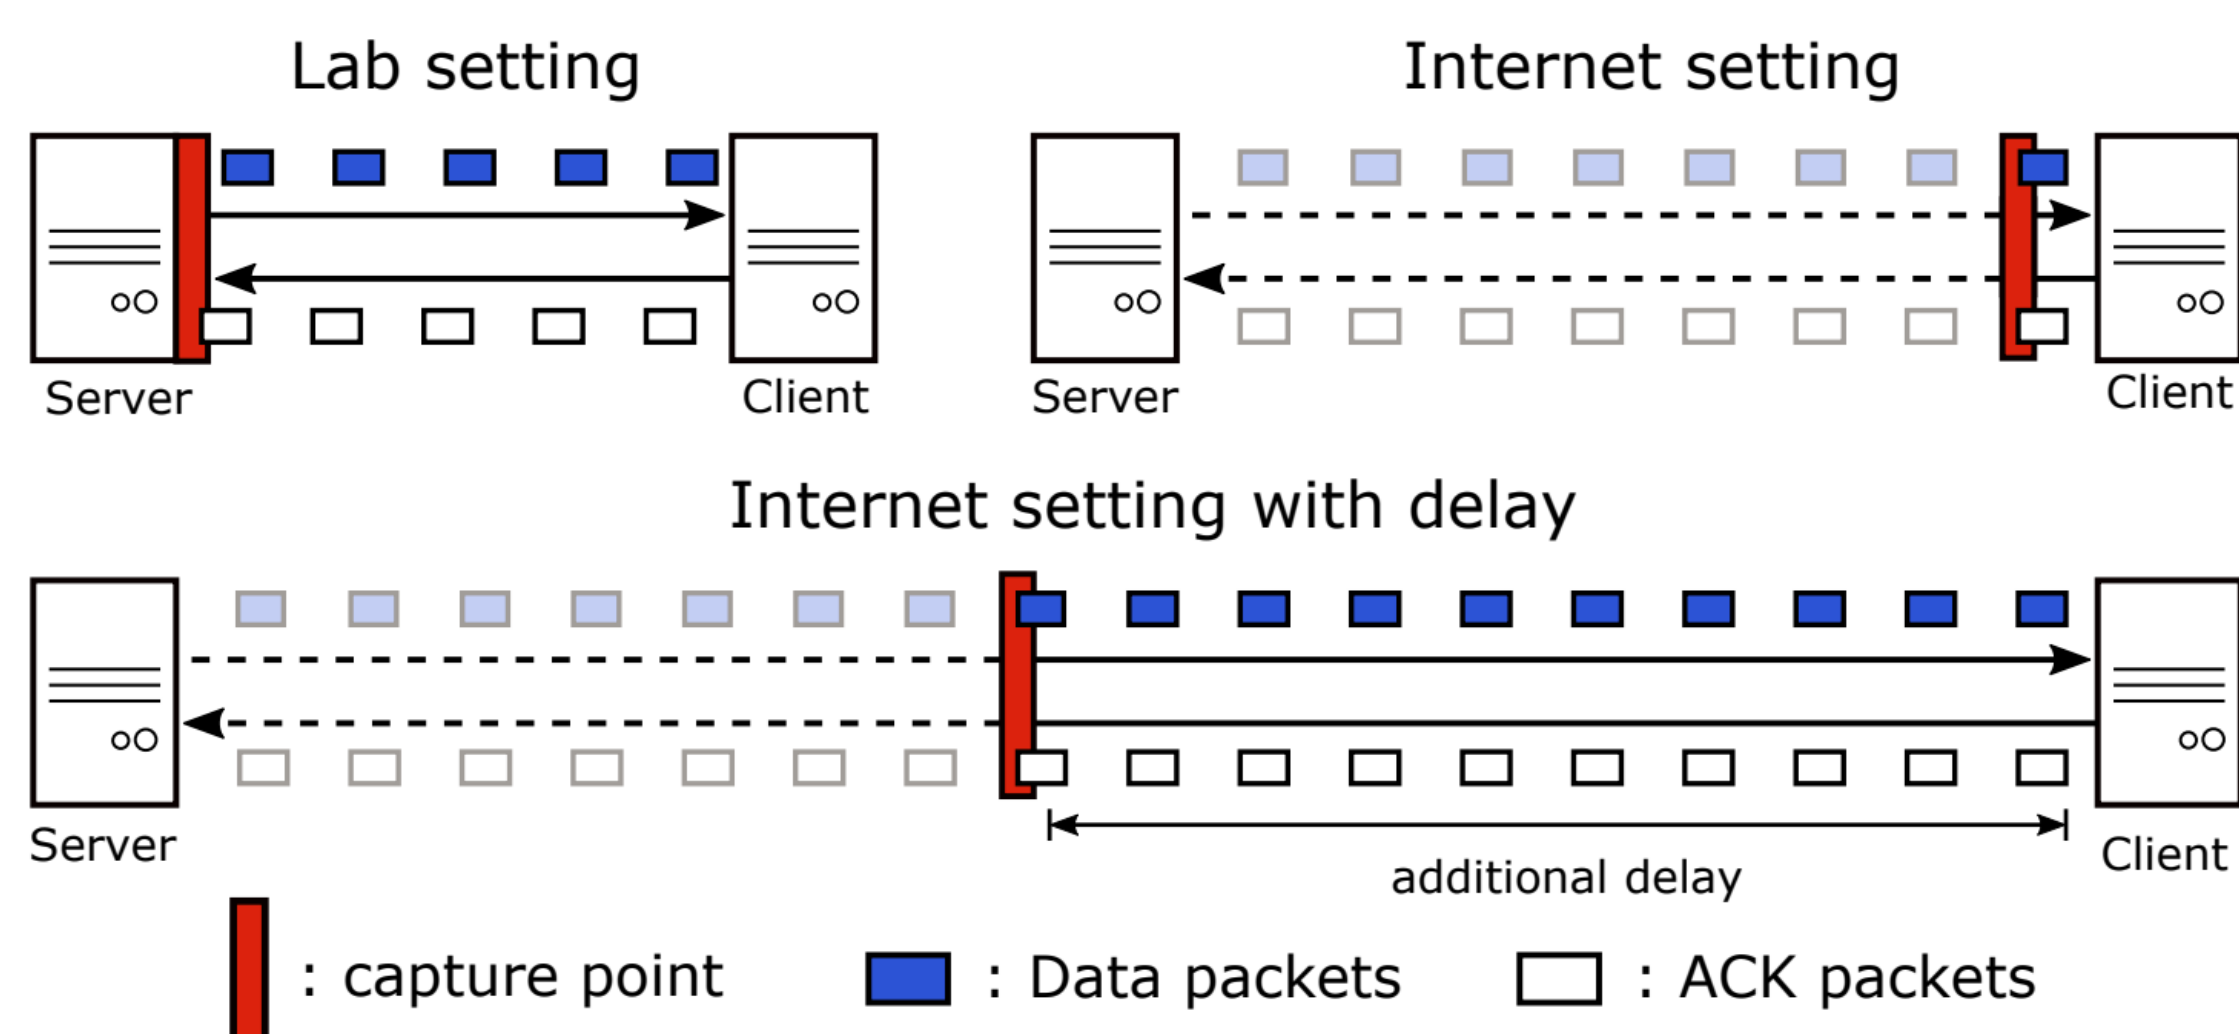
Measurement studies to study the deployment of different CCAs on the Internet have been done several times in the past, but every time we've had to re-invent the wheel because none of these methodologies have proven to be future-proof.

Our goal is to not only take stock of the Internet's CCA landscape in 2024 (5 years since the last study), but also do so in a future-proof way while meeting the following design goals:

- General and Extensible Measurement methodology**
The biggest challenge to a CCA identification tool is the continuous deployment of new CCAs on the Internet. For this reason, a CCA identification tool can't be ad-hoc
- Being client-agnostic**
A modern CCA identification tool also needs to work seamlessly with a large variety of clients in order to detect the deployment of different CCAs for different applications and application traffic
- Passive**
It also needs to be indistinguishable from a normal connection and be as passive as possible

Nebby's measurement methodology

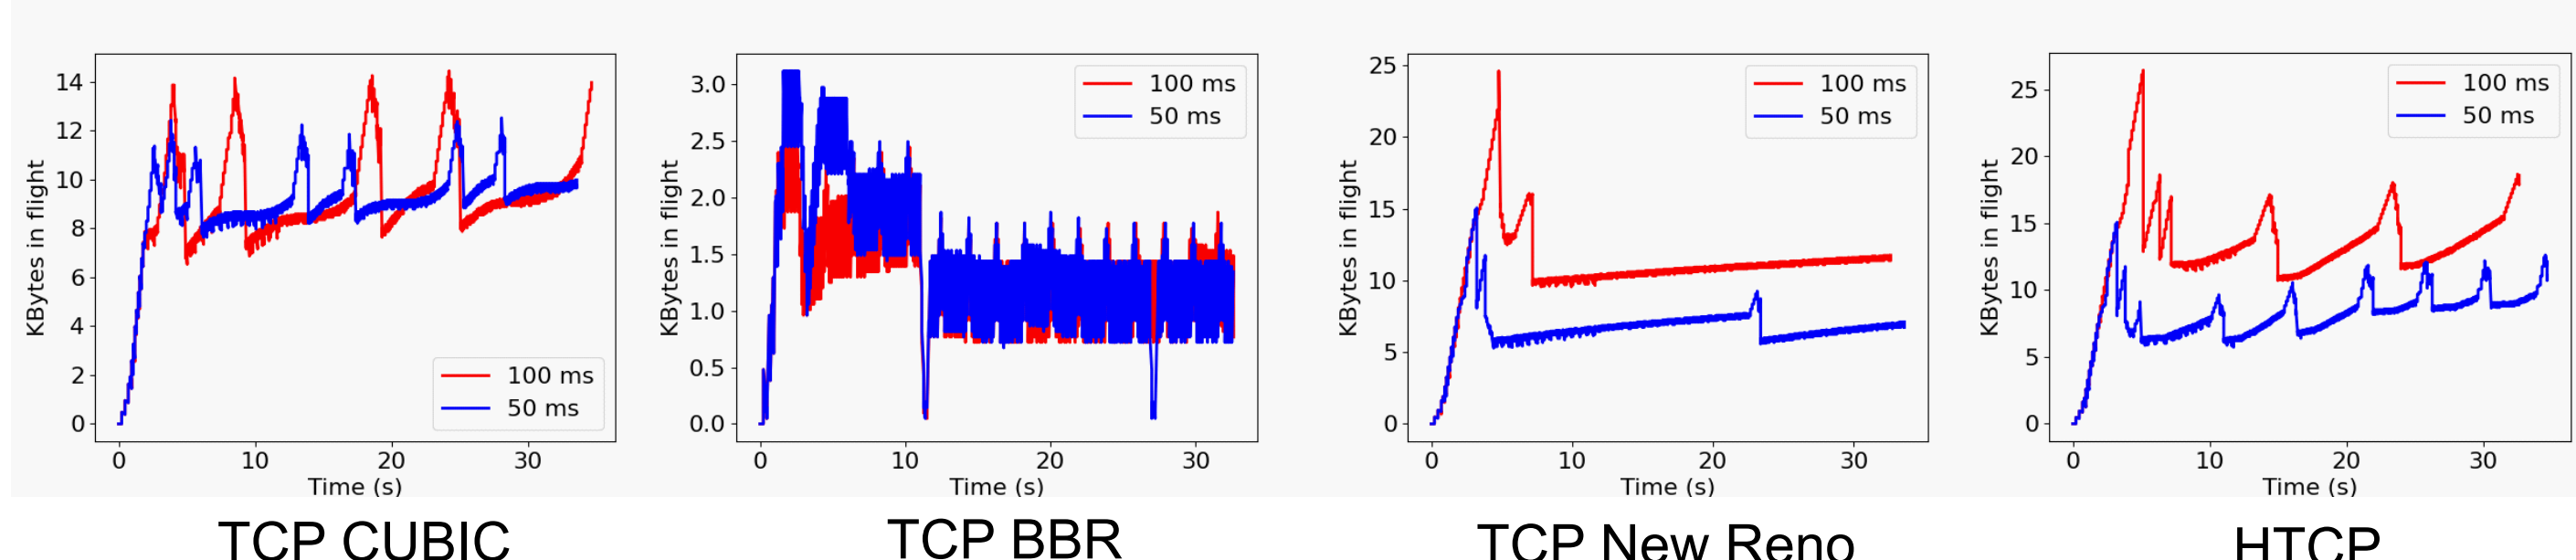
Nebby elects to measure a server's bytes-in-flight (BiF) during a connection to identify since it represents the operating point of both window-based and rate-based congestion control algorithms.



Using additional delay to increase the fraction of *visible* in-flight packets

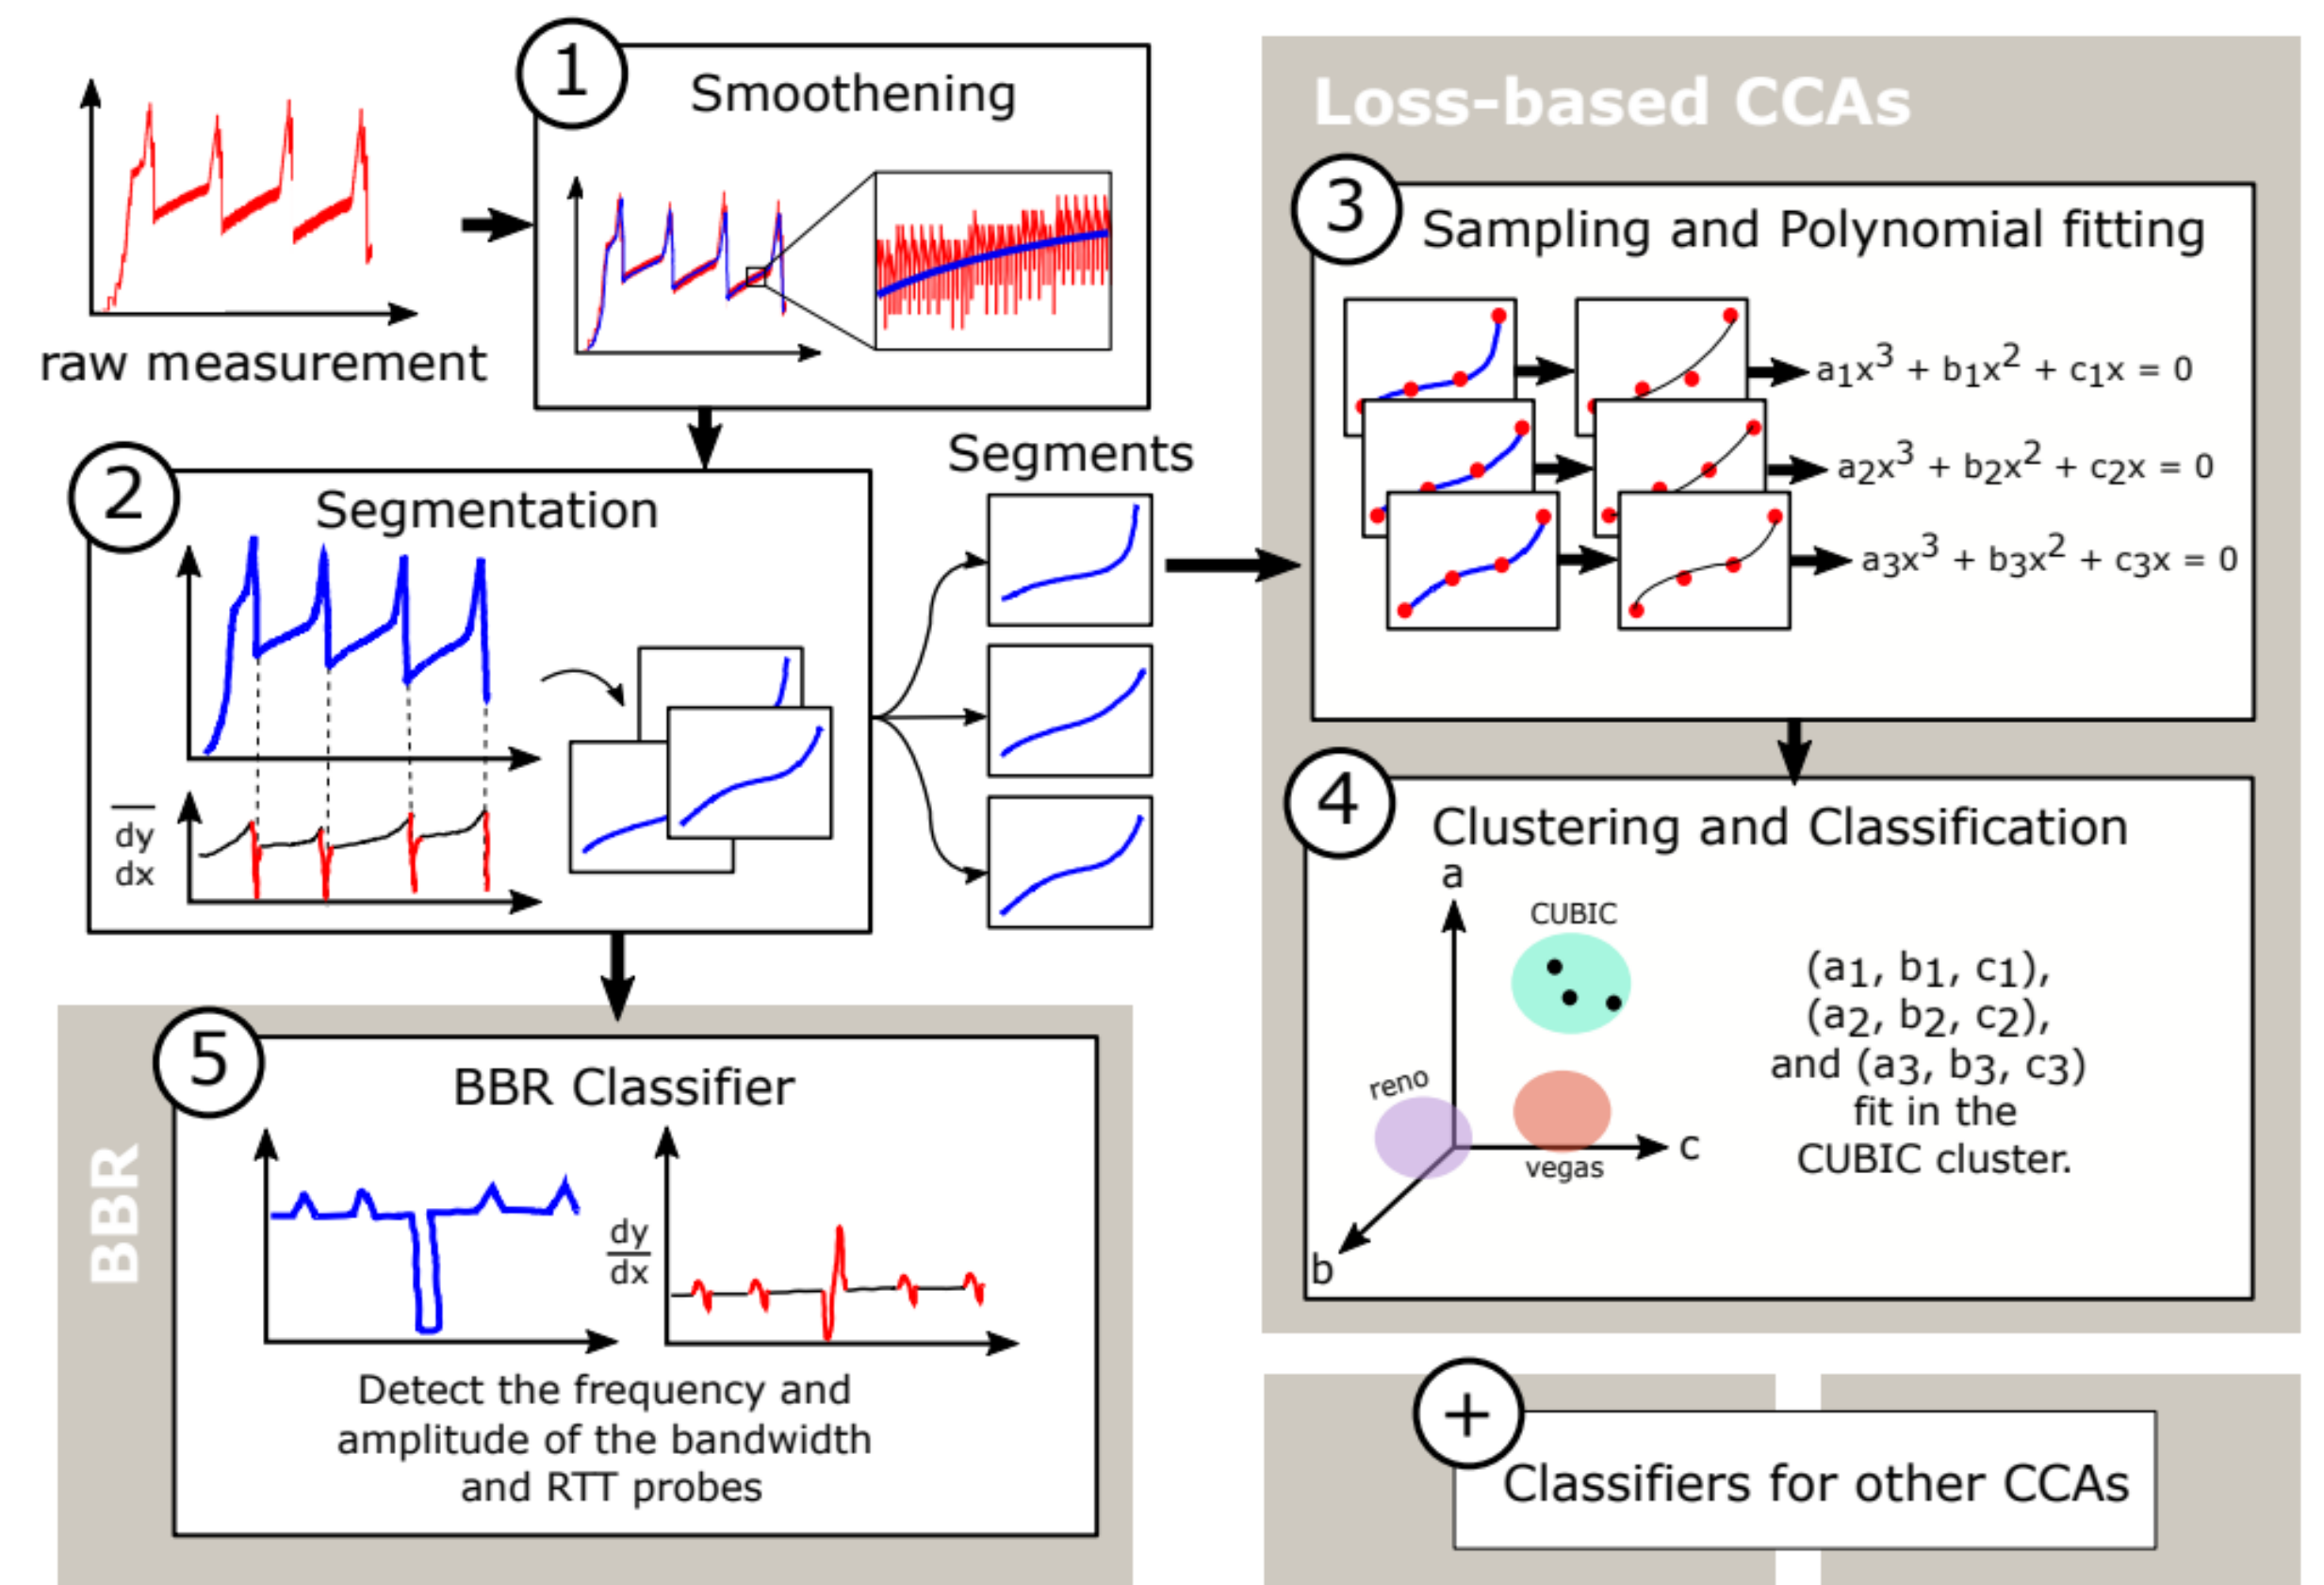
Nebby estimates a server's BiF by artificially delaying the packets after the capture point, giving it visibility into a large fraction of the network pipe. Because traffic on the Internet tends to be bursty, an additional delay 100ms was sufficient to view the entire BiF for a connection.

Nebby's BiF measurements for CCAs



Nebby is able to capture visually distinct BiF traces for most CCAs in the Linux kernel and BBRv2

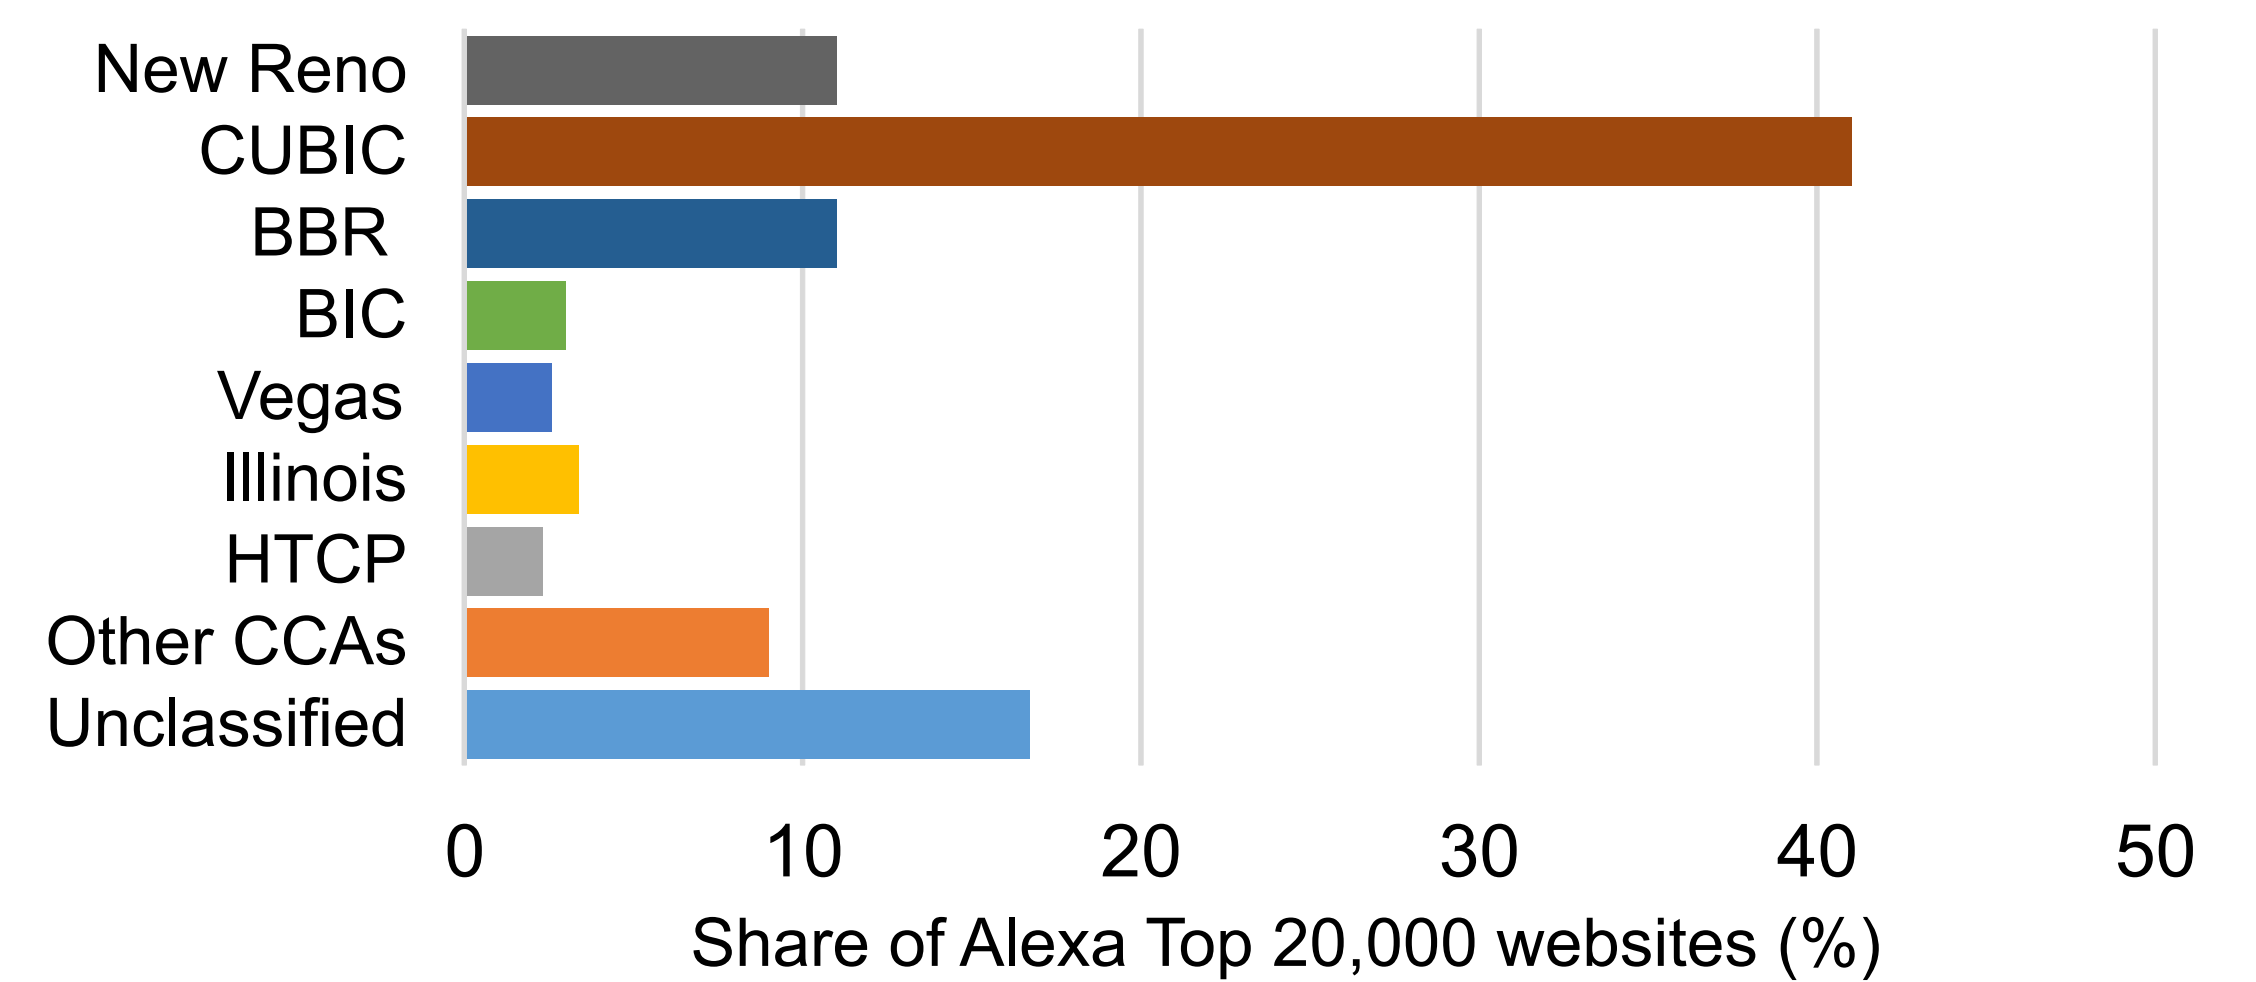
Classifying the BiF measurements



Nebby uses a shape-based classifier that can classify all CCAs in the linux kernel and BBRv2 with an average accuracy of 96% and is easily extensible to other CCAs as well. We have successfully been able to extend the classifier for CCAs like Copa, PCC, and AkamaiCC (a proprietary CCA deployed by Akamai)

Internet Measurement Results


Overall, Nebby's measurements have revealed that the Internet's congestion control landscape is more heterogeneous than it has ever been in the past.



Congestion Control Algorithm (CCA) diversity on the Internet according to a measurement study in 2023 [1]

Overall, we had 5 key findings

- CCA deployment can differ by region**
BBR is more popular in the US and Europe than South America.
- Slow migration to BBRv2**
Most websites that deployed BBRv1 in 2019 have not upgraded to BBRv2, which is supposed to be a fairer alternative to BBRv1



- Unknown CCAs on the Internet**
Nebby caught the deployment of a proprietary CCA by Akamai as well as an early deployment of BBRv3 before it was publicly announced.
- QUIC still has a long way to go**
Only about 8% of the measured websites responded to a QUIC handshake.
- Deployment can differ by asset type**
Some websites like appletv.com and twitch.tv use different CCAs to deliver different assets. Overall, BBR is preferred for streaming video traffic.