

# Researchers find Google's new congestion control algorithm does not treat data fairly

22 October 2019, by Daniel Tkacik



Carnegie Mellon Ph.D. student Ranysha Ware presents her study at the Internet Measurement Conference in Amsterdam. Credit: Carnegie Mellon University

If the Internet had its own superhero, it might be the congestion control algorithm (CCA). CCAs are an essential piece of code Internet giants use to ensure that the Internet doesn't cripple amid a massive data traffic jam. They've been used since the 1980s to slow data transfers when they sense that a network is becoming overloaded.

Like any great superhero, CCAs try to work fairly; when the network is becoming overloaded, they won't prioritize one company's services over another.

However, new research out of Carnegie Mellon shows that a new CCA called BBR, developed by Google, can be unfair competing with other services in overloaded networks. Those findings are being presented this week at the [Internet Measurement Conference](#) in Amsterdam.

"In a given network, our model shows that BBR would take up 40 percent of the bandwidth, leaving the remaining 60 percent to be split between the rest of the parties on the network," says Justine Sherry, a CyLab faculty member and an assistant professor in the Computer Science Department (CSD) at Carnegie Mellon University. This goes against the concept of Internet fairness."

What does this mean for users? Imagine your home uses a 50 megabit per second (Mbps) connection provided by an Internet service provider. Most CCAs try to split the bandwidth evenly when many users want to use the network. If two users are each connected to a different Internet service, the CCA should try to give 25 Mbps to one user and 25 Mbps to the other.

CSD Ph.D. student Ranysha Ware, who leads the research project on Internet fairness, was surprised when she ran experiments modeling network links and saw BBR exhibit very different behavior.

"When only two users are sharing the network, BBR's share is more than fair at 40 percent," Ware says. "But, as we added more users to the network, BBR did not give up any bandwidth as more users joined the [network](#); it kept using 40 percent."

Imagine six people want to share the same 50 Mbps connection. A user connected to a service using BBR would get 20 Mbps of bandwidth, leaving the remaining 30 Mbps to be split between the other five [users](#). Each user would get only 5 Mbps to work with. For video, this difference in bandwidth could be the difference between ultra-high definition video and standard definition.

In 2017, when Google first announced their algorithm, they claimed its design was fundamentally different from most current CCAs.

"People told us that it would be too hard to say anything mathematically provable about BBR

because it works differently from traditional CCAs," says Sherry. But her team found that, indeed, BBR could be compared with other existing CCAs in terms of how it treats data using a mathematical approach based on congestion control windowing.

Is BBR going to harm Internet performance for its competitors?

"Only in the most congested links," Sherry says. "At my house, I have a 1 Gbps connection and it would be very hard to generate the kind of congestion that would make BBR hurt its competitors."

"BBR is a new and evolving algorithm," Sherry says. "We believe that BBR will probably change because of these findings."

Other authors on the study included CSD Department Head Srinivasan Seshan and Nefeli Networks software engineer and CSD alumnus Matthew Mukerjee.

Provided by Carnegie Mellon University

APA citation: Researchers find Google's new congestion control algorithm does not treat data fairly (2019, October 22) retrieved 20 September 2021 from <https://techxplore.com/news/2019-10-google-congestion-algorithm.html>

*This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.*