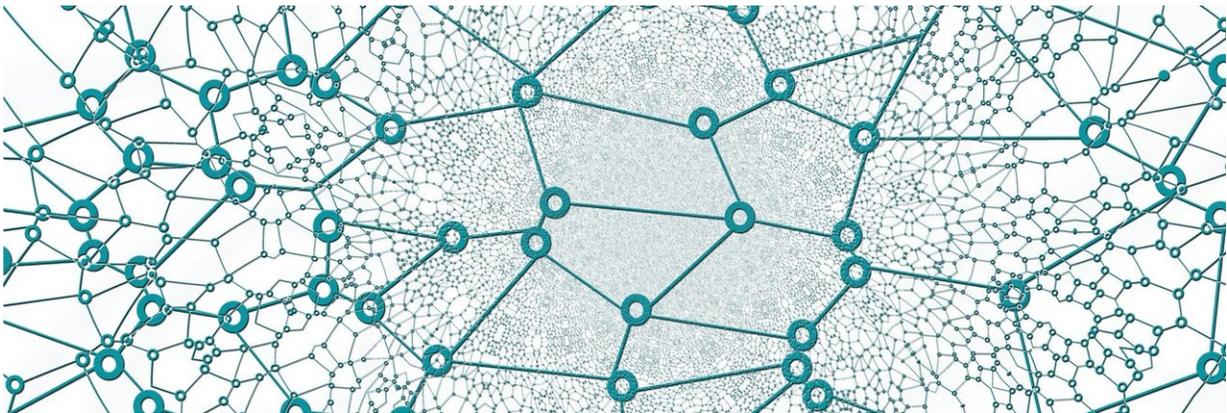


Researchers discover major roadblock in alleviating network congestion

August 4 2022, by Adam Zewe



Credit: CC0 Public Domain

When users want to send data over the internet faster than the network can handle, congestion can occur—the same way traffic congestion snarls the morning commute into a big city.

Computers and devices that transmit data over the internet break the data down into smaller packets and use a special algorithm to decide how fast to send those packets. These [congestion](#) control algorithms seek to fully discover and utilize available network capacity while sharing it fairly with other users who may be sharing the same network. These algorithms try to minimize delay caused by data waiting in queues in the network.

Over the past decade, researchers in industry and academia have developed several algorithms that attempt to achieve high rates while controlling delays. Some of these, such as the BBR algorithm developed by Google, are now widely used by many websites and applications.

But a team of MIT researchers has discovered that these algorithms can be deeply unfair. In a new study, they show there will always be a network scenario where at least one sender receives almost no bandwidth compared to other senders, a problem known as starvation cannot be avoided.

"What is really surprising about this paper and the results is that when you take into account the real-world complexity of network paths and all the things they can do to data packets, it is basically impossible for delay-controlling congestion control algorithms to avoid starvation using current methods," says Mohammad Alizadeh, associate professor of electrical engineering and [computer science](#) (EECS).

While Alizadeh and his co-authors weren't able to find a traditional congestion control algorithm that could avoid starvation, there may be algorithms in a different class that could prevent this problem. Their analysis also suggests that changing how these algorithms work, so that they allow for larger variations in delay, could help prevent starvation in some network situations.

Alizadeh wrote the paper with first author and EECS graduate student Venkat Arun and senior author Hari Balakrishnan, the Fujitsu Professor of Computer Science and Artificial Intelligence. The research will be presented at the ACM Special Interest Group on Data Communications (SIGCOMM) conference.

Controlling congestion

Congestion control is a [fundamental problem](#) in networking that researchers have been trying to tackle since the 1980s.

A user's computer does not know how fast to send data packets over the network because it lacks information, such as the quality of the network connection or how many other senders are using the network. Sending packets too slowly makes poor use of the available bandwidth. But sending them too quickly can overwhelm the network, and in doing so, packets will start to get dropped. These packets must be resent, which leads to longer delays. Delays can also be caused by packets waiting in queues for a long time.

Congestion control algorithms use packet losses and delays as signals to infer congestion and decide how fast to send data. But the internet is complicated, and packets can be delayed and lost for reasons unrelated to network congestion. For instance, data could be held up in a queue along the way and then released with a burst of other packets, or the receiver's acknowledgement might be delayed. The authors call delays that are not caused by congestion "jitter."

Even if a congestion control algorithm measures delay perfectly, it can't tell the difference between delay caused by congestion and delay caused by jitter. Delay caused by jitter is unpredictable and confuses the sender. Because of this ambiguity, users start estimating delay differently, which causes them to send packets at unequal rates. Eventually, this leads to a situation where starvation occurs and someone gets shut out completely, Arun explains.

"We started the project because we lacked a theoretical understanding of congestion control behavior in the presence of jitter. To place it on a firmer theoretical footing, we built a [mathematical model](#) that was simple enough to think about, yet able to capture some of the complexities of the internet. It has been very rewarding to have math tell

us things we didn't know and that have practical relevance," he says.

Studying starvation

The researchers fed their mathematical model to a computer, gave it a series of commonly used congestion control algorithms, and asked the computer to find an algorithm that could avoid starvation, using their model.

"We couldn't do it. We tried every algorithm that we are aware of, and some new ones we made up. Nothing worked. The computer always found a situation where some people get all the bandwidth and at least one person gets basically nothing," Arun says.

The researchers were surprised by this result, especially since these algorithms are widely believed to be reasonably fair. They started suspecting that it may not be possible to avoid starvation, an extreme form of unfairness. This motivated them to define a class of algorithms they call "delay-convergent algorithms" that they proved will always suffer from starvation under their network model. All existing congestion control algorithms that control delay (that the researchers are aware of) are delay-convergent.

The fact that such simple failure modes of these widely used algorithms remained unknown for so long illustrates how difficult it is to understand algorithms through empirical testing alone, Arun adds. It underscores the importance of a solid theoretical foundation.

But all hope is not lost. While all the algorithms they tested failed, there may be other algorithms which are not delay-convergent that might be able to avoid starvation. This suggests that one way to fix the problem might be to design congestion control algorithms that vary the delay range more widely, so the range is larger than any delay that might occur

due to jitter in the [network](#).

"To control delays, algorithms have tried to also bound the variations in delay about a desired equilibrium, but there is nothing wrong in potentially creating greater [delay](#) variation to get better measurements of congestive delays. It is just a new design philosophy you would have to adopt," Balakrishnan adds.

Now, the researchers want to keep pushing to see if they can find or build an [algorithm](#) that will eliminate starvation. They also want to apply this approach of mathematical modeling and computational proofs to other thorny, unsolved problems in networked systems.

"We are increasingly reliant on computer systems for critical things, and we need to put their reliability on a firmer conceptual footing. We've shown the surprising things you can discover when you put in the time to come up with these formal specifications of what the problem actually is," says Alizadeh.

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Researchers discover major roadblock in alleviating network congestion (2022, August 4) retrieved 26 April 2024 from <https://techxplore.com/news/2022-08-major-roadblock-alleviating-network-congestion.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.