

System allocates data center bandwidth more fairly, so no part of a webpage lags behind others

March 28 2017, by Larry Hardesty



At the Usenix Symposium on Networked Systems Design and Implementation, researchers from MIT's Computer Science and Artificial Intelligence Laboratory will present a new system for allocating bandwidth in data center networks. Credit: Massachusetts Institute of Technology

A webpage today is often the sum of many different components. A user's home page on a social-networking site, for instance, might display the latest posts from the users' friends; the associated images, links, and comments; notifications of pending messages and comments on the user's own posts; a list of events; a list of topics currently driving online discussions; a list of games, some of which are flagged to indicate that it's the user's turn; and of course the all-important ads, which the site depends on for revenues.

With increasing frequency, each of those components is handled by a different program running on a different server in the website's data center. That reduces processing time, but it exacerbates another problem: the equitable allocation of network bandwidth among programs.

Many websites aggregate all of a page's components before shipping them to the user. So if just one program has been allocated too little bandwidth on the data center network, the rest of the page—and the user—could be stuck waiting for its component.

At the Usenix Symposium on Networked Systems Design and Implementation this week, researchers from MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) are presenting a new system for allocating bandwidth in data center networks. In tests, the system maintained the same overall data [transmission](#) rate—or network "throughput"—as those currently in use, but it allocated bandwidth much more fairly, completing the download of all of a page's components up to four times as quickly.

"There are easy ways to maximize throughput in a way that divides up the resource very unevenly," says Hari Balakrishnan, the Fujitsu Professor in Electrical Engineering and Computer Science and one of two senior authors on the paper describing the new system. "What we

have shown is a way to very quickly converge to a good allocation."

Joining Balakrishnan on the paper are first author Jonathan Perry, a graduate student in [electrical engineering](#) and computer science, and Devavrat Shah, a professor of electrical engineering and computer science.

Central authority

Most networks regulate data traffic using some version of the [transmission control protocol](#), or TCP. When traffic gets too heavy, some packets of data don't make it to their destinations. With TCP, when a sender realizes its packets aren't getting through, it halves its transmission rate, then slowly ratchets it back up. Given enough time, this procedure will reach an equilibrium point at which [network bandwidth](#) is optimally allocated among senders.

But in a big website's data center, there's often not enough time. "Things change in the network so quickly that this is inadequate," Perry says. "Frequently it takes so long that [the transmission rates] never converge, and it's a lost cause."

TCP gives all responsibility for traffic regulation to the end users because it was designed for the public internet, which links together thousands of smaller, independently owned and operated networks. Centralizing the control of such a sprawling network seemed infeasible, both politically and technically.

But in a [data center](#), which is controlled by a single operator, and with the increases in the speed of both data connections and computer processors in the last decade, centralized regulation has become practical. The CSAIL researchers' system is a centralized system.

The system, dubbed Flowtune, essentially adopts a market-based solution to bandwidth allocation. Operators assign different values to increases in the transmission rates of data sent by different programs. For instance, doubling the transmission rate of the image at the center of a webpage might be worth 50 points, while doubling the transmission rate of analytics data that's reviewed only once or twice a day might be worth only 5 points.

Supply and demand

As in any good market, every link in the network sets a "price" according to "demand"—that is, according to the amount of data that senders collectively want to send over it. For every pair of sending and receiving computers, Flowtune then calculates the transmission rate that maximizes total "profit," or the difference between the value of increased transmission rates—the 50 points for the picture versus the 5 for the analytics data—and the price of the requisite bandwidth across all the intervening links.

The maximization of profit, however, changes demand across the links, so Flowtune continually recalculates prices and on that basis recalculates maximum profits, assigning the resulting [transmission rates](#) to the servers sending data across the [network](#).

The paper also describes a new procedure that the researchers developed for allocating Flowtune's computations across cores in a multicore computer, to boost efficiency. In experiments, the researchers compared Flowtune to a widely used variation on TCP, using data from real data centers. Depending on the data set, Flowtune completed the slowest 1 percent of data requests nine to 11 times as rapidly as the existing system.

More information: Paper: "Flowtune: Flowlet control for datacenter

networks" [dspace.mit.edu/bitstream/handle... SAIL-TR-2016-011.pdf](https://dspace.mit.edu/bitstream/handle/1890/10600/1/SAIL-TR-2016-011.pdf)

*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: System allocates data center bandwidth more fairly, so no part of a webpage lags behind others (2017, March 28) retrieved 10 April 2024 from
<https://techxplore.com/news/2017-03-allocates-center-bandwidth-webpage-lags.html>

<p>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.</p>
--