

Programmable network routers provide more flexible traffic management without sacrificing speed

23 August 2016, by Larry Hardesty



"You need to have the ability for researchers and engineers to try out thousands of ideas," says Hari Balakrishnan, the Fujitsu Professor in Electrical Engineering and Computer Science at MIT. "With this platform, you become constrained not by hardware or technological limitations, but by your creativity. You can innovate much more rapidly." Credit: Massachusetts Institute of Technology

Like all data networks, the networks that connect servers in giant server farms, or servers and workstations in large organizations, are prone to congestion. When network traffic is heavy, packets of data can get backed up at network routers or dropped altogether.

Also like all [data networks](#), big private networks have control algorithms for managing [network traffic](#) during periods of congestion. But because the routers that direct traffic in a server farm need to be superfast, the [control algorithms](#) are hardwired into the routers' circuitry. That means that if someone develops a better algorithm, network operators have to wait for a new generation of hardware before they can take

advantage of it.

Researchers at MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) and five other organizations hope to change that, with routers that are programmable but can still keep up with the blazing speeds of modern data networks. The researchers outline their system in a pair of papers being presented at the annual conference of the Association for Computing Machinery's Special Interest Group on Data Communication.

"This work shows that you can achieve many flexible goals for managing traffic, while retaining the high performance of traditional routers," says Hari Balakrishnan, the Fujitsu Professor in Electrical Engineering and Computer Science at MIT. "Previously, programmability was achievable, but nobody would use it in production, because it was a factor of 10 or even 100 slower."

"You need to have the ability for researchers and engineers to try out thousands of ideas," he adds. "With this platform, you become constrained not by hardware or technological limitations, but by your creativity. You can innovate much more rapidly."

The first author on both papers is Anirudh Sivaraman, an MIT graduate student in [electrical engineering](#) and [computer science](#), advised by both Balakrishnan and Mohammad Alizadeh, the TIBCO Career Development Assistant Professor in Electrical Engineering and Computer Science at MIT, who are coauthors on both papers. They're joined by colleagues from MIT, the University of Washington, Barefoot Networks, Microsoft Research, Stanford University, and Cisco Systems.

Different strokes

Traffic management can get tricky because of the

different types of data traveling over a network, and the different types of performance guarantees offered by different services. With Internet phone calls, for instance, delays are a nuisance, but the occasional dropped packet—which might translate to a missing word in a sentence—could be tolerable. With a large data file, on the other hand, a slight delay could be tolerable, but missing data isn't.

Similarly, a network may guarantee equal bandwidth distribution among its users. Every router in a data network has its own memory bank, called a buffer, where it can queue up packets. If one user has filled a router's buffer with packets from a single high-definition video, and another is trying to download a comparatively tiny text document, the network might want to bump some of the video packets in favor of the text, to help guarantee both users a minimum data rate.

A router might also want to modify a packet to convey information about network conditions, such as whether the packet encountered congestion, where, and for how long; it might even want to suggest new transmission rates for senders.

Computer scientists have proposed hundreds of traffic management schemes involving complex rules for determining which packets to admit to a router and which to drop, in what order to queue the packets, and what additional information to add to them—all under a variety of different circumstances. And while in simulations many of these schemes promise improved network performance, few of them have ever been deployed, because of hardware constraints in routers.

The MIT researchers and their colleagues set themselves the goal of finding a set of simple computing elements that could be arranged to implement diverse traffic management schemes, without compromising the operating speeds of today's best routers and without taking up too much space on-chip.

To test their designs, they built a compiler—a program that converts high-level program instructions into low-level hardware instructions—which they used to compile seven

experimental traffic-management algorithms onto their proposed [circuit elements](#). If an algorithm wouldn't compile, or if it required an impractically large number of circuits, they would add new, more sophisticated circuit elements to their palette.

Assessments

In one of the two new papers, the researchers provide specifications for seven circuit types, each of which is slightly more complex than the last. Some simple [traffic management](#) algorithms require only the simplest circuit type, while others require more complex types. But even a bank of the most complex circuits would take up only 4 percent of the area of a router chip; a bank of the least complex types would take up only 0.16 percent.

Beyond the seven algorithms they used to design their circuit elements, the researchers ran several other algorithms through their compiler and found that they compiled to some combination of their simple circuit elements.

"We believe that they'll generalize to many more," says Sivaraman. "For instance, one of the circuits allows a programmer to track a running sum—something that is employed by many algorithms."

In the second paper, they describe the design of their scheduler, the circuit element that orders packets in the router's queue and extracts them for forwarding. In addition to queuing packets according to priority, the scheduler can also stamp them with particular transmission times and forward them accordingly. Sometimes, for instance, it could be useful for a router to slow down its transmission rate, in order to prevent bottlenecks elsewhere in the network, or to help ensure equitable bandwidth distribution.

Finally, the researchers drew up specifications for their circuits in Verilog, the language electrical engineers typically use to design commercial chips. Verilog's built-in analytic tools verified that a router using the researchers' circuits would be fast enough to support the packet rates common in today's high-speed networks, forwarding a packet of data every nanosecond.

"There are a lot of problems in computer networking we've never been able to solve at the speed that traffic actually flows through the network, because there wasn't support directly in the network devices to analyze the traffic or act on the traffic as it arrives," says Jennifer Rexford, a professor of computer science at Princeton University. "What's exciting about both of these works is that they really point to next-generation switch hardware that will be much, much more capable—and more importantly, more programmable, so that we can really change how the network functions without having to replace the equipment inside the network."

"At the edge of the network, applications change all the time," she adds. "Who knew Pokémon Go was going to happen? It's incredibly frustrating when applications' needs evolve years and years more quickly than the equipment's ability to support it. Getting the time scale of innovation inside the [network](#) to be closer to the time scale of innovation in applications is, I think, quite important."

More information: Packet Transactions: High-Level Programming for Line-Rate Switches.

web.mit.edu/domino/

Programmable Packet Scheduling at Line Rate.

web.mit.edu/pifo/

Provided by Massachusetts Institute of Technology

APA citation: Programmable network routers provide more flexible traffic management without sacrificing speed (2016, August 23) retrieved 21 September 2020 from

<https://techxplore.com/news/2016-08-programmable-network-routers-flexible-traffic.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.