

New approach speeds up testing of traffic management solutions for data center networks

March 22 2017, by Larry Hardesty



MIT researchers have developed a system for testing new traffic management protocols. The system requires no alteration to network hardware but still works at realistic speeds — 20 times as fast as networks of software-controlled routers. Credit: Massachusetts Institute of Technology

The transmission control protocol, or TCP, which manages traffic on the Internet, was first proposed in 1974. Some version of TCP still regulates data transfer in most major data centers, the huge warehouses of servers maintained by popular websites.

That's not because TCP is perfect or because computer scientists have had trouble coming up with possible alternatives; it's because those alternatives are too hard to test. The routers in data center networks have their traffic management protocols hardwired into them. Testing a new protocol means replacing the existing [network](#) hardware with either reconfigurable chips, which are labor-intensive to program, or software-controlled routers, which are so slow that they render large-scale testing impractical.

At the Usenix Symposium on Networked Systems Design and Implementation later this month, researchers from MIT's Computer Science and Artificial Intelligence Laboratory will present a system for testing new traffic management protocols that requires no alteration to network hardware but still works at realistic speeds—20 times as fast as networks of software-controlled routers.

The system maintains a compact, efficient computational model of a network running the new protocol, with virtual data packets that bounce around among virtual routers. On the basis of the model, it schedules transmissions on the real network to produce the same traffic patterns. Researchers could thus run real web applications on the network [servers](#) and get an accurate sense of how the new protocol would affect their performance.

"The way it works is, when an endpoint wants to send a [data] packet, it first sends a request to this centralized emulator," says Amy Ousterhout, a graduate student in electrical engineering and computer science (EECS) and first author on the new paper. "The emulator emulates in

software the scheme that you want to experiment with in your network. Then it tells the endpoint when to send the packet so that it will arrive at its destination as though it had traversed a network running the programmed scheme."

Ousterhout is joined on the paper by her advisor, Hari Balakrishnan, the Fujitsu Professor in Electrical Engineering and Computer Science; Jonathan Perry, a graduate student in EECS; and Petr Lapukhov of Facebook.

Traffic control

Each packet of data sent over a computer network has two parts: the header and the payload. The payload contains the data the recipient is interested in—image data, audio data, text data, and so on. The header contains the sender's address, the recipient's address, and other information that routers and end users can use to manage transmissions.

When multiple packets reach a [router](#) at the same time, they're put into a queue and processed sequentially. With TCP, if the queue gets too long, subsequent packets are simply dropped; they never reach their recipients. When a sending computer realizes that its packets are being dropped, it cuts its transmission rate in half, then slowly ratchets it back up.

A better protocol might enable a router to flip bits in packet headers to let end users know that the network is congested, so they can throttle back transmission rates before packets get dropped. Or it might assign different types of packets different priorities, and keep the transmission rates up as long as the high-priority traffic is still getting through. These are the types of strategies that computer scientists are interested in testing out on real networks.

Speedy simulation

With the MIT researchers' new system, called Flexplane, the emulator, which models a network running the new protocol, uses only packets' header data, reducing its computational burden. In fact, it doesn't necessarily use all the header data—just the fields that are relevant to implementing the new protocol.

When a server on the real network wants to transmit data, it sends a request to the emulator, which sends a dummy packet over a [virtual network](#) governed by the new protocol. When the dummy packet reaches its destination, the emulator tells the real server that it can go ahead and send its real packet.

If, while passing through the virtual network, a dummy packet has some of its header bits flipped, the real server flips the corresponding bits in the real packet before sending it. If a clogged router on the virtual network drops a dummy packet, the corresponding real packet is never sent. And if, on the virtual network, a higher-priority dummy packet reaches a router after a lower-priority packet but jumps ahead of it in the queue, then on the real network, the higher-priority packet is sent first.

The servers on the network thus see the same packets in the same sequence that they would if the real routers were running the new [protocol](#). There's a slight delay between the first request issued by the first server and the first transmission instruction issued by the emulator. But thereafter, the servers issue packets at normal network speeds.

The ability to use real servers running real web applications offers a significant advantage over another popular technique for testing new network management schemes: software simulation, which generally uses statistical patterns to characterize the applications' behavior in a computationally efficient manner.

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: New approach speeds up testing of traffic management solutions for data center networks (2017, March 22) retrieved 10 August 2024 from <https://techxplore.com/news/2017-03-approach-traffic-solutions-center-networks.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.