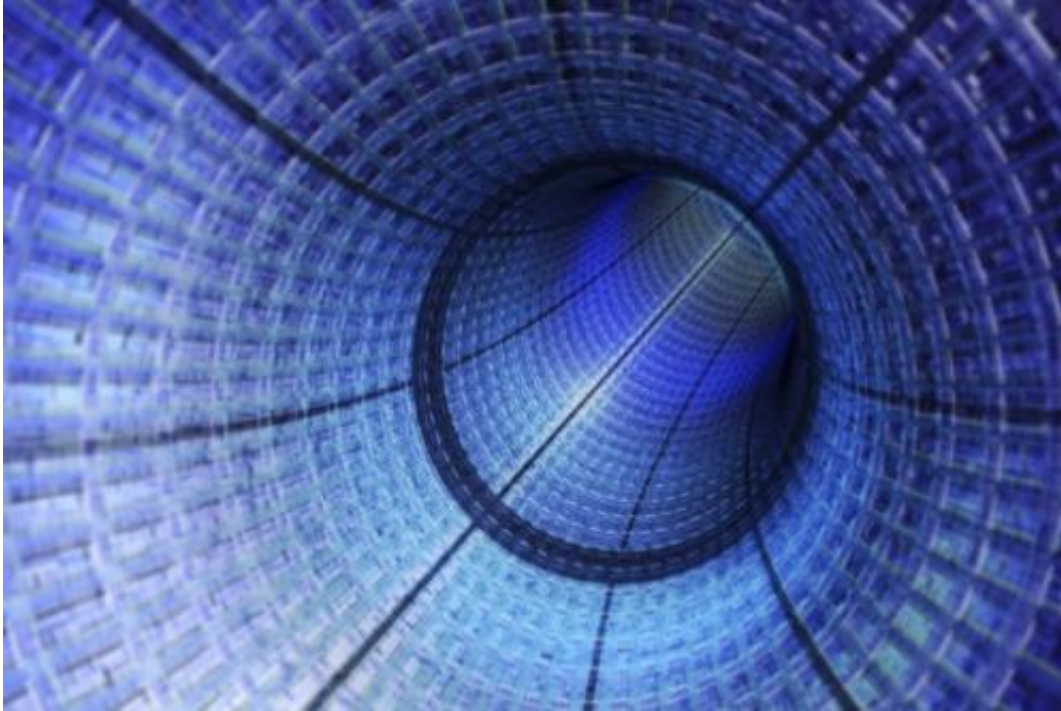# How to fix the Internet's plumbing problem

April 25 2014, by Angela Herring

Associate professor Edmund Yeh is developing an algorithm that will help solve the Internet's plumbing problem. Credit: Thinkstock.

(Phys.org) —Twenty-five years ago, an engineer at CERN, the European Organization for Nuclear Research in Geneva Switzerland, had an idea that would change the world. Tim Berners Lee sent a memo to his colleagues at the world's largest particle physics laboratory proposing a way for all of its computers—and eventually all the computers in the world—to talk to each other. This memo marks the birth of the World Wide Web.

Since then, the Internet has become increasingly essential to human activity—from reading the news and buying stocks to communicating and researching flu symptoms. But it still has some problems—namely, its plumbing, according to Edmund Yeh, an associate professor of electrical and computer engineering at Northeastern University.

Yeh is a principal investigator on a multi-university $7.9 million research project to unclog the Internet's plumbing and fix this problem. The Internet's basic architecture is highly inefficient, he said, because new connections to each piece of content must be created whenever a user wants access to the material.

"If you've got a lot of demand for a particular data type, it's like water building up," Yeh explained. That water, he said, can be managed in two ways—one is by getting it to its destination drain (i.e., the data server), the other is to drill a new drain somewhere along its journey (i.e., a caching point that temporarily stores the data).

Yeh's research team is building what it calls the Named Data Network—one of four projects funded by the federal government's Future Internet Program. NDN maintains a network of caching points, or "drains," that are strategically located around the core of the Internet as well as its periphery.

Content distribution networks such as Akamai are taking a similar approach, Yeh said, but their operations work only at the edge of the Internet and provide service only to large content providers such as CNN. "This," Yeh said, "is due to the inability of the current Internet architecture to handle caching."

With NDN, user requests would be directed to caching points located closer to them in the Internet, regardless of the type of content they're trying to access.

Still, a challenge remains: where to place the drains and how to best direct the "water" toward them. In collaboration with researchers at the California Institute of Technology, Yeh's lab at Northeastern recently [released a paper on *arXiv*](#) in which the team presents an algorithm called the VIP Algorithm. It simultaneously performs those two processes—both the drilling of drains and the directing of water. The processes work in a coordinated manner, Yeh said, with each relying on the other and both dynamically adapting to the prevailing demand for content.

Incidentally, the global network of high-energy physics researchers attempting to access [particle physics](#) data from the Large Hadron Collider at CERN—where the World Wide Web was born all those years ago—also has a similar plumbing problem. But physics researchers based at the California Institute of Technology believe Yeh's algorithms could help them unclog their network pipes. So the two groups have teamed up to test it out.

Indeed, Yeh's approach is a promising candidate for any scientific network in which a global community of users is requesting access to a vast body of data.

"We could make an impact first in these scientific big-data applications," he said. "That's a good testing ground for the broader Internet."

  More information: [arxiv.org/abs/1310.5569](#)

Provided by Northeastern University

Citation: How to fix the Internet's plumbing problem (2014, April 25) retrieved 11 May 2024