UC San Diego Researchers’ New Algorithm Significantly Boosts Routing Efficiency of Networks

August 18, 2008

Paul K. Mueller

A time-and-money-saving question shared by commuters in their cars and networks sharing ever-changing Internet resources is: "What's the best way to get from here to there?"

A new algorithm developed by computer scientists at the University of California, San Diego helps answer that question, at least for computer networks, and it promises to significantly boost the efficiency of network routing.

Called XL, for approximate link state, the algorithm increases network routing efficiency by suppressing updates from parts of the system - updates which force connected networks to continuously re-calculate the paths they use in the great matrix of the Internet.

"Routing in a static network is trivial," say the authors in their paper, which will be presented at this week's ACM SIGCOMM conference. "But most real networks are dynamic - network links go up and down - and thus some nodes need to recalculate their routes in response."

The traditional approach, said Stefan Savage, professor of computer science at UC San Diego, "is to tell everyone; flood the topology change throughout the network and have each node re-compute its table of best routes - but that requirement to universally communicate, and to act on each change, is a big problem."

What the team did with their new routing algorithm, according to Savage's student Kirill Levchenko, was to reduce the "communication overhead" of route computation - by an order of magnitude.

"Being able to adapt to hardware failures is one of the fundamental characteristics of the Internet," Levchenko said. "Our routing algorithm reduces the overhead of route re-computation after a network change, making it possible to support larger networks. The benefits are especially significant when networks are made up of low-power devices of slow links."

The real technical innovation of their work, said another of the authors, Geoffrey M. Voelker, "is in how information about changes in the network is propagated. The XL routing algorithm propagates only some updates, reducing the number of updates sent through the network."

They meet the "central challenge" of determining which updates are important and which can be suppressed by using three rules for update propagation, said team member Ramamohan Paturi. "The rules ensure that selected routes are nearly as good as if complete information about the network were available," he said, "but at a fraction of the overhead required for maintaining such a state of perfect knowledge."

The computer scientists also believe that there are "significant opportunities" to improve the efficiency of link-state routing even further. They look forward to discovering an algorithm that improves on their Approximate Link work with similar boosts in efficiency.

Grants from the National Science Foundation helped support the team's research.

Comment: Stefan Savage, 858-534-4616 Media Contact: Paul K. Mueller, 858-534-8564