Creating a high-speed internet lane for emergency situations

August 1 2017, by Nirmala Shenoy, Erik Golen And Jennifer Schneider

In an emergency, responders' telecommunications could get delayed by overloaded networks. Credit: City of Hampton, Virginia

During large disasters, like hurricanes, wildfires and terrorist attacks, people want emergency responders to arrive quickly and help people

deal with the crisis. In order to do their best, police, medics, firefighters and those who manage them need lots of information: Who is located where, needing what help? And what equipment and which rescuers are available to intervene? With all of the technology we have, it might seem that gathering and sharing lots of information would be pretty simple. But communicating through a disaster is much more challenging than it appears.

The event itself can make communications worse, damaging networks and phone systems or cutting electricity to an area. And regular people often add to the problem as they overload mobile networks with calls, texts and other electronic messages checking on loved ones or seeking help.

As researchers about digital networks and emergency communications, we are developing a faster and more reliable way to send and receive large amounts of data through the internet in times of crisis. Working with actual responders and emergency managers, we have created a method for giving urgent information priority over other internet traffic, effectively creating a high-speed lane on the internet for use in emergencies. While a national emergency responder network initiative called FirstNet is beginning to get going, it requires building an all-new wireless network just for emergency services to use. By contrast, our system uses existing internet connections, while giving priority to rescue workers' data.

Connecting networks

At the moment, it's reasonably common for communication networks to become overloaded when disaster strikes. When lots of people try to make cellphone calls or use mobile data, the networks get too busy for calls to connect and messages to go through.

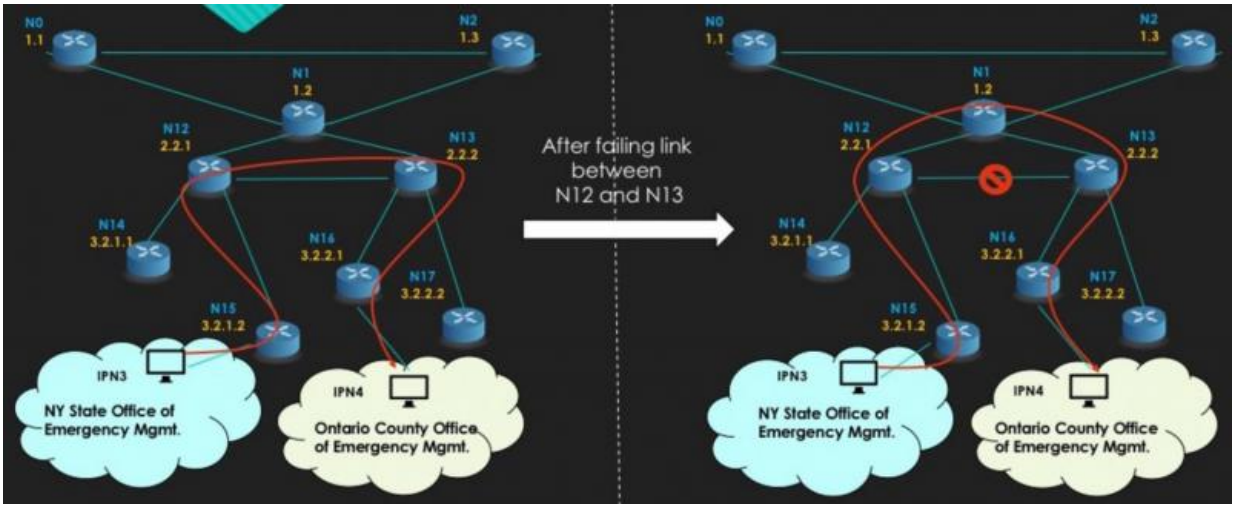

When a link fails, the network system must find a new connection between two communicating devices. Credit: Rochester Institute of Technology, CC BY-ND

The problem is that standard methods for routing traffic through the internet aren't always able to handle all those connections at one time. In technical terms, the internet is a collection of more than 54,000 smaller networks. Some of the networks that make up the internet are quite large, like those belonging to major internet service providers or large corporations, but many of them are fairly small. No matter their size, each of these networks has equipment that lets it route traffic to each of the others.

Computer networks don't all connect directly to each other. And their digital addresses don't help much – we humans assume 12 Main Street and 14 Main Street are next door, but computers with similar numeric addresses may not be physical neighbors to each other.

As a result, the router connecting each of these 54,000 networks to the rest of the internet must keep a list of every one of its counterparts, and

the most efficient way to reach each of them. This is like needing a list of written directions for every place in the world you might want to go.

This system, governed by the rules set out in the "border gateway protocol," works well most of the time. But when it fails, there can be long delays in communications. In fact, on average, 150 seconds (two and a half minutes) can go by before a failure is identified. In that time, the data just wait in an information traffic jam, not moving. Online, milliseconds matter – hundreds of seconds are effectively an eternity.

When one router detects a network failure, it has to let all the others know what's happened, and how to reroute their traffic. This is like having just one traffic cop try to coordinate rush hour around a major bottleneck. The process takes at least several minutes, and sometimes several hours. Until then, data in transit can be delayed or lost entirely. In an emergency, that could mean the difference between life and death.

Developing the emergency protocol

Working with students from Rochester Institute of Technology's Golisano College of Computing and Information Sciences, we have created a new traffic control system tailored specifically to emergency response networks. It runs without affecting other protocols on the internet. We call it the multi-node label routing protocol.

Rather than requiring every router to keep track of the best directions to every other one, we divide possible routes for internet traffic into hierarchies. These mirror existing emergency response plans: An individual responder sends information to a local commander, who combines several responders' data and passes the data on to regional managers, who assemble a wider picture they pass on to state or federal response coordinators.

Our routing plan makes direct network connections mirror this real-world emergency response hierarchy. When routers are allowed to connect only with their immediate neighbors in the hierarchy, they can notice when links fail and reroute traffic much more quickly.

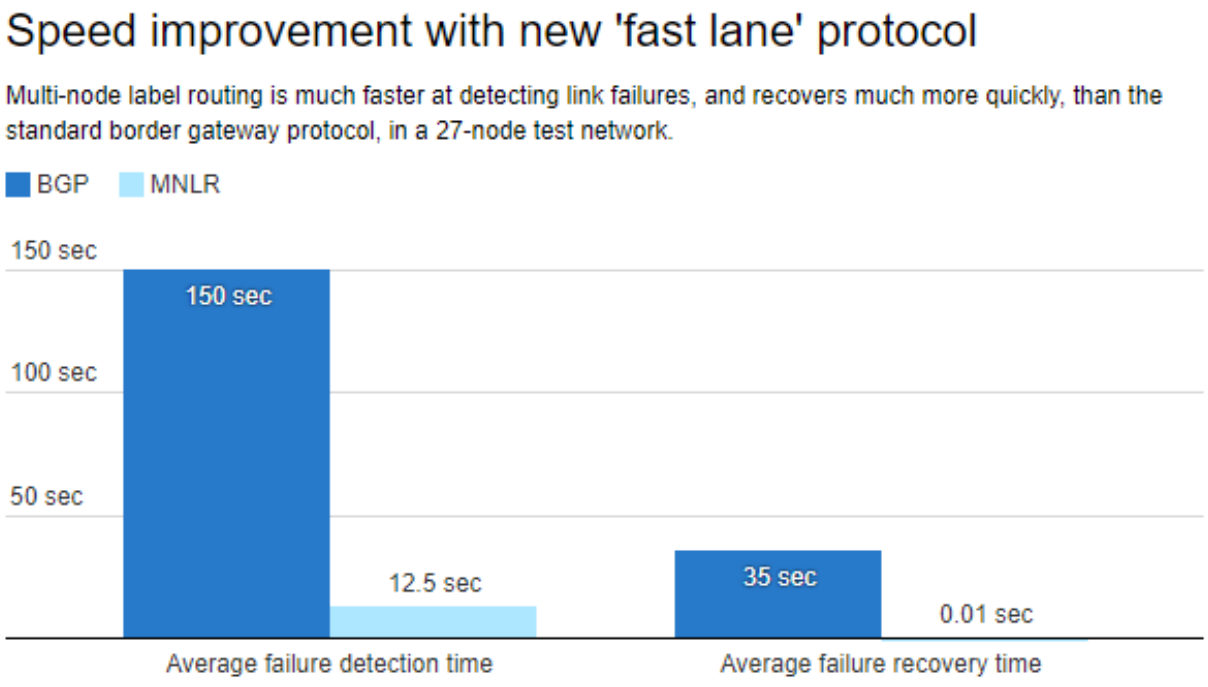

The Conversation, CC-BY-ND

Source: Rochester Institute of Technology

Testing in the real world

Our system is designed to operate over the same internet as everyone else, and without affecting other traffic. We tested our system on the

National Science Foundation's Global Environment for Network Innovations, a collaborative effort among many universities around the U.S. that allows researchers to develop networking protocols and systems using real computers and networking equipment located across the country. In our case, we connected 27 computers together for our tests, devised by RIT environmental, health and safety students, many of whom are volunteer emergency responders.

Our test – which we did in front of real emergency commanders and personnel – compared our system to the standard border gateway protocol. When we broke links in the 27-node network, multi-node label routing communications resumed within 12.5 seconds, which is 12 times faster than the regular border gateway protocol's recovery speed. We can shorten that delay even more by changing settings in our protocol's configuration.

Our system can easily be installed across a much wider area than just 27 test machines, specifically because of how it simplifies the paths information takes between routers. This means incident commanders and managers get information more quickly, and are better able to allocate responders and equipment to meet needs as they develop. In this way, our work supports the efforts of those who support us in our hour of need.

This article was originally published on The Conversation. Read the original article.

Provided by The Conversation