

Computing's invisible challenge

April 8 2014, by Angela Herring



Assistant professor of electrical and computer engineering Ningfan Mi was recently named one of 42 recipients of the Air Force Office of Scientific Research Young Investigator Award. Photo by Brooks Canaday

To us, it may not seem like a big deal: CNN's website is taking too long to load. The day's most popular YouTube video won't stop buffering. "Twitter is over capacity." While these little hiccups in usability may frustrate end users, they merely scratch the surface of the enormous technical challenge that's confronting the backend.

Northeastern University assistant professor of electrical and computer engineering Ningfang Mi recently learned she was one of 42 early-career researchers to win a Young Investigator Award from the Air Force Office of Scientific Research. They will receive the grants over a three-year period.

She plans to use award to figure out a better way to manage the vast amount of information sharing that takes place online—and push that massive [technical challenge](#) even further into the background for end users.

These days most of the data we request online is stored in the so-called "cloud"—a series of virtual computers distributed on physical servers around the world. For instance, Google has 12 data centers across four continents. The 20,000 emails sitting in my Gmail inbox aren't actually stored on my computer—they're stored in Google's cloud, which exists on all those remote servers. Every time I look at one of my emails, I am requesting access to it from one of those servers.

Now consider YouTube. Its billions of hours of video aren't all sitting on the same physical server; rather, they are stored remotely in the cloud. In this case, I am just one of millions of users requesting the same video in a given moment. And that, Mi explained, is where things get challenging.

Her research is focused on modeling performance in different scenarios and figuring out the best ways to manage resources based on the outcomes of those models. This will give her a sense of the workloads and number of traffic requests that remote servers are likely to have to handle.

"Based on this kind of information," she said, "how can I find the best configuration for the platform in order to provide the highest quality of service?"

There are two options: She can either move information around on a single server or move information between servers. The best choice will depend on the situation at hand.

"Before predictions were based more on average load or traffic, but now we know that in reality the workload changes," Mi said. "The term I use here is 'burstiness' or 'spikes.'"

Indeed, it all depends on the burstiness of human behavior. Some online phenomena are predictable, Mi said. For instance, you're likely to see a burst in email activity on the East Coast every weekday at around 9 a.m. EST. Similarly, the Internet is likely to be all-a-flurry across a range of websites on election night as people world over discuss the race on Twitter, stream acceptance speeches on NBC, and read about the results in The New York Times.

But what about when a celebrity unexpectedly passes away or makes a comment that goes viral? Or when a boy in a balloon suddenly becomes one of the biggest news stories on the Internet? No one can predict events like that, so no amount of resource management preparation could ready YouTube for the associated activity spikes.

Mi, for her part, is developing models that will help detect those bursts with more immediacy—and in some cases even predict them a couple hours in advance. So while we may not know when the next media hoax will drive traffic from millions of curious viewers, at least our computers will be able to handle it better.

Provided by Northeastern University

Citation: Computing's invisible challenge (2014, April 8) retrieved 27 April 2024 from <https://phys.org/news/2014-04-invisible.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.