


Using Wall Street secrets to reduce the cost of cloud infrastructure

August 16 2019



Credit: CC0 Public Domain

Stock market investors often rely on financial risk theories that help them maximize returns while minimizing financial loss due to market fluctuations. These theories help investors maintain a balanced portfolio to ensure they'll never lose more money than they're willing to part with at any given time.

Inspired by those theories, MIT researchers in collaboration with Microsoft have developed a "risk-aware" [mathematical model](#) that could improve the performance of cloud-computing networks across the globe. Notably, cloud infrastructure is extremely expensive and consumes a lot of the world's energy.

Their model takes into account failure probabilities of links between data centers worldwide—akin to predicting the volatility of stocks. Then, it runs an optimization engine to allocate traffic through optimal paths to minimize loss, while maximizing overall usage of the network.

The model could help major cloud-service providers—such as Microsoft, Amazon, and Google—better utilize their infrastructure. The conventional approach is to keep links idle to handle unexpected traffic shifts resulting from link failures, which is a waste of energy, bandwidth, and other resources. The new model, called TeaVar, on the other hand, guarantees that for a target percentage of time—say, 99.9 percent—the network can handle all [data traffic](#), so there is no need to keep any links idle. During that 0.01 percent of time, the model also keeps the data dropped as low as possible.

In experiments based on real-world data, the model supported three times the traffic throughput as traditional traffic-engineering methods, while maintaining the same high level of network availability. A paper describing the model and results will be presented at the ACM SIGCOMM conference this week.

Better network utilization can save service providers millions of dollars, but benefits will "trickle down" to consumers, says co-author Manya Ghobadi, the TIBCO Career Development Assistant Professor in the MIT Department of Electrical Engineering and Computer Science and a researcher at the Computer Science and Artificial Intelligence Laboratory (CSAIL).

"Having greater utilized infrastructure isn't just good for cloud services—it's also better for the world," Ghobadi says. "Companies don't have to purchase as much infrastructure to sell services to customers. Plus, being able to efficiently utilize datacenter resources can save enormous amounts of energy consumption by the cloud infrastructure. So, there are benefits both for the users and the environment at the same time."

Joining Ghobadi on the paper are her students Jeremy Bogle and Nikhil Bhatia, both of CSAIL; Ishai Menache and Nikolaj Bjorner of Microsoft Research; and Asaf Valadarsky and Michael Schapira of Hebrew University.

On the money

Cloud service providers use networks of fiber optical cables running underground, connecting data centers in different cities. To route traffic, the providers rely on "traffic engineering" (TE) software that optimally allocates data bandwidth—amount of data that can be transferred at one time—through all network paths.

The goal is to ensure maximum availability to users around the world. But that's challenging when some links can fail unexpectedly, due to drops in optical signal quality resulting from outages or lines cut during construction, among other factors. To stay robust to failure, providers keep many links at very low utilization, lying in wait to absorb full data loads from downed links.

Thus, it's a tricky tradeoff between network availability and utilization, which would enable higher data throughputs. And that's where traditional TE methods fail, the researchers say. They find optimal paths based on various factors, but never quantify the reliability of links. "They don't say, 'This link has a higher probability of being up and

running, so that means you should be sending more traffic here," Bogle says. "Most links in a network are operating at low utilization and aren't sending as much traffic as they could be sending."

The researchers instead designed a TE model that adapts core mathematics from "conditional value at risk," a risk-assessment measure that quantifies the average loss of money. With investing in stocks, if you have a one-day 99 percent conditional value at risk of \$50, your expected loss of the worst-case 1 percent scenario on that day is \$50. But 99 percent of the time, you'll do much better. That measure is used for investing in the stock market—which is notoriously difficult to predict.

"But the math is actually a better fit for our cloud infrastructure setting," Ghobadi says. "Mostly, link failures are due to the age of equipment, so the probabilities of failure don't change much over time. That means our probabilities are more reliable, compared to the stock market."

Risk-aware model

In networks, data bandwidth shares are analogous to invested "money," and the network equipment with different probabilities of failure are the "stocks" and their uncertainty of changing values. Using the underlying formulas, the researchers designed a "risk-aware" model that, like its financial counterpart, guarantees data will reach its destination 99.9 percent of time, but keeps traffic loss at minimum during 0.1 percent worst-case failure scenarios. That allows cloud providers to tune the availability-utilization tradeoff.

The researchers statistically mapped three years' worth of network signal strength from Microsoft's networks that connects its [data centers](#) to a probability distribution of link failures. The input is the network topology in a graph, with source-destination flows of data connected through lines (links) and nodes (cities), with each link assigned a

bandwidth.

Failure probabilities were obtained by checking the signal quality of every link every 15 minutes. If the signal quality ever dipped below a receiving threshold, they considered that a link failure. Anything above meant the link was up and running. From that, the model generated an average time that each link was up or down, and calculated a failure probability—or "risk"—for each link at each 15-minute time window. From those data, it was able to predict when risky links would fail at any given window of time.

The researchers tested the model against other TE software on simulated traffic sent through networks from Google, IBM, ATT, and others that spread across the world. The researchers created various failure scenarios based on their probability of occurrence. Then, they sent simulated and real-world data demands through the network and cued their models to start allocating bandwidth.

The researchers' model kept reliable links working to near full capacity, while steering data clear of riskier links. Over traditional approaches, their [model](#) ran three times as much data through the [network](#), while still ensuring all data got to its destination. The code is freely available on GitHub.

More information: Jeremy Bogle et al, TEAVAR, *Proceedings of the ACM Special Interest Group on Data Communication - SIGCOMM '19* (2019). [DOI: 10.1145/3341302.3342069](https://doi.org/10.1145/3341302.3342069)

Provided by Massachusetts Institute of Technology

Citation: Using Wall Street secrets to reduce the cost of cloud infrastructure (2019, August 16)

retrieved 21 June 2024 from <https://techxplore.com/news/2019-08-wall-street-secrets-cloud-infrastructure.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.