Using machine learning to hunt down cybercriminals

October 9 2019, by Adam Conner-Simons

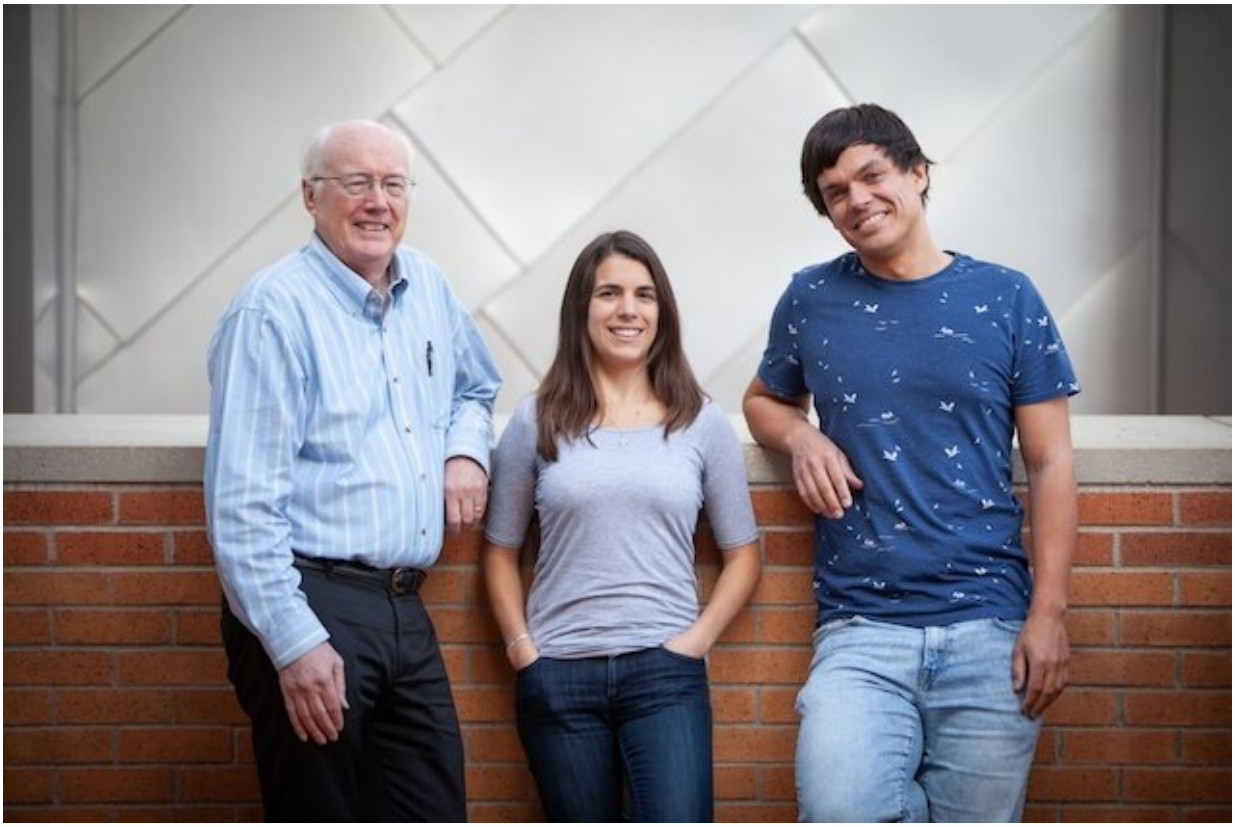

Left to right: senior research scientist David Clark, graduate student Cecilia Testart, and postdoc Philipp Richter. Credit: Jason Dorfman, MIT CSAIL

Hijacking IP addresses is an increasingly popular form of cyber-attack. This is done for a range of reasons, from sending spam and malware to

stealing Bitcoin. It's estimated that in 2017 alone, routing incidents such as IP hijacks affected more than 10 percent of all the world's routing domains. There have been major incidents at Amazon and Google and even in nation-states—a study last year suggested that a Chinese telecom company used the approach to gather intelligence on western countries by rerouting their internet traffic through China.

Existing efforts to detect IP hijacks tend to look at specific cases when they're already in process. But what if we could predict these incidents in advance by tracing things back to the hijackers themselves?

That's the idea behind a new machine-learning system developed by researchers at MIT and the University of California at San Diego (UCSD). By illuminating some of the common qualities of what they call "serial hijackers," the team trained their system to be able to identify roughly 800 suspicious networks—and found that some of them had been hijacking IP addresses for years.

"Network operators normally have to handle such incidents reactively and on a case-by-case basis, making it easy for cybercriminals to continue to thrive," says lead author Cecilia Testart, a graduate student at MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) who will present the paper at the ACM Internet Measurement Conference in Amsterdam on Oct. 23. "This is a key first step in being able to shed light on serial hijackers' behavior and proactively defend against their attacks."

The paper is a collaboration between CSAIL and the Center for Applied Internet Data Analysis at UCSD's Supercomputer Center. The paper was written by Testart and David Clark, an MIT senior research scientist, alongside MIT postdoc Philipp Richter and data scientist Alistair King as well as research scientist Alberto Dainotti of UCSD.

The nature of nearby networks

IP hijackers exploit a key shortcoming in the Border Gateway Protocol (BGP), a routing mechanism that essentially allows different parts of the internet to talk to each other. Through BGP, networks exchange routing information so that data packets find their way to the correct destination.

In a BGP hijack, a malicious actor convinces nearby networks that the best path to reach a specific IP address is through their network. That's unfortunately not very hard to do, since BGP itself doesn't have any security procedures for validating that a message is actually coming from the place it says it's coming from.

"It's like a game of Telephone, where you know who your nearest neighbor is, but you don't know the neighbors five or 10 nodes away," says Testart.

In 1998 the U.S. Senate's first-ever cybersecurity hearing featured a team of hackers who claimed that they could use IP hijacking to take down the Internet in under 30 minutes. Dainotti says that, more than 20 years later, the lack of deployment of security mechanisms in BGP is still a serious concern.

To better pinpoint serial attacks, the group first pulled data from several years' worth of network operator mailing lists, as well as historical BGP data taken every five minutes from the global routing table. From that, they observed particular qualities of malicious actors and then trained a machine-learning model to automatically identify such behaviors.

The system flagged networks that had several key characteristics, particularly with respect to the nature of the specific blocks of IP addresses they use:

- Volatile changes in activity: Hijackers' address blocks seem to disappear much faster than those of legitimate networks. The average duration of a flagged network's prefix was under 50 days, compared to almost two years for legitimate networks.
- Multiple address blocks: Serial hijackers tend to advertise many more blocks of IP addresses, also known as "network prefixes."
- IP addresses in multiple countries: Most networks don't have foreign IP addresses. In contrast, for the networks that serial hijackers advertised that they had, they were much more likely to be registered in different countries and continents.

Identifying false positives

Testart said that one challenge in developing the system was that events that look like IP hijacks can often be the result of human error, or otherwise legitimate. For example, a network operator might use BGP to defend against distributed denial-of-service attacks in which there's huge amounts of traffic going to their network. Modifying the route is a legitimate way to shut down the attack, but it looks virtually identical to an actual hijack.

Because of this issue, the team often had to manually jump in to identify false positives, which accounted for roughly 20 percent of the cases identified by their classifier. Moving forward, the researchers are hopeful that future iterations will require minimal human supervision and could eventually be deployed in production environments.

"The authors' results show that past behaviors are clearly not being used to limit bad behaviors and prevent subsequent attacks," says David Plonka, a senior research scientist at Akamai Technologies who was not involved in the work. "One implication of this work is that network operators can take a step back and examine global Internet routing across years, rather than just myopically focusing on individual incidents."

As people increasingly rely on the Internet for critical transactions, Testart says that she expects IP hijacking's potential for damage to only get worse. But she is also hopeful that it could be made more difficult by new security measures. In particular, large backbone networks such as AT&T have recently announced the adoption of resource public key infrastructure (RPKI), a mechanism that uses cryptographic certificates to ensure that a network announces only its legitimate IP addresses.

"This project could nicely complement the existing best solutions to prevent such abuse that include filtering, antispoofing, coordination via contact databases, and sharing routing policies so that other networks can validate it," says Plonka. "It remains to be seen whether misbehaving networks will continue to be able to game their way to a good reputation. But this work is a great way to either validate or redirect the network operator community's efforts to put an end to these present dangers."

More information: Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table. Internet Measurement Conference (IMC '19), October 21–23, 2019, Amsterdam, Netherlands. doi.org/10.1145/3355369.3355581

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology