Researchers devise new way to evaluate cybersecurity methods

June 28 2023, by Adam Zewe

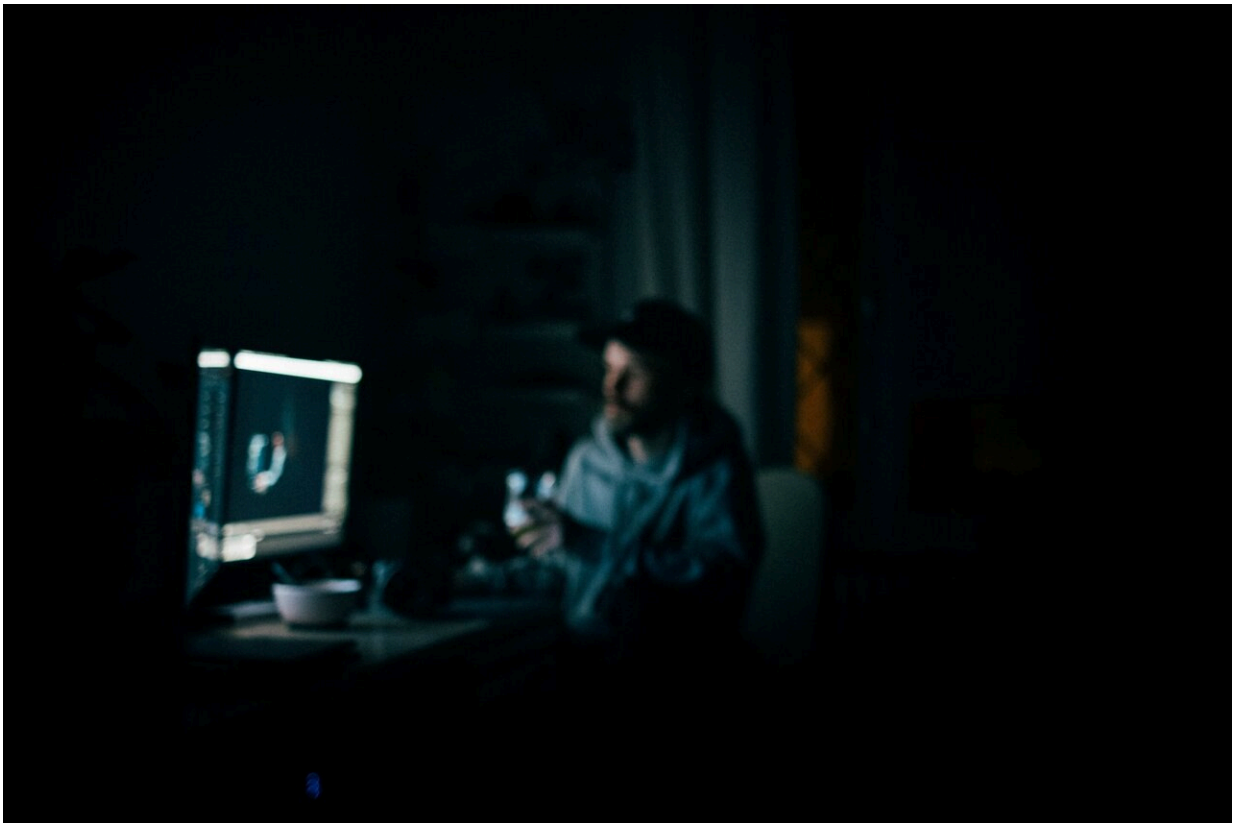

Credit: Unsplash/CC0 Public Domain

A savvy hacker can obtain secret information, such as a password, by observing a computer program's behavior, like how much time that program spends accessing the computer's memory.

Security approaches that completely block these "side-channel attacks" are so computationally expensive that they aren't feasible for many real-world systems. Instead, engineers often apply what are known as obfuscation schemes that seek to limit, but not eliminate, an attacker's ability to learn secret information.

To help engineers and scientists better understand the effectiveness of different obfuscation schemes, MIT researchers created a framework to quantitatively evaluate how much information an attacker could learn from a victim program with an obfuscation scheme in place.

Their framework, called Metior, allows the user to study how different victim programs, attacker strategies, and obfuscation scheme configurations affect the amount of sensitive information that is leaked. The framework could be used by engineers who develop microprocessors to evaluate the effectiveness of multiple security schemes and determine which architecture is most promising early in the chip design process.

"Metior helps us recognize that we shouldn't look at these security schemes in isolation. It is very tempting to analyze the effectiveness of an obfuscation scheme for one particular victim, but this doesn't help us understand why these attacks work. Looking at things from a higher level gives us a more holistic picture of what is actually going on," says Peter Deutsch, a graduate student and lead author of an open-access paper on Metior.

Deutsch's co-authors include Weon Taek Na, an MIT graduate student in electrical engineering and computer science; Thomas Bourgeat Ph.D. '23, an assistant professor at the Swiss Federal Institute of Technology (EPFL); Joel Emer, an MIT professor of the practice in computer science and electrical engineering; and senior author Mengjia Yan, the Homer A. Burnell Career Development Assistant Professor of Electrical

Engineering and Computer Science (EECS) at MIT and a member of the Computer Science and Artificial Intelligence Laboratory (CSAIL). The research was presented last week at the International Symposium on Computer Architecture in Orlando, Florida.

Illuminating obfuscation

While there are many obfuscation schemes, popular approaches typically work by adding some randomization to the victim's behavior to make it harder for an attacker to learn secrets. For instance, perhaps an obfuscation scheme involves a program accessing additional areas of the computer memory, rather than only the area it needs to access, to confuse an attacker. Others adjust how often a victim accesses memory or another a shared resource so an attacker has trouble seeing clear patterns.

But while these approaches make it harder for an attacker to succeed, some amount of information from the victim still "leaks" out. Yan and her team want to know how much.

They had previously developed CaSA, a tool to quantify the amount of information leaked by one particular type of obfuscation scheme. But with Metior, they had more ambitious goals. The team wanted to derive a unified model that could be used to analyze any obfuscation scheme—even schemes that haven't been developed yet.

To achieve that goal, they designed Metior to map the flow of information through an obfuscation scheme into random variables. For instance, the model maps the way a victim and an attacker access shared structures on a computer chip, like memory, into a mathematical formulation.

One Metior derives that mathematical representation, the framework

uses techniques from information theory to understand how the attacker can learn information from the victim. With those pieces in place, Metior can quantify how likely it is for an attacker to successfully guess the victim's secret information.

"We take all of the nitty-gritty elements of this microarchitectural side-channel and map it down to, essentially, a math problem. Once we do that, we can explore a lot of different strategies and better understand how making small tweaks can help you defend against information leaks," Deutsch says.

Surprising insights

They applied Metior in three case studies to compare attack strategies and analyze the information leakage from state-of-the-art obfuscation schemes. Through their evaluations, they saw how Metior can identify interesting behaviors that weren't fully understood before.

For instance, a prior analysis determined that a certain type of side-channel attack, called probabilistic prime and probe, was successful because this sophisticated attack includes a preliminary step where it profiles a victim system to understand its defenses.

Using Metior, they show that this advanced attack actually works no better than a simple, generic attack and that it exploits different victim behaviors than researchers previously thought.

Moving forward, the researchers want to continue enhancing Metior so the framework can analyze even very complicated obfuscation schemes in a more efficient manner. They also want to study additional obfuscation schemes and types of victim programs, as well as conduct more detailed analyses of the most popular defenses.

Ultimately, the researchers hope this work inspires others to study microarchitectural security evaluation methodologies that can be applied early in the chip design process.

"Any kind of microprocessor development is extraordinarily expensive and complicated, and design resources are extremely scarce. Having a way to evaluate the value of a security feature is extremely important before a company commits to microprocessor development. This is what Metior allows them to do in a very general way," Emer says.

More information: Peter W. Deutsch et al, Metior: A Comprehensive Model to Evaluate Obfuscating Side-Channel Defense Schemes, *Proceedings of the 50th Annual International Symposium on Computer Architecture* (2023). DOI: 10.1145/3579371.3589073

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