Researchers mitigate potential side-channel attack vulnerability in multicore processors

August 11 2022, by Adam Zewe

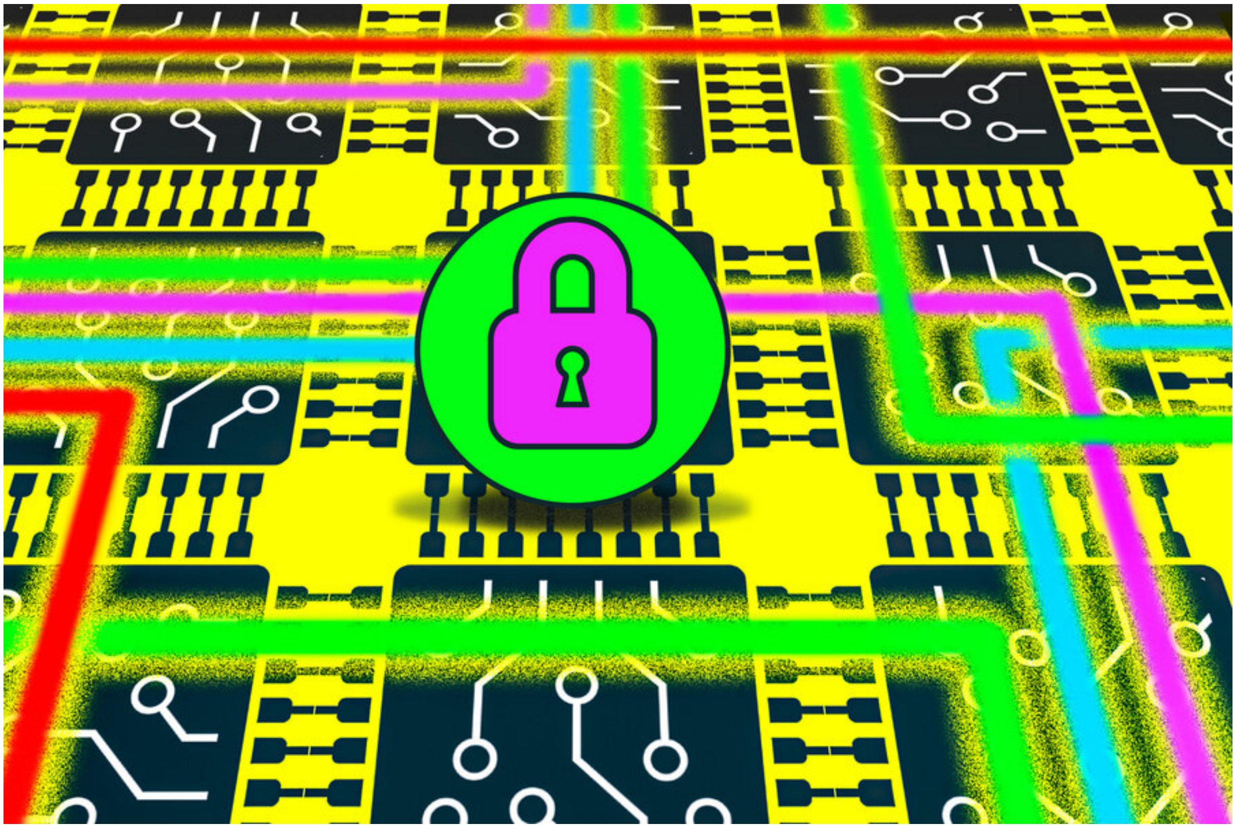

MIT researchers have shown that a component of modern computer processors that enables different areas of the chip to communicate with each other is susceptible to a side-channel attack. Credit: Jose-Luis Olivares, MIT

A component of computer processors that connects different parts of the chip can be exploited by malicious agents who seek to steal secret information from programs running on the computer, MIT researchers have found.

Modern computer processors contain many computing units, called cores, which share the same hardware resources. The on-chip interconnect is the component that enables these cores to communicate with each other. But when programs on multiple cores run simultaneously, there is a chance they can delay one another when they use the interconnect to send data across the chip at the same time.

By monitoring and measuring these delays, a malicious agent could conduct what is known as a "side-channel attack" and reconstruct secret information that is stored in a program, such as a cryptographic key or password.

MIT researchers reverse-engineered the on-chip interconnect to study how this kind of attack would be possible. Drawing on their discoveries, they built an analytical model of how traffic flows between the cores on a processor, which they used to design and launch surprisingly effective side-channel attacks. Then they developed two mitigation strategies that enable a user to improve security without making any physical changes to the computer chip.

"A lot of current side-channel defenses are ad hoc—we see a little bit of leakage here and we patch it. We hope our approach with this analytical model pushes more systematic and robust defenses that eliminate whole classes of attacks at the same time," says co-lead author Miles Dai, MEng '21.

Dai wrote the paper with co-lead author Riccardo Paccagnella, a graduate student at the University of Illinois at Urbana-Champaign; Miguel Gomez-Garcia '22; John McCalpin, a research scientist at Texas

Advanced Computing Center; and senior author Mengjia Yan, the Homer A. Burnell Career Development Assistant Professor of Electrical Engineering and Computer Science (EECS) and a member of the Computer Science and Artificial Intelligence Laboratory (CSAIL). The research is being presented at the USENIX Security Conference.

Probing processors

A modern processor is like a two-dimensional grid, with multiple cores laid out in rows and columns. Each core has its own cache where data are stored, and there is also a larger cache that is shared across the entire processor. When a program located on one core needs to access data in a cache that is on another core or in the shared cache, it must use the on-chip interconnect to send this request and retrieve the data.

Though it is a large component of the processor, the on-chip interconnect remains understudied because it is difficult to attack, Dai explains. A hacker needs to launch the attack when traffic from two cores is actually interfering with each other, but since traffic spends so little time in the interconnect, it is difficult to time the attack just right. The interconnect is also complex, and there are multiple paths traffic can take between cores.

To study how traffic flows on the interconnect, the MIT researchers created programs that would intentionally access memory caches located outside their local cores.

"By testing out different situations, trying different placements, and swapping out locations of these programs on the processor, we can understand what the rules are behind traffic flows on the interconnect," Dai says.

They discovered that the interconnect is like a highway, with multiple lanes going in every direction. When two traffic flows collide, the interconnect uses a priority arbitration policy to decide which traffic flow gets to go first. More "important" requests take precedence, like those from programs that are critical to a computer's operations.

Using this information, the researchers built an analytical model of the processor that summarizes how traffic can flow on the interconnect. The model shows which cores would be most vulnerable to a side-channel attack. A core would be more vulnerable if it can be accessed through many different lanes. An attacker could use this information to select the best core to monitor to steal information from a victim program.

"If the attacker understands how the interconnect works, they can set themselves up so the execution of some sensitive code would be observable through interconnect contention. Then they can extract, bit by bit, some secret information, like a cryptographic key," Paccagnella explains.

Effective attacks

When the researchers used this model to launch side-channel attacks, they were surprised by how quickly the attacks worked. They were able to recover full cryptographic keys from two different victim programs.

After studying these attacks, they used their analytical model to design two mitigation mechanisms.

In the first strategy, the system administrator would use the model to identify which cores are most vulnerable to attacks and then schedule sensitive software to run on less vulnerable cores. For the second mitigation strategy, the administrator could reserve cores located around a susceptible program and run only trusted software on those cores.

The researchers found that both mitigation strategies were able to significantly reduce the accuracy of side-channel attacks. Neither requires the user to make any changes to the physical hardware, so the mitigations would be relatively easy to implement, Dai says.

Ultimately, they hope their work inspires more researchers to study the security of on-chip interconnects, Paccagnella says.

"We hope this work highlights how the on-chip interconnect, which is such a large component of computer processors, remains an overlooked attack surface. In the future, as we build systems that have stronger isolation properties, we should not ignore the interconnect," he adds.

More information: Don't Mesh Around: Side-Channel Attacks and Mitigations on Mesh Interconnects. people.csail.mit.edu/mengjia/d … Attack_USENIX_22.pdf

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

Citation: Researchers mitigate potential side-channel attack vulnerability in multicore processors (2022, August 11) retrieved 20 March 2024 from https://techxplore.com/news/2022-08-mitigate-potential-side-channel-vulnerability-multicore.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.