


Scientists discover new vulnerability affecting computers globally

30 April 2021, by Audra Book



Credit: Pixabay/CC0 Public Domain

In 2018, industry and academic researchers revealed a potentially devastating hardware flaw that made computers and other devices worldwide vulnerable to attack.

Researchers named the vulnerability [Spectre](#) because the flaw was built into modern computer processors that get their speed from a technique called "speculative execution," in which the [processor](#) predicts instructions it might end up executing and preps by following the predicted path to pull the instructions from memory. A Spectre attack tricks the processor into executing instructions along the wrong path. Even though the processor recovers and correctly completes its task, hackers can access confidential data while the processor is heading the wrong way.

Since Spectre was discovered, the world's most talented computer scientists from industry and academia have worked on software patches and hardware defenses, confident they've been able to protect the most vulnerable points in the speculative execution process without slowing down computing speeds too much.

They will have to go back to the drawing board.

A team of University of Virginia School of Engineering computer science researchers has uncovered a line of attack that breaks all Spectre defenses, meaning that billions of computers and other devices across the globe are just as vulnerable today as they were when Spectre was first announced. The team reported its discovery to international chip makers in April and will present the new challenge at a worldwide computing architecture conference in June.

The researchers, led by Ashish Venkat, William Wulf Career Enhancement Assistant Professor of Computer Science at UVA Engineering, found a whole new way for hackers to exploit something called a "micro-op cache," which speeds up computing by storing simple commands and allowing the processor to fetch them quickly and early in the speculative execution process. Micro-op caches have been built into Intel computers manufactured since 2011.

Venkat's team discovered that hackers can steal data when a processor fetches commands from the micro-op cache.

"Think about a hypothetical airport security scenario where TSA lets you in without checking your boarding pass because (1) it is fast and efficient, and (2) you will be checked for your boarding pass at the gate anyway," Venkat said. "A computer processor does something similar. It predicts that the check will pass and could let instructions into the pipeline. Ultimately, if the prediction is incorrect, it will throw those instructions out of the pipeline, but this might be too late because those instructions could leave side-effects while waiting in the pipeline that an attacker could later exploit to infer secrets such as a password."

Because all current Spectre defenses protect the processor in a later stage of speculative execution,

they are useless in the face of Venkat's team's new attacks. Two variants of the attacks the team discovered can steal speculatively accessed information from Intel and AMD processors.

"Intel's suggested defense against Spectre, which is called LFENCE, places sensitive code in a waiting area until the security checks are executed, and only then is the sensitive code allowed to execute," Venkat said. "But it turns out the walls of this waiting area have ears, which our attack exploits. We show how an attacker can smuggle secrets through the micro-op cache by using it as a covert channel."

Venkat's team includes three of his computer science graduate students, Ph.D. student Xida Ren, Ph.D. student Logan Moody and master's degree recipient Matthew Jordan. The UVA team collaborated with Dean Tullsen, professor of the Department of Computer Science and Engineering at the University of California, San Diego, and his Ph.D. student Mohammadkazem Taram to reverse-engineer certain undocumented features in Intel and AMD processors.

They have detailed the findings in their paper: "I See Dead ?ops: Leaking Secrets via Intel/AMD Micro-Op Caches."

This newly discovered vulnerability will be much harder to fix.

"In the case of the previous Spectre attacks, developers have come up with a relatively easy way to prevent any sort of attack without a major performance penalty" for computing, Moody said. "The difference with this attack is you take a much greater performance penalty than those previous attacks."

"Patches that disable the micro-op cache or halt speculative execution on legacy hardware would effectively roll back critical performance innovations in most modern Intel and AMD processors, and this just isn't feasible," Ren, the lead student author, said.

"It is really unclear how to solve this problem in a way that offers high performance to legacy

hardware, but we have to make it work," Venkat said. "Securing the micro-op cache is an interesting line of research and one that we are considering."

Venkat's team has disclosed the vulnerability to the product security teams at Intel and AMD. Ren and Moody gave a tech talk at Intel Labs worldwide April 27 to discuss the impact and potential fixes. Venkat expects computer scientists in academia and industry to work quickly together, as they did with Spectre, to find solutions.

The team's paper has been accepted by the highly competitive International Symposium on Computer Architecture, or ISCA. The annual ISCA conference is the leading forum for new ideas and research results in computer architecture and will be held virtually in June.

Venkat is also working in close collaboration with the Processor Architecture Team at Intel Labs on other microarchitectural innovations, through the [National Science Foundation/Intel Partnership on Foundational Microarchitecture Research Program](#).

Venkat was well prepared to lead the UVA research team into this discovery. He has forged a long-running partnership with Intel that started in 2012 when he interned with the company while he was a computer science graduate student at the University of California, San Diego.

This research, like other projects Venkat leads, is funded by the National Science Foundation and Defense Advanced Research Projects Agency.

Venkat is also one of the university researchers who co-authored a paper with collaborators Mohammadkazem Taram and Tullsen from UC San Diego that introduce a more targeted microcode-based defense against Spectre. Context-sensitive fencing, as it is called, allows the processor to patch running code with speculation fences on the fly.

Introducing one of just a handful more targeted microcode-based defenses developed to stop Spectre in its tracks, "[Context-Sensitive Fencing: Securing Speculative Execution via Microcode Customization](#)" was published at the *ACM*

International Conference on Architectural Support for Programming Languages and Operating Systems in April 2019. The paper was also selected as a top pick among all computer architecture, [computer](#) security, and VLSI design conference papers published in the six-year period between 2014 and 2019.

The new Spectre variants Venkat's team discovered even break the context-sensitive fencing mechanism outlined in Venkat's award-winning paper. But in this type of research, breaking your own defense is just another big win. Each security improvement allows researchers to dig even deeper into the hardware and uncover more flaws, which is exactly what Venkat's research group did.

Provided by University of Virginia School of Engineering and Applied Science

APA citation: Scientists discover new vulnerability affecting computers globally (2021, April 30) retrieved 11 May 2021 from <https://techxplore.com/news/2021-04-scientists-vulnerability-affecting-globally.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.