Computer scientists unveil novel attacks on cybersecurity

April 27 2024, by Katie E. Ismael

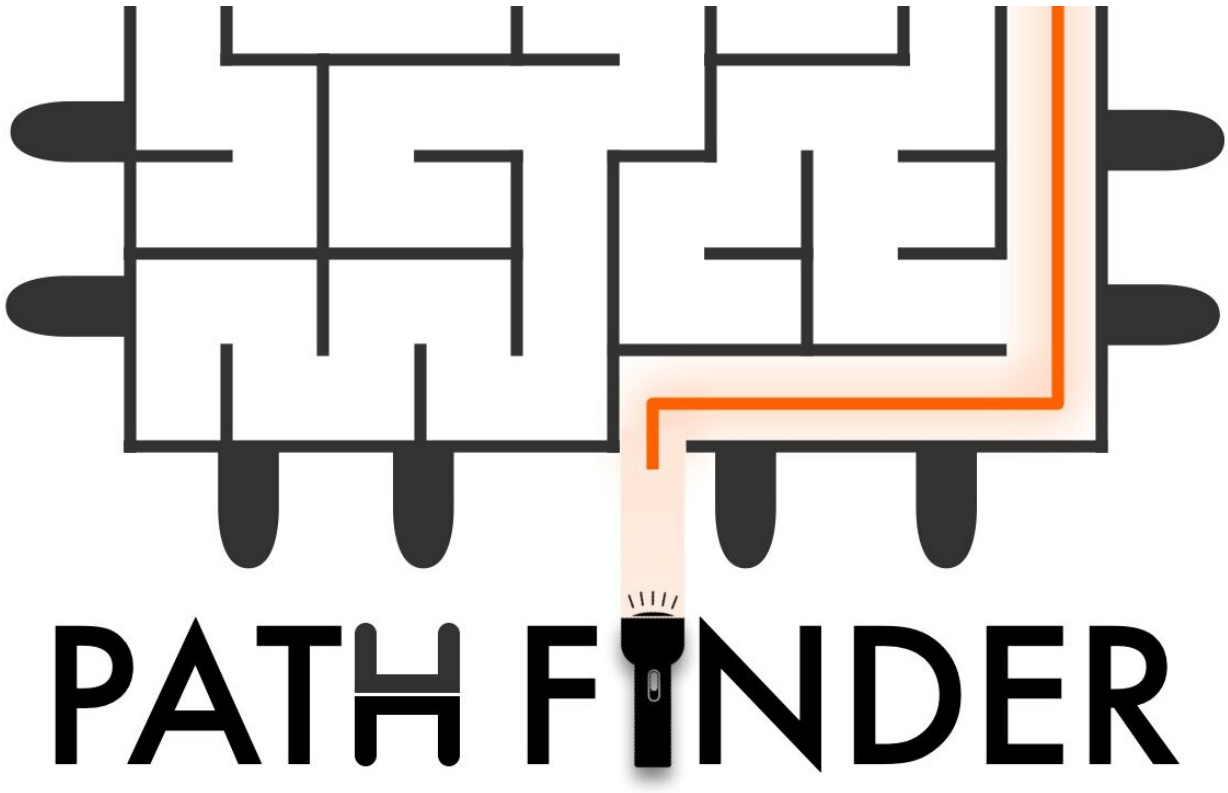

The new paper, "Pathfinder: High-Resolution Control-Flow Attacks Exploiting the Conditional Branch Predictor," details two novel attacks that could compromise the billions of Intel processors in use. Credit: Hosein Yavarzadeh

Researchers have found two novel types of attacks that target the conditional branch predictor found in high-end Intel processors, which

could be exploited to compromise billions of processors currently in use.

The multi-university and industry research team led by computer scientists at University of California San Diego will present their work at the 2024 ACM ASPLOS Conference that begins tomorrow. The paper, "Pathfinder: High-Resolution Control-Flow Attacks Exploiting the Conditional Branch Predictor," is based on findings from scientists from UC San Diego, Purdue University, Georgia Tech, the University of North Carolina Chapel Hill and Google.

They discover a unique attack that is the first to target a feature in the branch predictor called the Path History Register, which tracks both branch order and branch addresses. As a result, more information with more precision is exposed than with prior attacks that lacked insight into the exact structure of the branch predictor.

Their research has resulted in Intel and Advanced Micro Devices (AMD) addressing the concerns raised by the researchers and advising users about the security issues. Today, Intel is set to issue a Security Announcement, while AMD will release a Security Bulletin.

In software, frequent branching occurs as programs navigate different paths based on varying data values. The direction of these branches, whether "taken" or "not taken," provides crucial insights into the executed program data. Given the significant impact of branches on modern processor performance, a crucial optimization known as the "branch predictor" is employed. This predictor anticipates future branch outcomes by referencing past histories stored within prediction tables. Previous attacks have exploited this mechanism by analyzing entries in these tables to discern recent branch tendencies at specific addresses.

In this new study, researchers leverage modern predictors' utilization of a Path History Register (PHR) to index prediction tables. The PHR

records the addresses and precise order of the last 194 taken branches in recent Intel architectures. With innovative techniques for capturing the PHR, the researchers demonstrate the ability to not only capture the most recent outcomes but also every branch outcome in sequential order. Remarkably, they uncover the global ordering of all branches. Despite the PHR typically retaining the most recent 194 branches, the researchers present an advanced technique to recover a significantly longer history.

"We successfully captured sequences of tens of thousands of branches in precise order, utilizing this method to leak secret images during processing by the widely used image library, libjpeg," said Hosein Yavarzadeh, a UC San Diego Computer Science and Engineering Department Ph.D. student and lead author of the paper.

The researchers also introduce an exceptionally precise Spectre-style poisoning attack, enabling attackers to induce intricate patterns of branch mispredictions within victim code. "This manipulation leads the victim to execute unintended code paths, inadvertently exposing its confidential data," said UC San Diego computer science Professor Dean Tullsen.

"While prior attacks could misdirect a single branch or the first instance of a branch executed multiple times, we now have such precise control that we could misdirect the 732nd instance of a branch taken thousands of times," said Tullsen.

The team presents a proof-of-concept where they force an encryption algorithm to transiently exit earlier, resulting in the exposure of reduced-round ciphertext. Through this demonstration, they illustrate the ability to extract the secret AES encryption key.

"Pathfinder can reveal the outcome of almost any branch in almost any

victim program, making it the most precise and powerful microarchitectural control-flow extraction attack that we have seen so far," said Kazem Taram, an assistant professor of computer science at Purdue University and a UC San Diego computer science Ph.D. graduate.

In addition to Dean Tullsen and Hosein Yavarzadeh, other UC San Diego co-authors are. Archit Agarwal and Deian Stefan. Other co-authors include Christina Garman and Kazem Taram, Purdue University; Daniel Moghimi, Google; Daniel Genkin, Georgia Tech; Max Christman and Andrew Kwong, University of North Carolina Chapel Hill.

Researchers communicated the security findings outlined in the paper to both Intel and AMD in November 2023. Intel has informed other affected hardware/software vendors about the issues. Both Intel and AMD plan to address the concerns raised in the paper today through a Security Announcement and a Security Bulletin (AMD-SB-7015), respectively. The findings have been shared with the Vulnerability Information and Coordination Environment (VINCE), Case VU#157097: Class of Attack Primitives Enable Data Exposure on High End Intel CPUs.

More information: Hosein Yavarzadeh et al, Pathfinder: High-Resolution Control-Flow Attacks Exploiting the Conditional Branch Predictor, *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3* (2024). DOI: 10.1145/3620666.3651382

Provided by University of California - San Diego