

Planting ideas in a computer's head: Researchers find new attack on AMD computer chips

August 9 2023, by Oliver Morsch



The hardware used by the ETH researchers with one of the computer chips that are susceptible to the Inception attack. Credit: Kaveh Razavi / ETH Zurich

Everyone has, at one time or another, experienced how dreams can influence our moods and actions. However, putting an idea in somebody else's head while they are dreaming in order to make them do something

specific once they wake up is still the stuff of science fiction. In the 2010 movie "Inception," Leonardo di Caprio's character tries to get the heir of a wealthy businessman to break up his father's empire. To do so, he shares a dream with the heir, in which through clever manipulation, the heir's convictions about his father are subtly altered, leading him to abandon his late father's business.

While sharing dreams and planting such ideas is impossible in reality, something very similar has recently been achieved in the world of computers. A team of researchers at ETH Zurich led by Kaveh Razavi, professor in the Department of Information Technology and Engineering, has demonstrated a serious vulnerability of certain CPUs (central processing units) whereby an attacker can plant the equivalent of an idea in a victim's CPU, coax it into executing certain commands, and thus retrieve information. Razavi and his colleagues present [their research](#) at the conference [USENIX Security 2023](#) this week.

A complex attack

While Razavi's research paper contains names that are reminiscent of James Bond and disaster movies—"Spectre" and "Meltdown" make an appearance—the bulk of it is intricate computer science.

"In fact, much like the movie of the same name, the Inception attack is particularly complex and difficult to explain," says master's student Daniël Trujillo, who found this new attack during his thesis work in Razavi's group, supervised by Ph.D. student Johannes Wikner.

"Still," Wikner adds, "the crux of the matter with all these attacks is rather simple—namely, the fact that a computer's CPU has to make guesses all the time, and those guesses can be tampered with."

In modern computers, guesses are needed because during the execution

of a program—a game, say, or a web browser—the CPU has to make hundreds of millions of decisions per second. At certain points in the execution, the following command may depend on a choice made based on some information that has to be retrieved from the computer's memory. CPUs have become incredibly fast in recent years, but the speed at which data can be transferred from the memory (DRAM) to the CPU has not been able to keep up with that acceleration. As a result, the CPU would have to spend a lot of its time waiting for fresh data in order to make a decision.

Speeding up by guessing

This is where guessing comes into play: Based on past experience, the CPU creates a kind of look-up table and uses it to make a guess as to the most likely next step, which it then executes. In the vast majority of cases, the CPU is right and thus saves a lot of valuable computing time. Occasionally, however, it makes the wrong guess, and such a misprediction can be exploited by an attacker to gain access to sensitive information.

"The Spectre attack, which was discovered in 2018, is based on such mispredictions," says Razavi, "but initially it seemed that manufacturers had found ways to mitigate it." In fact, chip manufacturers have provided features for partly deleting the look-up table when switching between security contexts (that is, when the sensitive kernel of the computer is accessed) or adding a bit of information that tells the CPU whether or not a prediction in the look-up table was created in the kernel, and can therefore be trusted.

Planting an idea in a CPU

Nevertheless, Razavi and his co-workers set out to test whether even

with the new security measures an attack could be launched. After a lengthy search, they stumbled upon something strange: "It looked as though we could make the CPUs manufactured by AMD believe that they had seen certain instructions before, whereas in reality that had never happened," says Trujillo. Just like in the movie, the researchers could plant an idea in the CPU while it was—in a sense—dreaming.

As a consequence, the look-up table—which the CPU continuously creates from past instructions—could, once again, be manipulated. Since the CPU was convinced that the entries in the look-up table originated from instructions it had seen before, the security feature that was meant to ensure that only trustworthy predictions are taken into consideration was bypassed. In this way, the ETH researchers were able to leak data from anywhere in the computer's memory, including [sensitive information](#) such as the hash of the root password.

Serious vulnerability

That, of course, is a very serious [security vulnerability](#), so Razavi informed AMD in February to make sure they had time to come up with a patch before the research paper was published (AMD assigned the number CVE-2023-20569 to the vulnerability).

"We have shown this concept of a new class of dangerous attacks, which is particularly relevant in the context of cloud computing, where several customers share the same hardware," Razavi says, "It also raises questions for the future." For instance, he wants to find out if there are other, similar [attacks](#) and whether an "Inception"-like attack is also possible on CPUs from other manufacturers.

More information: Trujillo D, Wikner J, Razavi K: external page [Inception: Exposing New Attack Surfaces with Training in Transient Execution](#). 32nd Usenix Security Symposium, 2023.

<https://www.usenix.org/conference/usenixsecurity23/presentation/trujillo>

Provided by ETH Zurich

Citation: Planting ideas in a computer's head: Researchers find new attack on AMD computer chips (2023, August 9) retrieved 28 April 2024 from <https://techxplore.com/news/2023-08-ideas-amd-chips.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.