AXIADO

# The Security Hole that Caused Everyone to Give Up

by Axel Kloth

Early in 2018, researchers in the US and Europe revealed two massive hardware security holes dubbed Meltdown and Spectre. At first, the world panicked. The processor industry leapt into the fray creating patches to these holes, and the panic subsided.

Then something interesting happened. Users started complaining of significant performance reductions. Customer service call centers at Apple, Dell and Microsoft started telling people how to turn off the performance-sucking features, which, in turn, made the systems vulnerable again. Let's take a step back and see how we got here.

## A Short History of a Bad Decision

Multi-core processors were introduced to deal with, among other issues, limits in the performance of single-core designs. The design allowed the processor to offload repetitive or time-intensive operations to supplemental cores and speed up overall processing.

A problem arose, however, when it was discovered that doubling the number of processor cores did not mean a doubling of performance, because a computationally intensive process that takes a set amount of time on a single processor cannot go faster than that. This is an oversimplified description of Amdahl’s Law, but appropriate for this discussion. A follow-up to Amdahl’s Law by Dr. John Gustafson, often called Gustafson’s Law of Parallel Speed-Up, provided a more optimistic view of multi-core performance enhancements, but still showed a diminishing set of gains with more cores on a processor chip.

A little over 20 years ago, users started noticing the effects of Amdahl’s and Gustafson’s laws and started complaining: “Just what are we paying for?” The answer from the industry was a switch from in-order processing to out-of-order processing.

Users were frustrated with the progress of multi-core processors—they wanted a noticeable increase in performance. The most obvious way to do that was by allowing the processor to execute instructions out-of-order; i.e., reorder them when the individual results were available, and combine them in the result.

This improved performance but came with a caveat: When a processor executes instructions out of order, it may encounter dependencies or a branch without knowing which part of the branch to retrieve data from, so it has to execute data speculatively to both parts of the branch. The more nested branches there are, the more speculative execution and the data cached until instructions are available. Now the cache algorithm has to speculate based on a jump or a branch that hasn't yet been determined.

This process leaves a large trail of data and instructions in the cache, and a simple cache dump would now reveal all these instructions and all of the data without encryption.

A cache dump is always allowed by the operating system kernel and by a debugger. Those are the two pieces of software that are allowed to dump the cache contents to the DRAM.

Whenever an operating system kernel goes into panic mode, it dumps the cache so developers can examine the cache contents and compare the expected contents versus the actual contents in order to debug the hardware and the software.

Getting a kernel into panic mode isn't that difficult, and any hacker can do it, which gives them access to all the data in the cache.

This is, in essence, the massive security hole that makes up Meltdown and Spectre. While there are no available commercial tools to exploit these holes against a target, it is not hard to imagine that state actors, criminals and terrorist organizations have the knowledge to exploit any system based on its weakness. Of course, they would not be published, unless the intrusion itself was discovered as it was with Stuxnet. However, Stuxnet was only discovered years after the fact.

## Good News and Bad News

First the good news. The main corporate players, including Apple, Intel and Arm, all issued patches to plug this massive security hole. Those patches essentially reduced the degree to which speculative branching of out-of-order processing is allowed.

Now the bad news. Those patches sacrifice the 10-15 percent performance gain realized over an in-order machine.

When the patches were applied, it wasn't long before customer service teams started fielding calls about the reduction in performance and the fix was simple—just go to the backdoor and start turning off the new security features. Surprise! Performance returned. Of course, and so did the Meltdown and Spectre holes. What’s more, hackers can easily turn them off as well, even if the user did not.

So, the cure was neither permanent nor truly effective. It does not matter whether you are dealing with a mobile device, a desktop computer or a public or private data center, this security hole elevates the attacker to the privilege level of a super user or root. Once they have taken control of a system they can pretty much do whatever they want with it.

## Foundational Weakness

Meltdown and Spectre are both made possible by the foundational architecture of the processor and its ability to execute instructions out of order. That is true for all commercial processors, including those made by Intel, AMD and those based on Arm IP cores. This is a basic design choice to improve on the single-core, single-thread performance of their processors. Now it is apparent that it was a bad choice: trading security for any performance increase is an undesirable consequence and one that cannot be fixed in the current processor design paradigm.

At the core of the despair now sweeping the electronics industry is the knowledge that no matter what is done, any security measure can be bypassed with a modicum of technical expertise. The security solutions put in place are little more than placebos and companies are resigned to paying for breaches that impact users with identity protection insurance. Some flat-out state that the security of the user is on the user, and wash their hands of any responsibility—this in spite of the fact that they sell a product they know is inherently insecure. Nothing short of the elimination of the current technology can resolve the problem.

Luckily, however, that is not out of the question.

## Hope on the Horizon

At the Hot Chips Conference in August 2018, Professor Mark Hill of the University of Wisconsin–Madison said the answer might lie in specialized cores, flushing caches on context switches and business ideas like charging more for exclusive virtual machines. Ultimately, however, he and several other panelists said it might require a new Architecture 2.0.

That is exactly what Axiado hopes to accomplish. The Axiado firewall processor restores the in-order paradigm on a core so efficient that it is at least as fast as the most powerful out-of-order processor on the market. Because it uses in-order processing, the cache dump is protected from unauthorized access, so software developers can verify the correctness of their code. It is simply not vulnerable to an attack, together with a branch target injection or other means, that would fool the operating system into a cache dump into the DRAM where it is accessible by a user.

Modern processors all have enough cores, so adding more cores to make them faster is unnecessary. The issue is one of software and the ability of the operating system to fairly distribute loads onto many, many cores. Axiado has taken numerous precautions that allows them to do that more effectively and efficiently than what is currently on the market.

Axiado's processor cycle per instruction (CPI) values are similar to those found in other processors. Since their architecture uses an in-order processor pipeline, it is 10 to 15 percent below the performance of an out-of-order processor. However, there are a variety of reasons that it outperforms the so-called faster processors.

First, it allows operating systems to make better use of all the cores and accelerates those functions in targeted applications for the most computationally intensive programs and subroutines.

Hardware is more power effective, delivers better performance and provides lower latency than software-executed functions. Axiado’s firewall processor executes as many instructions as possible in hardware, not software, resulting in a device that outperforms a software-only based solution on any machine, including those with out-of-order processors.

For example, if encryption and decryption is executed in software versus hardware, hardware units can easily outperform a software solution by a factor of 10 or 20. If 10 percent of an application load is this kind of mathematically demanding software, then outsourcing that 10 percent to an accelerator will more than overcome the issues with in-order versus out-of-order machines. Even if the computational load is 10 percent and 10 percent is encryption and decryption, the Axiado firewall processor is already on par with an out-of-order machine. If the requirement for encryption and decryption makes up 50 percent of the performance on an out-of-order machine, then Axiado's solution will outperform the out-of-order machine. At that point, Axiado's firewall processor will be twice as fast, with twice the throughput compared to that out-of-order machine.

## Security is Not an Impossibility

If the industry continues on its current path, digital security is a quixotic effort. Simply put, we want to change the path. The Axiado firewall  processor is the foundation for a truly secure digital future. It offers the performance increase the user world clamors for and slams the door shut on malefactors at all levels.

It is time to rethink everything.