Securing Modern Processors: Secure TLBs and Beyond

1. Understand All Possible Timing Attacks

Cache timing channels use timing differences of memory related operations to derive information about sensitive data. Mostly-used two kinds of operations in the attacks:

- Cache access: cache hit (fast) vs. cache miss (slow)
- Cache flush: data is in the cache (slow) vs. not (fast)

Attack examples: Prime + Probe, Flush + Reload Attack

However, **any cache-like structure with varying timing can be vulnerable to timing attacks.**

Defenders need to deal with all types of attacks.

- Attacker just needs to find one that works.
- Defenders need to protect all types of attacks in anticipation of what the attacker can do.

2. Three-Step Modeling Approach

$Step1$ ⇝	$Step2$ ⇝	$Step3$ (fast/slow)
The initial state of the cache block set by a memory operation	Memory operation alters the state of the cache	Final memory operations and timing observation (fast/slow)

Derivation of all the vulnerabilities:

- Explicit enumerate all three-steps + different timings
- Implement cache three-step simulator + reduction rules to derive all the effective vulnerabilities

There are 88 possible strong cache timing attack types and 24 possible strong TLB timing attack types.

3. Projects Overview

This series of work focuses on the timing-based side-channel attacks due to the Translation Look-aside Buffers (TLBs) and caches. Our contributions are:

- Present the first **three-step modeling** approach that can be used to reason about all hit-based and miss-based timing-based attacks on TLBs and caches.

TLB:

- Design the first hardware defenses for TLB attacks: a **new Static-Partition (SP) TLB** and new **Random-Fill (RF) TLB** in a RISC-V Rocket Core processor.
- Evaluate the security of TLBs using our micro security benchmarks running on RISC-V simulation and show that results match with theoretical analysis.
- Test performance by running RSA tests alongside SPEC 2006 benchmarks under Linux on FPGAs.

The proposed secure TLBs can prevent not only prior attacks but also new timing-based attacks in TLBs found using our new three-step model.

Cache:

- Develop **first automatically-generated benchmark suite to evaluate processor caches' vulnerability** to all possible timing attacks.
- Test on **Intel** and **AMD** processors in the lab and on Amazon's Elastic Compute Cloud (Amazon EC2) servers and generate scores for different processors to understand which attacks they are vulnerable to, and to be able to evaluate and customize defenses.

4. Detailed Evaluation Results

Figure 1 shows the general flow and relationship of attacks, three-step modeling approach and the generated benchmark. Figure 2 shows the hardware microarchitecture of secure SP TLB. Figure 3 shows the evaluation result of 88 Strong types of cache vulnerabilities on different machines.

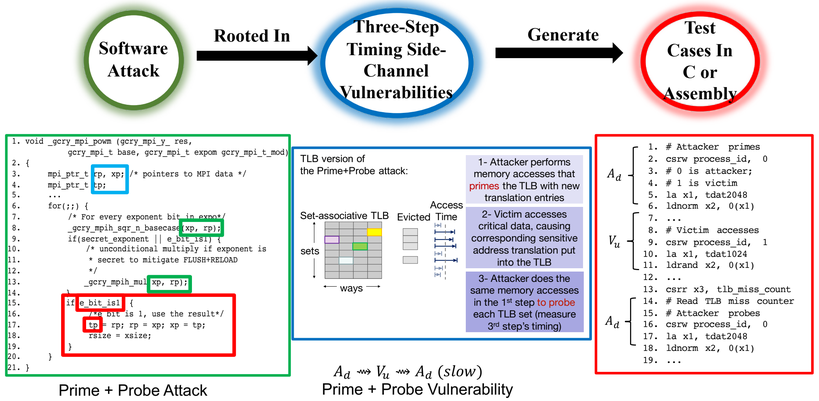

Figure 1. Relationship of attacks, three-step model, and the benchmark

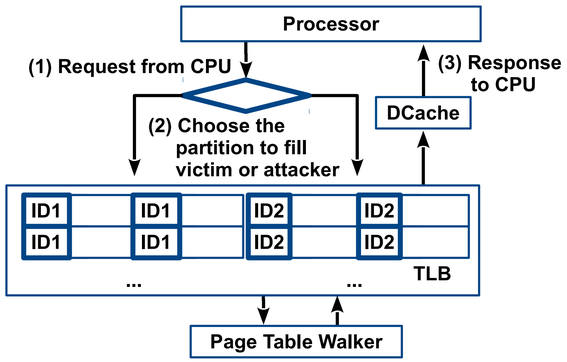

Figure 2. Sample block diagram of the SP TLB with victim (ID1) and attacker (ID2) ways being allocated 50% of TLB space.

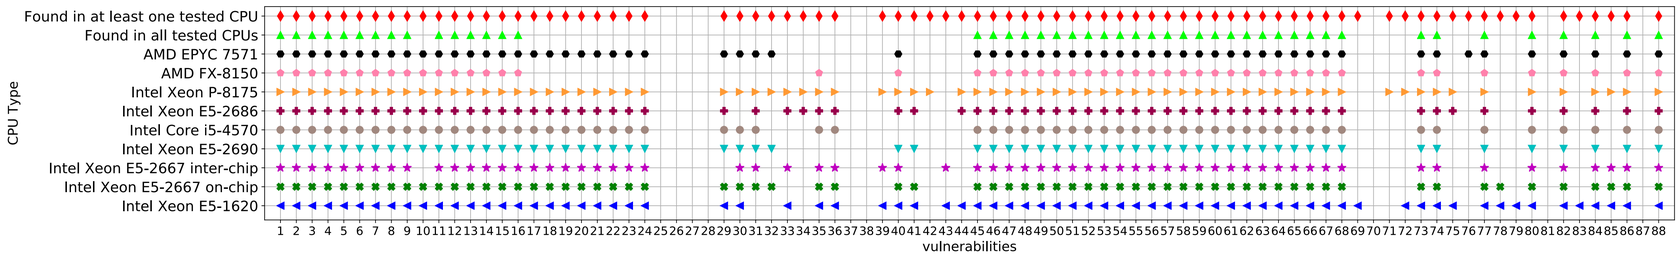

Figure 3. Evaluation of 88 Strong types of cache vulnerabilities on different machines.

Reference:

1. **Deng, Shuwen**, Wenjie Xiong, and Jakub Szefer. "Secure tlbs." 2019 ACM/IEEE 46th Annual International Symposium on Computer Architecture (ISCA). IEEE, 2019.
2. **Deng, Shuwen**, Wenjie Xiong, and Jakub Szefer. "A Benchmark Suite for Evaluating Caches' Vulnerability to Timing Attacks." Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems. 2020.