INTRODUCTION

Microarchitectural attacks have plunged Computer Architecture into a security crisis. Yet, as the slowing of Moore's law justifies the use of ever more exotic microarchitecture, it is likely we have only seen the tip of the iceberg.

- Uncovers seven classes of microarchitectural **optimizations with novel security implications**1. DATA MEMORY-DEPENDENT4. COMPUTATION REUSE
 - Prefetchers
- 5. SILENT STORES
- 2. Computation Simplification 6. Value Prediction 3. PIPELINE COMPRESSION
 - 7. Register-File Compression
- Proposes a conceptual framework through which to study microarchitectural optimizations
- Demonstrates **several Proofs-of-Concept** to show their efficacy

This paper's goal is to perform an early analysis to inform secure and performant development moving forward.

MOTIVATING EXAMPLE: DATA MEMORY DEPENDENT PREFETCHERS

Data Memory-Dependent Prefetchers (DMPs) are **effective** in cases where stride prefetchers fail, e.g., in applications dominated by indirections or "pointer chasing."

for (i=0...N) Multi-Level Indirection access pattern. The DMP prefetches entries in Y by using the *values* **X[Y[Z[i]]]** of **Z**, and entries in **X** by using the *values* of **Y**

DMPs in an Adversarial Setting Attacker Goal: read memory outside of a software sandbox.



1) Attacker activates the prefetcher, 3) The DMP reads any **attacker** with the now-malicious program for(i=0...N): X[Y[Z[i]]]

chosen private value **y=Y[z]**, out of bounds of Y

2) Attacker tricks the prefetcher 4) Finally leaks that value over a into reading attacker-controlled data $z=Z[i+\Delta]$, out of bounds of Z

traditional cache covert channel vis. the final prefetch for **&X[y]**

AUGURY: FOLLOW UP WORK AT OAKLAND'22 We demonstrate the existence of a pointer- chasing **DMP on recent Apple processors**, including the A14 and M1. We then *reverse engineer the details of this DMP* to determine the opportunities for and restrictions it places on attackers using it. Finally, we demonstrate several **basic attack primitives** capable of leaking pointer values using the DMP.

Read more about this follow up work at: **prefetchers.info**

Opening Pandora's Box: A Systematic Study of New Ways Microarchitecture Can Leak Private Data

Jose Rodrigo Sanchez Vicarte, **Pradyumna Shome**, Nandeeka Nayak, Caroline Trippel*, Adam Morrison_o, David Kohlbrenner•, Christopher W. Fletcher University of Illinois at Urbana-Champaign, *Stanford University, °Tel Aviv University, •University of Washington



Data at rest Register File ____ Data Memory

The paper develops a novel conceptual framework and abstraction, called Microarchitectural Leakage Descriptors (MLDs), that characterizes privacy leakage precisely. An MLD is a map that indicates when a microarchitectural optimization changes a program's **observable execution** (e.g., based on timing or hardware resource usage) as a function of what changes to what program data.

EXAMPLE MLD: SILENT STORES The silent stores optimization **skips stores which do not** change the value already in memory, reducing pressure on the memory system.

Silent stores can be thought of as having two inputs: The **in-flight store data** (in use) and the **data already in memory** (at rest). *Depending on these values*, **silent** stores maps to two, distinct, observable outcomes. • Active attack: If the attacker controls either input, they can **perform a search over multiple replays** to leak the other operand.

It's non-trivial to *convert a sin*gle silent store into a timing *difference*. We develop an **am**plification gadget. Attacker induces a silent store, conditioned on the **data left behind** in memory by the victim's prior encryption operation.

We found a range of microarchitectura optimizations with novel security implications

MICROARCHITECTURAL LEAKAGE DESCRIPTORS (MLDs)



- Passive Attack: The program can leak its own data. All stores become potentially latent gadgets.

SILENT STORES POC ON BITSLICE AES128 ENCRYPTION



CONCLUSION

Proactively understanding the security implications of 'exotic' microarchitecture is of fundamental importance for building holistic, long-term, defenses and anticipating future attacks.

This paper **performed a systematic study of the** computer architecture literature through a security lens. We found a range of microarchitectural optimizations with novel security implications—ranging from ones as devastating as Spectre/Meltdown (but without relying on speculative execution) to ones that render constant-time programming ineffective, or in need of overhaul.

Microarchitectural Leakage Descriptors (MLDs) present a starting point towards communicating exactly the information essential for software to implement performant and secure defenses, while hiding second-order details about the hardware.

ACKNOWLEDGEMENTS

This work was partially funded by an Intel ISRA grant, NSF grants #1816226 and #1942888, and ISF grant #2005/17.



EXAMPLE MLD: COMPUTATION REUSE

Minor, performance-guided, decisions can result in significantly different security implications.

V1. One variant of *Computation Reuse* skips computations (forwards the result) by using operand *values* as memoization table keys. This *leaks data values!*

V2. A second variant uses *operand register IDs*. This only leaks control flow.

> The Grainger College of Engineering UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN





TEL AVIV UNIVERSIT